Oracle® Cloud

Using Oracle Autonomous Database Serverless

G31644-13
June 2025

Oracle Cloud Using Oracle Autonomous Database Serverless,

G31644-13

Copyright © 2024, 2025, Oracle and/or its affiliates.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle®, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface

Audience xxx
Documentation Accessibility xxx
Related Documents xxx
Conventions xxxi

1 Overview

What's New 1-1
- June 2025 1-2
- May 2025 1-3
- April 2025 1-6
- March 2025 1-7
- February 2025 1-9
- January 2025 1-9

About Autonomous Database 1-10
- What is Oracle Autonomous Database? 1-10
- Key features 1-11
- About Autonomous Database Workload Types 1-14
  - About Oracle Autonomous Data Warehouse 1-15
  - About Oracle Autonomous Transaction Processing 1-15
  - About Autonomous JSON Database 1-16
  - About Oracle APEX Application Development 1-17
- Autonomous Database Region Availability 1-17

Get Help and Contact Support 1-17
- Post Questions on Forums 1-18
- Open a Support Ticket to Create a Service Request 1-18
- Create a Service Request in My Oracle Cloud Support 1-19
- Register to My Oracle Cloud Support as a New User 1-20

Service Level Objectives (SLOs) 1-21
- RTO and RPO Objectives 1-22
- Built-In Tool Availability 1-22
- Zero-Regression Service Level Objective 1-23

How Is Autonomous Database Billed? 1-24

Autonomous Database Billing Summary 1-24
Compute Models in Autonomous Database 1-26
ECPU Compute Model Billing Information 1-27
OCPU Compute Model Billing Information 1-34
Oracle Autonomous Database Serverless Features Billing 1-39
Oracle Autonomous Database Serverless Billing for Database Tools 1-45
Autonomous Database for Developers Billing and Tenancy Service Limit 1-47
Always Free Autonomous Database 1-48
Upgrade an Always Free Autonomous Database to a Paid Instance 1-51
Autonomous Database for Developers 1-52
Resource Restrictions with Autonomous Database for Developers 1-53
Oracle Database Versions with Autonomous Database for Developers 1-53
Supported Features with Autonomous Database for Developers 1-53
Provisioning and Cloning Workflow with Autonomous Database for Developers 1-54
Upgrade Autonomous Database for Developers Instance to a Full Paid Instance 1-54
Features Not Available in Autonomous Database for Developers 1-55

2 Quick Start

Before You Begin 2-1
Autonomous Database 15 Minute Quick Start 2-2

3 Tasks

Provision an Autonomous Database 3-2
Connect to Autonomous Database 3-2
About Connecting to an Autonomous Database Instance 3-3
Secure Connections to Autonomous Database with mTLS or with TLS 3-4
Connect to Autonomous Database Through a Firewall 3-7
Using Application Continuity 3-7
Connect to Autonomous Database Using a Client Application 3-8
About Connecting to Autonomous Database Using a Client Application 3-9
Prepare for Oracle Call Interface, ODBC, and JDBC OCI Connections with Wallets (mTLS) 3-9
Prepare for Oracle Call Interface, ODBC, and JDBC OCI Connections Using TLS Authentication 3-12
Prepare for JDBC Thin Connections 3-13
Connect Microsoft .NET, Visual Studio Code, and Visual Studio with a Wallet (mTLS) 3-14
Connect Microsoft .NET, Visual Studio Code, and Visual Studio Without a Wallet 3-15
Connect Node.js and other Scripting Languages (mTLS) 3-16
Connect Node.js, and Other Scripting Languages Without a Wallet 3-19
Connect Power BI and Microsoft Data Tools to Autonomous Database 3-21
Connect Python Applications to Autonomous Database 3-22

Connect Go Applications to Autonomous Database 3-32
Download Database Connection Information 3-38
Download Client Credentials (Wallets) 3-39
Wallet README File 3-41
View TNS Names and Connection Strings for an Autonomous Database Instance 3-43
Connect to Autonomous Database Using Oracle Database Tools 3-44
Connect Oracle SQL Developer with a Wallet (mTLS) 3-44
Connect Oracle SQL Developer Without a Wallet 3-46
Connect SQL*Plus with a Wallet (mTLS) 3-48
Connect SQL*Plus Without a Wallet 3-49
Connect Oracle SQLcl Cloud with a Wallet (mTLS) 3-50
Connect Oracle SQLcl Cloud Without a Wallet 3-52
Connect with Built-In Oracle Database Actions 3-54
About Database Actions (SQL Developer Web) 3-54
Access Database Actions as ADMIN 3-55
Provide Database Actions Access to Database Users 3-56
Required Roles to Access Tools from Database Actions 3-57
IAM Policies to Access Database Actions and Tools from the OCI Console 3-57
Access Database Actions, Oracle APEX, Oracle REST Data Services, and Developer Tools Using a Vanity URL 3-58
Connect with JDBC Thin Driver 3-58
JDBC Thin Connections with a Wallet (mTLS) 3-59
JDBC Thin Connections Without a Wallet (TLS) 3-67
Preparing for Oracle Call Interface Connections 3-68
Oracle Call Interface (OCI) Connections and Wallets (mTLS) 3-69
Oracle Call Interface (OCI) Connections with TLS Authentication 3-69
Predefined Database Service Names for Autonomous Database 3-69
Connect with Oracle Analytics Desktop 3-70
Connect with Oracle Analytics Cloud 3-70
Connection and Networking Options and Features 3-70
Using ACLs, VCNs, and Private Endpoints with Autonomous Database 3-71
Connect with Oracle Cloud Infrastructure FastConnect 3-74
Access Autonomous Database with VCN Transit Routing 3-74
Access Autonomous Database with Service Gateway 3-74
Use Database Resident Connection Pooling with Autonomous Database 3-75
Use TCP Fast Open (TFO) to Connect Clients to Autonomous Database 3-75
Best Practices for Low Latency Connections with Autonomous Database 3-79
View Network Information on the OCI Console 3-86
Manage Users 3-88
Create Users on Autonomous Database 3-89
Create Users on Autonomous Database with Database Actions 3-89
Create Users on Autonomous Database - Connecting with a Client Tool 3-91

Unlock User Accounts on Autonomous Database 3-92
About User Passwords on Autonomous Database 3-92
Manage the Administrator Account on Autonomous Database 3-93
Remove Users 3-95
Manage User Profiles with Autonomous Database 3-95
Manage Password Complexity on Autonomous Database 3-97
Gradual Database Password Rollover for Applications 3-98
Manage User Roles and Privileges 3-99
Manage Users and User Roles on Autonomous Database - Connecting with Database Actions 3-99
Manage User Privileges on Autonomous Database - Connecting with a Client Tool 3-100
Create and Update User Accounts for Oracle Machine Learning 3-101
Create User 3-102
Add Existing Database User Account to Oracle Machine Learning Components 3-103
Use Identity and Access Management (IAM) Authentication with Autonomous Database 3-105
About Identity and Access Management (IAM) Authentication with Autonomous Database 3-106
Prerequisites for Identity and Access Management (IAM) Authentication on Autonomous Database 3-107
Enable Identity and Access Management (IAM) Authentication on Autonomous Database 3-108
Create Identity and Access Management (IAM) Groups and Policies for IAM Users 3-109
Add IAM Users on Autonomous Database 3-111
Add IAM Roles on Autonomous Database 3-112
Create IAM Database Password for IAM Users 3-113
Connect to Autonomous Database with Identity and Access Management (IAM) Authentication 3-113
Configure IAM Proxy Authentication 3-119
Disable Identity and Access Management (IAM) Authentication on Autonomous Database 3-121
Accessing a Database Cross-Tenancy Using an IAM Integration 3-121
Notes for Using Autonomous Database Tools with Identity and Access Management (IAM) Authentication 3-126
Use Microsoft Entra ID with Autonomous Database 3-127
About Integrating Oracle Autonomous Database with Microsoft Entra ID 3-127
Enable Microsoft Entra ID Authentication on Autonomous Database 3-132
Map Oracle Database Schemas and Roles for Azure AD 3-141
Microsoft Entra ID Client Configuration and Access for Autonomous Database 3-143
Testing the Accessibility of the Entra ID Endpoint 3-143
Use Microsoft Active Directory with Autonomous Database 3-145
Prerequisites to Configure CMU with Microsoft Active Directory on Autonomous Database 3-146
Configure CMU with Microsoft Active Directory on Autonomous Database 3-147
Kerberos Authentication for CMU with Microsoft Active Directory 3-149

Add Microsoft Active Directory Roles on Autonomous Database 3-153
Add Microsoft Active Directory Users on Autonomous Database 3-154
Tools Restrictions with Active Directory on Autonomous Database 3-155
Connect to Autonomous Database with Active Directory User Credentials 3-155
Verify Active Directory User Connection Information with Autonomous Database 3-156
Remove Active Directory Users and Roles on Autonomous Database 3-157
Disable Active Directory Access on Autonomous Database 3-157
Configure Kerberos Authentication with Autonomous Database 3-158
About Kerberos Authentication 3-158
Components of the Kerberos Authentication System 3-159
Enable Kerberos Authentication on Autonomous Database 3-160
Disable Kerberos Authentication on Autonomous Database 3-162
Notes for Kerberos Authentication on Autonomous Database 3-162
Create Credentials or Configure Policies and Roles 3-162
Manage Credentials 3-163
Create Credentials to Access Cloud Services 3-163
Create Credentials to Access Cloud Services in Cloud Shell 3-164
List Credentials 3-170
Delete Credentials 3-170
Use Vault Secret Credentials 3-170
Use Vault Secret Credentials with Oracle Cloud Infrastructure Vault 3-171
Use Vault Secret Credential with Azure Key Vault 3-174
Use Vault Secret Credential with AWS Secrets Manager 3-176
Use Vault Secret Credential with GCP Secret Manager 3-178
Refresh Vault Secret Credentials 3-180
Update Vault Secret Credentials 3-181
Configure Policies and Roles to Access Resources 3-181
Use Resource Principal to Access Oracle Cloud Infrastructure Resources 3-182
Use Amazon Resource Names (ARNs) to Access AWS Resources 3-188
Use Azure Service Principal to Access Azure Resources 3-196
Use Google Service Account to Access Google Cloud Platform Resources 3-201
Load Data 3-210
About Data Loading 3-211
Load Data with Oracle Database Actions 3-212
Load Data from Files in the Cloud 3-212
Create Credentials and Copy Data into an Existing Table 3-213
Create Credentials and Load Data Pump Dump Files into an Existing Table 3-215
Load JSON on Autonomous Database 3-216
Monitor and Troubleshoot Loads 3-228
Examples for Loading Data into Autonomous Database 3-229
Import Data Using Oracle Data Pump on Autonomous Database 3-241
Export Your Existing Oracle Database to Import into Autonomous Database 3-242

Import Data with Oracle Data Pump Latest Version or an Older Version 3-244
Access Log Files for Data Pump Import 3-252
Oracle Data Pump Import and Table Compression 3-252
Notes for Importing with Oracle Data Pump 3-253
Load Data from Local Files with Oracle Database Actions 3-254
Load Data into Existing Autonomous Database Table with Oracle Database Actions 3-254
Load Data or Query Data from Files in a Directory 3-255
Load Data from Local Files Using SQL*Loader 3-256
Use Data Pipelines for Continuous Load and Export 3-257
About Data Pipelines on Autonomous Database 3-257
Create and Configure Pipelines 3-261
Test Pipelines 3-266
Control Pipelines (Start, Stop, Drop, or Reset a Pipeline) 3-267
Monitor and Troubleshoot Pipelines 3-270
Use Oracle Maintained Pipelines 3-274
The Data Load Page 3-275
The Data Load Page 3-275
Augment Data using Table AI Assist Tool 3-407
The Table AI Assist Tool 3-407
Transform Data Using Data Studio 3-420
The Data Transforms Page 3-421
Data Marketplace 3-582
The Data Marketplace Tool 3-582
View Data Pump Jobs and Import Data with Data Pump 3-595
Link Data 3-595
Link to Object Storage Data 3-596
Query External Data 3-597
Access ORC, Parquet or Avro file types 3-600
Query External Tables with Implicit Partitioning 3-606
Access partitioned object storage data 3-612
Querying External Partitioned Data 3-625
Query Data that Uses Table Hyperlink Source 3-627
Query Hybrid Partitioned Data 3-629
Query External Data Pump Dump Files 3-631
Query Big Data Service Hadoop (HDFS) Data from Autonomous Database 3-632
Integrate Data Lake metadata with Data Catalog 3-632
Query External Data with AWS Glue Data Catalog 3-648
Query Apache Iceberg Tables 3-655
Validate External Data 3-660
Validate External Partitioned Data 3-661
Validate Hybrid Partitioned Data 3-662
View Logs for Data Validation 3-663

Use Database Links 3-664
Create Database Links from Autonomous Database to Another Autonomous Database 3-665
Create Database Links to an Oracle Database that is not an Autonomous Database 3-678
Create Database Links to Non-Oracle Databases 3-690
Create Database Links from Other Databases to Autonomous Database 3-711
Drop Database Links 3-712
Call Web Services 3-713
Submit an HTTP Request to a Public Host 3-714
Submit an HTTP Request to a Private Host 3-714
Submit an HTTP Request to Private Site with a Proxy 3-715
Use Credential Objects to Set HTTP Authentication 3-717
Notes for Submitting HTTP Requests with Oracle APEX or Database Actions 3-718
Move Files 3-719
Create and Manage Directories 3-719
Create Directory in Autonomous Database 3-720
Drop Directory in Autonomous Database 3-721
List Contents of Directory in Autonomous Database 3-722
Access Network File System from Autonomous Database 3-722
Load Data from Directories in Autonomous Database 3-728
Copy Files Between Object Store and a Directory in Autonomous Database 3-730
Bulk Operations for Files in the Cloud 3-730
About Bulk File Operations 3-731
Bulk Copy Files in Cloud Object Storage 3-732
Bulk Move Files Across Cloud Object Storage 3-733
Bulk Download Files from Cloud Object Storage 3-734
Bulk Upload Files to Cloud Object Storage 3-734
Bulk Delete Files from Cloud Object Storage 3-735
Monitor and Troubleshoot Bulk File Loads 3-736
Notes for Bulk File Operations 3-738
Replicate Data 3-738
Replicate Data with Oracle GoldenGate 3-738
Export Data 3-739
Export Data to Object Store as CSV, JSON, or XML using EXPORT_DATA 3-740
Export JSON Data to Cloud Object Storage 3-740
Export Data as CSV to Cloud Object Storage 3-742
Export Data as Parquet to Cloud Object Storage 3-744
Export Data as XML to Cloud Object Storage 3-746
File Naming for Text Output (CSV, JSON, Parquet, or XML) 3-747
Export Data to a Directory Location 3-750
Export Data as CSV to a Directory 3-751
Export Data as JSON a Directory 3-752

Export Data as Parquet to a Directory 3-754
Export Data as XML to a Directory 3-756
Export Data to a Directory as Oracle Data Pump Files 3-757
Export Data with Data Pump to an Autonomous Database Directory 3-759
Use Data Pump to Create a Dump File Set on Autonomous Database 3-760
Move Dump File Set from Autonomous Database to Your Cloud Object Store 3-762
Export Data with Data Pump to Object Store 3-763
Use Oracle Data Pump to Export Data to Object Store Using CREDENTIAL Parameter (Version 19.9 or Later) 3-764
Use Oracle Data Pump to Export Data to Object Store Setting DEFAULT_CREDENTIAL Property 3-767
Export Data to Object Store as Oracle Data Pump Files using EXPORT_DATA 3-769
Download Dump Files, Run Data Pump Import, and Clean Up Object Store 3-771
Access Log Files for Data Pump Export 3-772
Oracle GoldenGate Capture for Oracle Autonomous Database 3-773
Data Sharing 3-773
Use Cloud Links for Read Only Data Access on Autonomous Database 3-773
About Cloud Links on Autonomous Database 3-774
Grant Cloud Links Access for Database Users 3-778
Register a Data Set 3-781
Find Data Sets and Use Cloud Links 3-795
Use Cloud Links Consumer Options 3-796
Monitor and View Cloud Links Information 3-800
Define a Virtual Private Database Policy to Secure a Registered Data Set 3-800
Notes for Cloud Links 3-802
Use Table Hyperlinks for Read Only Data Access on Autonomous Database 3-804
About Table Hyperlinks on Autonomous Database 3-804
Generate, List, Invalidate, and Monitor Table Hyperlinks 3-806
Use Table Hyperlink to Access Data 3-815
Use Table Hyperlinks to Access Data in Table Format 3-818
Use Table Hyperlinks to Create an External Table 3-825
Share Data with Data Studio 3-825
The Data Share Tool 3-826
Develop 3-867
Create Applications with Oracle APEX 3-868
About Oracle APEX 3-869
Access Oracle APEX Administration Services 3-870
Create Oracle APEX Workspaces in Autonomous Database 3-872
Access Oracle APEX App Builder 3-873
Create Oracle APEX Developer Accounts 3-873
Load Data from the Cloud into Oracle APEX 3-874
Use JSON Data with Oracle APEX 3-875

Use Web Services with Oracle APEX 3-876
Send Email from Oracle APEX 3-880
Control Oracle APEX Upgrades 3-883
Access Oracle APEX, Oracle REST Data Services, and Developer Tools Using a Vanity URL 3-887
Oracle APEX Limitations on Autonomous Database 3-887
Use JSON with Autonomous Database 3-889
Work with Simple Oracle Document Access (SODA) in Autonomous Database 3-890
Work with JSON Documents Using SQL and PL/SQL APIs on Autonomous Database 3-892
Load JSON Documents with Autonomous Database 3-892
Import SODA Collection Data Using Oracle Data Pump Version 19.6 or Later 3-892
Develop RESTful Services in Autonomous Database 3-896
About Oracle REST Data Services in Autonomous Database 3-896
Access RESTful Services and SODA for REST 3-897
Develop with Oracle REST Data Services on Autonomous Database 3-898
Use SODA for REST with Autonomous Database 3-898
Access Oracle REST Data Services, Oracle APEX, and Developer Tools Using a Vanity URL 3-904
About Customer Managed Oracle REST Data Services on Autonomous Database 3-904
Use Oracle Database API for MongoDB 3-905
Configure Access for MongoDB and Enable MongoDB 3-906
User Management for MongoDB 3-910
Create a Test Autonomous Database User for MongoDB 3-911
Connect MongoDB Applications to Autonomous Database 3-912
Send Email and Notifications on Autonomous Database 3-917
Send Email on Autonomous Database 3-918
Send Slack Notifications from Autonomous Database 3-927
Send Messages to OCI Notifications Service from Autonomous Database 3-929
Send Microsoft Teams Notifications from Autonomous Database 3-932
User Defined Notification Handler for Scheduler Jobs 3-936
About User Defined Notification Handler for Scheduler Jobs 3-936
Create a Job Notification Handler Procedure 3-936
Register the Job Handler Notification Procedure 3-938
Trigger the Job Handler Notification Procedure 3-939
De-Register the Job Handler Notification Procedure 3-941
Oracle Extensions for IDEs 3-941
Use Oracle Cloud Infrastructure Toolkit for Eclipse 3-942
Use Oracle Developer Tools for Visual Studio 3-942
Use Oracle Developer Tools for VS Code 3-943
Use Oracle Java on Autonomous Database 3-944
Enable Oracle Java 3-944
Check Oracle Java Version 3-944

Load Java classes and JAR Files into Autonomous Database 3-945
Notes for Oracle Java on Autonomous Database 3-945
Test Workloads with Oracle Real Application Testing 3-946
About Oracle Real Application Testing 3-946
Capture-Replay Workloads between Autonomous Databases 3-947
Capture-Replay Workloads between non-Autonomous and Autonomous Databases 3-954
Test Workloads Against an Upcoming Patch 3-956
Manage and Store Files in a Cloud Code Repository with Autonomous Database 3-959
About Cloud Code Repositories with Autonomous Database 3-959
Initialize a Cloud Code Repository 3-960
Create and Manage a Cloud Code Repository 3-961
Create and Manage Branches in a Cloud Code Repository 3-962
Export Schema Objects to the Cloud Code Repository Branch 3-964
Use File Operations with a Cloud Code Repository 3-965
Use SQL Install Operations with a Cloud Code Repository 3-966
Invoke User Defined Functions 3-968
About User Defined Functions 3-968
Steps to Invoke OCI Cloud Functions as SQL Functions 3-969
Steps to Invoke AWS Lambda as SQL Functions 3-974
Steps to Invoke Azure Functions as SQL Functions 3-979
Steps to Invoke Google Cloud Run Functions as SQL Functions 3-982
Invoke External Procedures as SQL Functions 3-985
Make External Calls Using a Customer-Managed Wallet 3-998
About Using a Customer-Managed Wallet with External Calls 3-999
Prerequisites to Use a Customer-Managed Wallet with External Calls 3-1001
Use a Customer-Managed Wallet for External Calls with UTL_HTTP 3-1001
Use a Customer-Managed Wallet with Scheduler Email Notifications 3-1003
Notes for Using a Customer-Managed Wallet with External Calls 3-1005
Use Cloud Tables to Store Logging and Diagnostic Information 3-1007
About Cloud Tables 3-1007
Create Cloud Tables 3-1008
Cloud Table Notes 3-1009
Analyze 3-1011
Build Reports and Dashboards Using Oracle Analytics 3-1011
Create Dashboards, Reports, and Notebooks 3-1011
Create Dashboards and Reports to Analyze and Visualize Your Data 3-1012
Create Notebooks, Workspaces, and Projects with Oracle Machine Learning Notebooks 3-1013
Spatial Analytics with Oracle Spatial 3-1015
About Oracle Spatial with Autonomous Database 3-1015
Oracle Spatial Limitations with Autonomous Database 3-1016
The Data Analysis Tool 3-1017

Searching and obtaining information about Analytic Views 3-1021
Creating Analytic Views 3-1022
Working with Analyses 3-1033
Viewing Analyses 3-1033
Workflow to build Analyses 3-1033
Format Numeric Columns in the Data Analysis Tool 3-1033
Creating Analyses 3-1034
Creating Reports 3-1035
Run Natural Language Query in the Data Analysis Tool 3-1046
Generate Natural Language explanation from SQL Query using AI Explain 3-1051
Vector Search in the Data Analysis Tool 3-1053
Using Calculation Templates 3-1058
Oracle Autonomous Database for Excel 3-1065
Oracle Autonomous Database for Google Sheets 3-1110
The Catalog Tool 3-1146
About the Catalog Page 3-1148
Browse and Search Catalogs 3-1152
Sorting and Filtering Entities 3-1153
Filter Entities 3-1153
Display Entity Search 3-1154
Catalog Settings 3-1155
Manage Catalogs 3-1158
Viewing Entity Details 3-1185
Registering Cloud Links to Access Data 3-1194
Export Data to Cloud 3-1195
Gathering Statistics for Tables 3-1196
Editing Tables 3-1197
The Data Insights Page 3-1205
About Insights 3-1206
Generate Insights and View Reports 3-1207
Use Oracle OLAP with Autonomous Database 3-1208
Enable Oracle OLAP on Autonomous Database 3-1208
Manage the Service 3-1209
Lifecycle Operations 3-1210
Provision an Autonomous Database Instance 3-1211
Start Autonomous Database 3-1216
Stop Autonomous Database 3-1216
Restart Autonomous Database 3-1217
Rename Autonomous Database 3-1217
Update the Display Name for an Autonomous Database Instance 3-1218
Terminate an Autonomous Database Instance 3-1219
Undelete a Terminated Autonomous Database Instance 3-1220

Change Autonomous Database Operation Mode to Read/Write Read-Only or Restricted 3-1222
Schedule Start and Stop Times for an Autonomous Database Instance 3-1224
Concurrent Operations on Autonomous Database 3-1225
Updating Compute and storage limits 3-1227
Add CPU or Storage Resources or Enable Auto Scaling 3-1227
Remove CPU or Storage Resources or Disable Auto Scaling 3-1229
Use Auto Scaling 3-1231
Update Billing Model, Choose Bring Your Own License, or Update Instance to a Paid Instance 3-1236
Choose Bring Your Own License Option While Provisioning or Cloning 3-1237
Choose Bring Your Own License on Autonomous Database (ECPU Compute Model) 3-1238
Choose Bring Your Own License on Autonomous Database (OCPU Compute Model) 3-1241
Update to ECPU Billing Model on Autonomous Database 3-1242
Update Always Free Instance to Paid with Autonomous Database 3-1244
Upgrade Autonomous JSON Database to Autonomous Transaction Processing 3-1245
Upgrade APEX Service to Autonomous Transaction Processing 3-1245
Clone, Move, or Upgrade Autonomous Database 3-1246
About Cloning on Autonomous Database 3-1247
Prerequisites for Cloning 3-1248
Clone an Autonomous Database Instance 3-1249
Clone an Autonomous Database from a Backup 3-1254
Cross Tenancy and Cross-Region Cloning 3-1261
Clone Autonomous Database to Change Workload Type 3-1267
Notes for Cloning Autonomous Database 3-1268
Move an Autonomous Database to a Different Compartment 3-1271
Upgrade an Autonomous Database Instance to Oracle Database 23ai 3-1271
Use Refreshable Clones with Autonomous Database 3-1273
About Refreshable Clones on Autonomous Database 3-1274
Prerequisites for Creating a Refreshable Clone 3-1280
Create and Manage Refreshable Clones 3-1280
View Refreshable Clones for an Autonomous Database Instance 3-1291
Use the API 3-1292
Refreshable Clone Notes 3-1293
Manage Autonomous Database Built-in Tools 3-1294
About Autonomous Database Built-in Tools 3-1295
About Configuring Built-in Tools Compute Resources and Limits 3-1296
View Autonomous Database Built-in Tools Status 3-1296
Configure Autonomous Database Built-in Tools 3-1297
Configure Autonomous Database Built-in Tools when you Provision or Clone an Instance 3-1300
Notes for Autonomous Database Built-in Tools 3-1301
Use Autonomous Database Events 3-1302

About Events Based Notification and Automation on Autonomous Database 3-1302
Critical Events on Autonomous Database 3-1303
Information Events on Autonomous Database 3-1305
Autonomous Database Lifecycle Events 3-1310
Get Notified of Autonomous Database Events 3-1313
Manage Time Zone File Updates on Autonomous Database 3-1315
About Time Zone File Update Options 3-1315
Use AUTO_DST_UPGRADE Time Zone File Option 3-1317
Use AUTO_DST_UPGRADE_EXCL_DATA Time Zone File Option 3-1318
Monitor Autonomous Database Availability 3-1320
Monitor Regional Availability of Autonomous Databases 3-1321
Manage Resource Availability with Compartment Quotas 3-1322
Manage and View Notifications 3-1324
About Work Requests 3-1325
View Patch and Maintenance Window Information, Set the Patch Level 3-1325
About Scheduled Maintenance and Patching 3-1326
View Maintenance Event History 3-1329
View Patch Details 3-1330
Set the Patch Level 3-1331
View Maintenance Status Notifications 3-1332
View and Manage Customer Contacts for Operational Issues and Announcements 3-1333
View Oracle Cloud Infrastructure Operations Actions 3-1335
DBA_OPERATOR_ACCESS View 3-1335
Monitor and Manage Performance 3-1336
Monitor the Performance of Autonomous Database with Database Actions 3-1336
Use Database Actions to Monitor Activity and Utilization 3-1337
Use Database Actions to Monitor Active Session History Analytics and SQL Statements 3-1344
Monitor Autonomous Database with Performance Hub 3-1345
Manage Concurrency and Priorities on Autonomous Database 3-1348
Database Service Names for Autonomous Database 3-1348
Idle Time Limits 3-1355
Service Concurrency 3-1355
Change MEDIUM Service Concurrency Limit (ECPU Compute Model) 3-1359
Change MEDIUM Service Concurrency Limit (OCPU Compute Model) 3-1363
Predefined Job Classes with Oracle Scheduler 3-1367
Manage CPU/IO Shares on Autonomous Database 3-1368
Manage Runaway SQL Statements on Autonomous Database 3-1370
Use Fast Ingest on Autonomous Database 3-1371
Enable Database Management Service to Monitor and Manage Autonomous Database 3-1372
Perform SQL Tracing on Autonomous Database 3-1373
Configure SQL Tracing on Autonomous Database 3-1374

Enable SQL Tracing on Autonomous Database 3-1375
Disable SQL Tracing on Autonomous Database 3-1376
View Trace File Saved to Cloud Object Store on Autonomous Database 3-1377
View Trace Data in SESSION_CLOUD_TRACE View on Autonomous Database 3-1378
Monitor Performance with Autonomous Database Metrics 3-1378
View Metrics for an Autonomous Database Instance 3-1379
View Metrics for Autonomous Databases in a Compartment 3-1381
Autonomous Database Metrics and Dimensions 3-1381
Use Custom Metrics on Autonomous Database 3-1381
Available Metrics: oci_autonomous_database 3-1388
Use Sample Data Sets in Autonomous Database 3-1396
Use the Autonomous Database Free Container Image 3-1397
About the Free Container Image 3-1398
Autonomous Database Free Container Image License 3-1398
Autonomous Database Free Container Image Features 3-1398
Autonomous Database Free Container Image Recommendations and Restrictions 3-1399
Container Registry Locations for Autonomous Database Free Container Image 3-1400
Start Free Container Image 3-1401
Perform Database Operations using adb-cli 3-1403
Connect to an Free Container Image 3-1404
Connect to ORDS, APEX, or Database Actions from an Autonomous Database Container Image 3-1404
Available TNS Aliases for Connecting to an Autonomous Database Container Image 3-1405
Connect to an Autonomous Database Container Image with a Wallet 3-1405
Connect to an Autonomous Database Container Image without a Wallet (TLS) 3-1406
Migrate Data Between Autonomous Database Free Containers 3-1407

4 Tutorials

Load and Analyze Your Data with Autonomous Database 4-1
Self-service Tools for Everyone Using Oracle Autonomous Database 4-1
Manage and Monitor Autonomous Database 4-2
Integrate, Analyze and Act on All Data using Autonomous Database 4-2
Load the Autonomous Database with Data Studio 4-2
QuickStart Java applications with Autonomous Database 4-3

5 Features

Data Lake 5-2
About Data Lake with Autonomous Database 5-2
What is a Data Lake? 5-2
Key Data Lake Features of Autonomous Database 5-3

Ready for Data Lake Workloads 5-4
Multi-Cloud Support 5-4
End-To-End Data Format Support 5-5
Enhanced Capabilities: Autonomous Database for Unstructured Data Management 5-6
Flexible Metadata Management 5-6
Federated Metadata Support 5-7
Collaboration 5-7
Broad Compatibility with Oracle Database Tools 5-8
Data Studio 5-8
High Availability 5-12
Use Application Continuity on Autonomous Database 5-13
About Application Continuity on Autonomous Database 5-13
Configure Application Continuity on Autonomous Database 5-15
Client Configuration for Continuous Availability on Autonomous Database 5-20
Notes for Application Continuity on Autonomous Database 5-25
Use Standby Databases with Autonomous Data Guard for Disaster Recovery 5-25
About Standby Databases 5-26
Enable Autonomous Data Guard 5-40
Update Standby to Use a Backup Copy Peer 5-43
Autonomous Data Guard Switchover and Failover Operations 5-43
Autonomous Data Guard Cross-Region and Cross Tenancy Operations 5-51
Events and Notifications for a Standby Database 5-76
Autonomous Database Cross-Region Paired Regions 5-77
Use the API 5-77
Autonomous Data Guard Notes 5-77
Backup and Restore Autonomous Database Instances 5-80
About Backup and Recovery on Autonomous Database 5-80
Edit Automatic Backup Retention Period on Autonomous Database 5-83
Create Long-Term Backups on Autonomous Database 5-85
View Backup Information and Backups on Autonomous Database 5-88
Restore and Recover your Autonomous Database 5-88
Backup and Restore Notes 5-90
Use OraMTS Feature 5-91
About OraMTS Recovery Service 5-92
Prerequisites to Enable OraMTS Recovery Service on Autonomous Database 5-92
Enable OraMTS Recovery Service on an Autonomous Database 5-93
Disable OraMTS Recovery Service on an Autonomous Database 5-94
Use Backup-Based Disaster Recovery 5-94
About Backup-Based Disaster Recovery 5-95
View Backup-Based Disaster Recovery Peer 5-98
Update Disaster Recovery Type 5-99
Backup-Based Disaster Recovery Switchover and Failover Operations 5-100

Backup-Based Disaster Recovery Cross-Region Operations 5-105
Backup-Based Disaster Recovery Events 5-112
Use the API 5-113
Track Table Changes with Flashback Time Travel 5-113
About Flashback Time Travel 5-114
Enable Flashback Time Travel for a Table 5-114
Disable Flashback Time Travel for a Table 5-115
Modify the Retention Time for Flashback Time Travel 5-115
Purge Historical Data for Flashback Time Travel 5-115
View Flashback Time Travel Information 5-116
Notes for Flashback Time Travel 5-117
Use OCI Full Stack Disaster Recovery with Autonomous Database 5-118
Prerequisite Steps for Autonomous Database with Full Stack Disaster Recovery 5-118
Enable Full Stack Disaster Recovery with Autonomous Database 5-118
Use Full Stack Disaster Recovery to Perform Switchover or Failover Operations 5-119
Notes for Full Stack Disaster Recovery 5-120
Security 5-121
Security Features Overview 5-122
Security and Authentication in Oracle Autonomous Database 5-123
Configuration Management 5-124
Data Encryption 5-124
Data Access Control 5-125
Auditing Overview on Autonomous Database 5-128
Assessing the Security of Your Database and its Data 5-129
Regulatory Compliance Certification 5-130
Security Testing Policies 5-131
Manage Encryption Keys on Autonomous Database 5-131
About Master Encryption Key Management on Autonomous Database 5-132
Manage Master Encryption Keys in OCI Vault 5-133
Manage Master Encryption Keys in Azure Key Vault 5-143
Manage Master Encryption Keys in AWS Key Management Service 5-150
Manage Master Encryption Keys in Oracle Key Vault 5-158
Switch to Oracle-Managed Encryption Keys on Autonomous Database 5-166
View History for Customer-Managed Encryption Keys on Autonomous Database 5-166
Notes for Using Customer-Managed Keys with Autonomous Database 5-167
Configure Network Access with Access Control Rules (ACLs) and Private Endpoints 5-171
Configure Network Access with Access Control Rules 5-171
Configure Network Access with Private Endpoints 5-177
Use Zero Trust Packet Routing on Autonomous Database 5-195
Update Network Options to Allow TLS or Require Only Mutual TLS (mTLS) Authentication on Autonomous Database 5-196
Audit Autonomous Database 5-200

About Auditing Autonomous Database 5-200
Register Oracle Data Safe on Autonomous Database 5-202
Extend Audit Record Retention with Oracle Data Safe on Autonomous Database 5-203
View and Manage Oracle Data Safe Audit Trails on Autonomous Database 5-204
View and Manage Audit Policies with Oracle Data Safe on Autonomous Database 5-204
Generate Audit Reports with Data Safe on Autonomous Database 5-204
Rotate Wallets for Autonomous Database 5-205
About Wallet Rotation 5-205
Rotate Wallets with Immediate Rotation 5-205
Rotate Wallets with Grace Period 5-206
Use Oracle Database Vault with Autonomous Database 5-209
Oracle Database Vault Users and Roles on Autonomous Database 5-209
Enable Oracle Database Vault on Autonomous Database 5-210
Disable Oracle Database Vault on Autonomous Database 5-211
Disable User Management with Oracle Database Vault on Autonomous Database 5-211
Enable User Management with Oracle Database Vault on Autonomous Database 5-212
IAM Policies for Autonomous Database 5-212
Policy Details for Autonomous Database 5-212
IAM Permissions and API Operations for Autonomous Database 5-218
Provide Specific Privileges in IAM Policies to Manage Autonomous Database 5-221
Use Oracle Data Safe 5-223
About Data Safe 5-223
Register Data Safe 5-224
Use Data Safe Features 5-225
Break Glass Access for SaaS on Autonomous Database 5-226
About Break Glass Access on Autonomous Database 5-226
Enable Break Glass Access 5-228
Disable Break Glass Access 5-230
Notes for Break Glass Access 5-231
Encrypt Data While Exporting or Decrypt Data While Importing 5-232
Encrypt Data While Exporting to Object Storage 5-232
Decrypt Data While Importing from Object Storage 5-237
Manage Oracle Cloud Operator Access 5-241
Enable Cloud Operator Access 5-242
Disable Cloud Operator Access 5-243
Performance Monitor and Management 5-243
Use OCI Ops Insights on Autonomous Database 5-243
Manage Optimizer Statistics on Autonomous Database 5-244
Manage Optimizer Statistics and Hints with Data Warehouse Workloads 5-244
Manage Optimizer Statistics and Hints with Transaction Processing and JSON Database Workloads 5-245
Manage Automatic Indexing on Autonomous Database 5-246

Manage Automatic Partitioning on Autonomous Database 5-246
About Automatic Partitioning 5-247
How Automatic Partitioning Works 5-248
Configure Automatic Partitioning 5-248
Use Automatic Partitioning 5-250
Generate Automatic Partitioning Reports 5-251
Example Automatic Partitioning Scenarios 5-252
Data Dictionary Views for Automatic Partitioning 5-257
Graph Analytics using Graph Studio 5-258
Use Oracle Graph Analytics 5-259
Build Applications Using Graph APIs 5-259
Use Oracle Machine Learning with Autonomous Database 5-260
OML Machine Learning 5-260
Creating Analytic Views 5-262
Use Select AI for Natural Language Interaction with your Database 5-273
About Select AI 5-274
Usage Guidelines 5-275
Supported Platforms 5-276
Select your AI Provider and LLMs 5-276
Select AI Concepts 5-284
Actions 5-284
AI Profile 5-285
AI Provider 5-285
Conversations 5-285
Database Credentials 5-285
Hallucination in LLM 5-285
IAM 5-285
Large Language Model (LLM) 5-286
Metadata 5-286
Metadata Clone 5-286
Natural Language Prompts 5-286
Network Access Control List (ACL) 5-286
Retrieval Augmented Generation (RAG) 5-286
Semantic Similarity Search 5-287
Vector Distance 5-287
Vector Index 5-287
Vector Store 5-287
Select AI Use Cases 5-287
Getting Started with Select AI 5-288
Manage AI Profiles 5-289
Use DBMS_CLOUD_AI to Configure AI Profiles 5-289
Perform Prerequisites for Select AI 5-289

Create and Set an AI Profile 5-296
Use AI Keyword to Enter Prompts 5-296
Enable Conversations to Enhance User Interaction 5-298
Select AI with Retrieval Augmented Generation (RAG) 5-299
Build your Vector Store 5-299
Use DBMS_CLOUD_AI to Create and Manage Vector Indexes 5-301
Use In-database Transformer Models 5-301
Benefits of Select AI RAG 5-302
Synthetic Data Generation 5-302
Benefits of Synthetic Data Generation 5-303
Generate Synthetic Data 5-303
Monitor and Troubleshoot Synthetic Data Generation 5-303
Examples of Using Select AI 5-305
Example: Select AI Actions 5-307
Example: Select AI with OCI Generative AI 5-308
Example: Select AI with OpenAI 5-318
Example: Select AI with Cohere 5-321
Example: Select AI with Azure OpenAI Service 5-323
Example: Select AI with Google 5-330
Example: Select AI with Anthropic 5-333
Example: Select AI with Hugging Face 5-337
Example: Enable Conversations in Select AI 5-341
Example: Set Up and Use Select AI with RAG 5-344
Example: Select AI with In-database Transformer Models 5-346
Example: Improve SQL Query Generation 5-353
Example: Generate Synthetic Data 5-358
Example: Enable or Disable Data Access 5-364
Example: Restrict Table Access in AI Profile 5-366
Example: Specify Case Sensitivity for Columns 5-367
JSON Document Stores 5-368
Cache Messages with Singleton Pipes and Using Pipes for Persistent Messaging 5-369
Cache Messages with Singleton Pipes 5-369
About Caching Messages with Singleton Pipes 5-369
Automatic Refresh of Cached Message with a Cache Function 5-372
Create an Explicit Singleton Pipe 5-373
Create an Explicit Singleton Pipe with a Cache Function 5-375
Use Persistent Messaging with Messages Stored in Cloud Object Store 5-376
About Persistent Messaging with DBMS_PIPE 5-377
Create an Explicit Persistent Pipe and Send a Message 5-379
Retrieve a Persistent Message on Same Database 5-381
Retrieve a Persistent Message by Creating a Pipe on a Different Database 5-382
Remove a Persistent Pipe 5-384

Full-Text Search Queries 5-385
Use Full-Text Search on Files in Object Storage 5-385
About Full-Text Search on Files in Object Storage 5-385
Create a Text Index on Object Storage Files 5-386
Drop an Index on the Cloud Storage Files 5-387
Text Index Reference Table 5-387
Monitor Text Index Creation 5-388
Use JSON Search on Documents in Object Storage 5-389
Use JSON Search on Documents in Object Storage 5-389
Create a JSON Search Index on JSON Documents in Files 5-390
Drop a JSON Search Index 5-391
Monitor JSON Search Index Creation 5-391
JSON Search Index Reference Table 5-392
JSON Search Index Reference View 5-393
Error Log Table 5-393
Use Oracle Workspace Manager on Autonomous Database 5-393
About Using Oracle Workspace Manager on Autonomous Database 5-394
Enable Oracle Workspace Manager on Autonomous Database 5-394
Disable Oracle Workspace Manager on Autonomous Database 5-395
Use and Manage Elastic Pools on Autonomous Database 5-395
About Elastic Pools 5-396
About Elastic Pool Billing 5-399
About Elastic Pool Billing with Autonomous Data Guard Enabled 5-407
About Elastic Pool Billing with Autonomous Data Guard Enabled with a Local Standby 5-407
About Elastic Pool Billing with a Cross-Region Standby 5-407
Create, Join or Manage an Elastic Pool 5-408
Create an Elastic Pool 5-408
Create an Elastic Pool While Provisioning or Cloning an Instance 5-410
Join an Existing Elastic Pool 5-411
Join an Elastic Pool While Provisioning or Cloning an Instance 5-412
Change the Elastic Pool Shape 5-414
List Elastic Pool Members 5-414
Offload Queries to Member Refreshable Clones 5-415
About Query Offloading 5-415
Enable Query Offload 5-418
Disable Query Offload 5-421
Query Offload from PL/SQL 5-421
Query Offload from a Scheduler Job 5-423
Notes for Offloading Queries 5-424
Remove Pool Members from an Elastic Pool 5-425
As Pool Leader Remove Members from an Elastic Pool 5-426

Notes for Leaving an Elastic Pool 5-426
Terminate an Elastic Pool 5-427
Elastic Pool Notes 5-427
Migrate to Autonomous Database 5-428
Migrate Oracle Databases to Autonomous Database 5-428
Migration Prerequisites 5-428
Recommended Migration Methods 5-429
Migrate Existing SQL Execution Plans to Autonomous Database to Mitigate the Risk of SQL Performance Regressions 5-431
Migrate Applications from Non-Oracle SQL to Oracle SQL 5-435
Migrate Applications from PostgreSQL Databases to Autonomous Database 5-436
Migrate Applications from MySQL Databases to Autonomous Database 5-444
Migrate Applications from SQL Server Databases to Autonomous Database 5-451
Use Database In-Memory with Autonomous Database 5-459
About Database In-Memory 5-459
Enabling Objects for In-Memory Population 5-460
Sizing the In-Memory Column Store 5-462
Disabling the In-Memory Column Store 5-462

6 Reference

Previous Feature Announcements 6-1
2024 What's New 6-2
December 2024 6-2
November 2024 6-3
October 2024 6-3
September 2024 6-5
August 2024 6-6
July 2024 6-7
June 2024 6-9
May 2024 6-11
April 2024 6-12
March 2024 6-14
February 2024 6-15
January 2024 6-16
2023 What's New 6-16
December 2023 6-17
November 2023 6-17
October 2023 6-18
September 2023 6-18
August 2023 6-19
July 2023 6-20

June 2023 6-21
May 2023 6-22
April 2023 6-22
March 2023 6-24
February 2023 6-25
January 2023 6-26
2022 What's New 6-26
December 2022 6-27
November 2022 6-27
October 2022 6-29
September 2022 6-29
August 2022 6-30
July 2022 6-30
June 2022 6-31
May 2022 6-32
April 2022 6-33
March 2022 6-34
February 2022 6-35
January 2022 6-35
2021 What's New 6-36
December 2021 6-37
November 2021 6-38
October 2021 6-38
September 2021 6-39
August 2021 6-39
July 2021 6-40
June 2021 6-40
May 2021 6-41
April 2021 6-42
March 2021 6-44
February 2021 6-45
January 2021 6-45
2020 What's New 6-45
December 2020 6-46
November 2020 6-46
October 2020 6-47
September 2020 6-47
August 2020 6-48
July 2020 6-48
June 2020 6-49
May 2020 6-49
April 2020 6-50

March 2020 6-51
February 2020 6-51
January 2020 6-52
2019 What's New 6-52
December 2019 6-53
November 2019 6-54
October 2019 6-54
September 2019 6-55
August 2019 6-55
July 2019 6-55
June 2019 6-56
May 2019 6-56
April 2019 6-57
March 2019 6-58
February 2019 6-59
January 2019 6-59
2018 What's New 6-59
October 2018 6-60
September 2018 6-60
August 2018 6-60
July 2018 6-60
June 2018 6-61
May 2018 6-61
April 2018 6-61
Sample Star Schema Benchmark (SSB) Queries and Analytic Views 6-62
Star Schema Benchmark Queries 6-62
Star Schema Benchmark Analytic Views 6-65
Autonomous Database Supplied Package Reference 6-67
DBMS_CLOUD Package 6-69
DBMS_CLOUD Endpoint Management 6-69
DBMS_CLOUD Subprograms and REST APIs 6-70
DBMS_CLOUD URI Formats 6-157
DBMS_CLOUD Package Format Options 6-167
DBMS_CLOUD Package Format Options for EXPORT_DATA 6-180
DBMS_CLOUD Avro, ORC, and Parquet Support 6-185
DBMS_CLOUD Exceptions 6-198
DBMS_CLOUD_ADMIN Package 6-198
Summary of DBMS_CLOUD_ADMIN Subprograms 6-199
DBMS_CLOUD_ADMIN Exceptions 6-243
DBMS_CLOUD_AI Package 6-243
DBMS_CLOUD_AI Overview 6-244
Summary of DBMS_CLOUD_AI Subprograms 6-244

DBMS_CLOUD_FUNCTION Package 6-276
Summary of DBMS_CLOUD_FUNCTION Subprograms 6-276
DBMS_CLOUD_LINK Package 6-286
DBMS_CLOUD_LINK Overview 6-287
Summary of DBMS_CLOUD_LINK Subprograms 6-287
DBMS_CLOUD_LINK_ADMIN Package 6-300
DBMS_CLOUD_LINK_ADMIN Overview 6-300
Summary of DBMS_CLOUD_LINK_ADMIN Subprograms 6-300
DBMS_CLOUD_MACADM Package 6-306
CONFIGURE_DATABASE_VAULT Procedure 6-306
DISABLE_DATABASE_VAULT Procedure 6-307
DISABLE_USERMGMT_DATABASE_VAULT Procedure 6-307
ENABLE_DATABASE_VAULT Procedure 6-308
ENABLE_USERMGMT_DATABASE_VAULT Procedure 6-308
DBMS_CLOUD_MIGRATION Package 6-309
Summary of DBMS_CLOUD_MIGRATION Subprograms 6-309
DBMS_CLOUD_NOTIFICATION Package 6-313
DBMS_CLOUD_NOTIFICATION Overview 6-313
Summary of DBMS_CLOUD_NOTIFICATION Subprograms 6-313
DBMS_CLOUD_PIPELINE Package 6-321
Summary of DBMS_CLOUD_PIPELINE Subprograms 6-322
DBMS_CLOUD_PIPELINE Attributes 6-327
DBMS_CLOUD_PIPELINE Views 6-329
DBMS_CLOUD_REPO Package 6-330
DBMS_CLOUD_REPO Overview 6-330
DBMS_CLOUD_REPO Data Structures 6-331
DBMS_CLOUD_REPO Subprogram Groups 6-331
Summary of DBMS_CLOUD_REPO Subprograms 6-333
DBMS_AUTO_PARTITION Package 6-353
CONFIGURE Procedure 6-353
VALIDATE_CANDIDATE_TABLE Function 6-356
RECOMMEND_PARTITION_METHOD Function 6-357
APPLY_RECOMMENDATION Procedure 6-359
REPORT_ACTIVITY Function 6-359
REPORT_LAST_ACTIVITY Function 6-360
DBMS_DATA_ACCESS Package 6-361
DBMS_DATA_ACCESS Overview 6-361
DBMS_DATA_ACCESS Security Model 6-361
Summary of DBMS_DATA_ACCESS Subprograms 6-362
DBMS_DCAT Package 6-372
Data Catalog Users and Roles 6-372
Required Credentials and IAM Policies 6-373

Summary of Connection Management Subprograms 6-376
Summary of Synchronization Subprograms 6-379
Summary of Data Catalog Views 6-387
DBMS_INMEMORY_ADMIN Package 6-401
Summary of DBMS_INMEMORY_ADMIN Subprograms 6-401
DBMS_MAX_STRING_SIZE Package 6-402
CHECK_MAX_STRING_SIZE Function 6-403
MODIFY_MAX_STRING_SIZE Procedure 6-403
DBMS_PIPE Package 6-404
DBMS_PIPE Overview for Singleton Pipes 6-405
Summary of DBMS_PIPE Subprograms for Singleton Pipes 6-405
DBMS_PIPE Overview for Persistent Messaging Pipes 6-413
Summary of DBMS_PIPE Subprograms for Persistent Messaging 6-414
DBMS_PROXY_SQL 6-423
DISABLE_READ_ONLY_OFFLOAD Procedure 6-423
ENABLE_READ_ONLY_OFFLOAD Procedure 6-424
DBMS_SHARE Package 6-425
Summary of Share Producer Subprograms 6-426
Summary of Share Consumer Subprograms 6-464
Summary of Share Producer Views 6-485
Summary of Share Consumer Views 6-502
DBMS_SHARE Constants 6-506
CS_RESOURCE_MANAGER Package 6-514
LIST_CURRENT_RULES Function 6-514
LIST_DEFAULT_RULES Function 6-515
REVERT_TO_DEFAULT_VALUES Procedure 6-515
UPDATE_PLAN_DIRECTIVE Procedure 6-516
CS_SESSION Package 6-518
SWITCH_SERVICE Procedure 6-518
Autonomous Database for Experienced Oracle Database Users 6-520
About Autonomous Database for Experienced Oracle Database Users 6-520
Data Warehouse Workload with Autonomous Database 6-521
Transaction Processing and JSON Database Workloads with Autonomous Database 6-523
Autonomous Database Views 6-524
Track Table and Partition Scan Access with Autonomous Database Views 6-525
Track Oracle Cloud Infrastructure Resources, Cost and Usage Reports with Autonomous Database Views 6-526
Pipeline Views 6-536
DBMS_CLOUD_AI Views 6-540
DBMS_CLOUD_FUNCTION Views 6-544
DBMS_CLOUD_LINK Views 6-545
DBMS_DATA_ACCESS Views 6-548

DBA_PROXY_SQL Views 6-549
Lockdown Errors View 6-549
Oracle Cloud Infrastructure Logging Interface Views 6-550
Real Application Testing Capture Replay Views 6-557
Stripe Views on Autonomous Database 6-558
Track DBMS_CLOUD Load Operations 6-571
Autonomous Database – Oracle Database Features 6-572
Autonomous Database with Oracle Database 23ai 6-573
Autonomous Database with Oracle Database 23ai Features 6-574
Autonomous Database with Oracle Database 23ai Notes 6-574
Autonomous Database RMAN Recovery Catalog 6-574
Use Autonomous Database as an RMAN Recovery Catalog 6-574
Notes for Users Migrating from Other Oracle Databases 6-575
Initialization Parameters 6-577
SQL Commands 6-587
Data Types 6-590
PL/SQL Package Notes for Autonomous Database 6-591
ADMIN User Roles and Privileges 6-596
Oracle XML DB 6-597
Oracle Text 6-598
Oracle Flashback 6-599
Oracle Database Real Application Security 6-599
Oracle LogMiner 6-600
Choose a Character Set for Autonomous Database 6-601
Query the Alert Log 6-602
Database Features Unavailable in Autonomous Database 6-603
Logical Partition Change Tracking and Materialized Views 6-603
About Logical Partition Change Tracking 6-603
Using Logical Partition Change Tracking 6-604
Example: Logical Partition Change Tracking 6-607
Migrating MySQL and Third-Party Databases to Autonomous Database 6-612
Migrating Amazon Redshift to Autonomous Database 6-613
Autonomous Database Redshift Migration Overview 6-613
Connect to Amazon Redshift 6-614
Connect to Autonomous Database 6-616
Start the Cloud Migration Wizard 6-617
Review and Finish the Amazon Redshift Migration 6-623
Use Generated Amazon Redshift Migration Scripts 6-624
Perform Post Migration Tasks 6-626
Migrating MySQL to Autonomous Database 6-626
Migrating Microsoft SQL Server or Sybase Adaptive Server to Autonomous Database 6-627
Migrating IBM DB2 (UDB) Database to Autonomous Database 6-627

Migrating Teradata Database to Autonomous Database 6-627
SODA Collection Metadata 6-627
SODA Collection Metadata on Autonomous Database 6-627
SODA Default Collection Metadata on Autonomous Database 6-628
SODA Customized Collection Metadata on Autonomous Database 6-629
Cloud Support and Identity Information 6-630
Obtain Tenancy Details 6-631
Database Name 6-633
Region 6-633
Tenancy OCID 6-633
Database OCID 6-633
Compartment OCID 6-633
Outbound IP Address 6-633
Public Domain Name 6-634
Tenant Account Name 6-634
Autoscalable Storage 6-634
Base Size 6-634
Infrastructure 6-634
Service 6-634
Applications 6-635
Availability Domain 6-635
Compute Model 6-635
Compute Count 6-635
Compute Auto Scaling 6-635
Cloud Identity Information for Autonomous Data Guard 6-635
Public IP Address Ranges for Autonomous Database 6-635

Preface

This document describes how to manage, monitor, and use Oracle Autonomous Database and provides references to related documentation.

- Audience
- Documentation Accessibility
- Related Documents
- Conventions

Audience

This document is intended for Oracle Cloud users who want to manage and monitor Oracle Autonomous Database.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

Oracle customer access to and use of Oracle support services will be pursuant to the terms and conditions specified in their Oracle order for the applicable services.

Related Documents

Depending on the region and when you provisioned your database, and in some cases depending on your provisioning choice, the Oracle Database version for your Autonomous Database is one of: Oracle Database 19c or Oracle Database 23ai.

If you have Oracle Database 19c, then many database concepts and features of this service are further documented here:

Oracle Database 19c

Oracle Database 23ai

If you are using Always Free Autonomous Database with Oracle Database 23ai, then many concepts and features of this service are further documented here:

Oracle Database 23ai

For additional information, see these Oracle resources:

- Welcome to Oracle Cloud Infrastructure

- *Oracle Cloud Infrastructure Object Storage*
- Get Started Using Autonomous JSON Database
- GoldenGate Real-Time Data Replication in Cloud
- *Getting Started with Oracle Analytics Cloud*
- *User's Guide for Oracle Analytics Desktop*

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

1

Overview

- What's New for Oracle Autonomous Database Serverless
  Here's a summary of the noteworthy Oracle Autonomous Database additions and enhancements.
- About Autonomous Database
  Provides information about Autonomous Database features, workload types, and regions.
- Get Help, Search Forums, and Contact Support
  When you use Oracle Autonomous Database, sometimes you need to get help from the community or to talk to someone in Oracle support. This topic provides more information about getting help by viewing and posting questions on forums and using Oracle Cloud Support to create a support request.
- Service Level Objectives (SLOs)
  Describes the Service Level Objectives (SLOs) for Oracle Autonomous Database Serverless.
- How Is Autonomous Database Billed?
  Provides details for Autonomous Database billing.
- Always Free Autonomous Database
  You have the option to create a limited number of Always Free Autonomous Databases that do not consume cloud credits. Always Free databases can be created in Oracle Cloud Infrastructure accounts that are in a trial period, have paying status, or are always free.
- Autonomous Database for Developers
  Autonomous Database for Developers provides you with low cost instances for developers and others to build and test new Autonomous Database applications.

What's New for Oracle Autonomous Database Serverless

Here's a summary of the noteworthy Oracle Autonomous Database additions and enhancements.

See Previous Feature Announcements for 2024 announcements and older announcements.

- June 2025
- May 2025
- April 2025
- March 2025
- February 2025
- January 2025

June 2025

Feature	Description
Table Hyperlink Network Access Checking with ACLs	You can generate a Table Hyperlink with ACL checking enabled. When you generate a Table Hyperlink and you enable ACL checking, Table Hyperlink data is only accessible to a consumer database after the incoming request's IP address is validated based on the ACLs defined on the producer database. See Generate, List, Invalidate, and Monitor Table Hyperlinks for more information.
Send Notifications to an OCI Notifications Service Topic	You can send messages from Autonomous Database to an OCI Notifications service topic. See Send Messages to OCI Notifications Service from Autonomous Database for more information.
Critical and Information Events with Event Category	Critical events and Information events include an event category attribute. Optionally you can subscribe to events and add an event category to receive notifications about events in a category, or to perform other types of automation for events by event type and event category. The event categories are: Incident, Maintenance, Performance, Security, and Other. See Use Autonomous Database Events for more information.
Data Studio Enhancements:	The Data Studio is updated with new features and changes: • **Use SQL Syntax to Query Remote Catalogs**: You can access external data by mounting catalogs and querying them using a simplified SQL syntax. This capability enables you to discover and analyze external data effortlessly. See Query Remote Catalogs and Databases for more information. • **Use Table Hyperlinks to Access Data using Spreadsheet Add-ins**: You can read and access data with a Table Hyperlink in a MS Excel Sheet using Oracle Autonomous Database Add-in for Excel. See Read and Access Data Using Table Hyperlinks in Excel Sheet for more information.
Select AI Enhancements	Select AI includes the following new features: • Select AI enables you to add table comments to the metadata of the LLM for a better SQL generation. You can achieve this by specifying `comments` to *true* in your AI profile. See "*Example: Improve SQL Generation with Table and Column Comments*" in Example: Improve SQL Query Generation. • Select AI integrates annotations in Oracle Database 23ai by adding it to the metadata that is sent to the LLM. You can achieve this by specifying `annotations` to *true* in your AI profile. See "*Example: Improve SQL Query Generation with Table and Column Annotations*" in Example: Improve SQL Query Generation • Select AI adds foreign key and referential key constraints into the metadata of the LLM to generate accurate `JOIN` conditions. You can achieve this by specifying `constraints` to *true* in your AI profile. See "*Example: Improve SQL Query Generation with Foreign Key and Referential Key Constraints*" in Example: Improve SQL Query Generation.

Feature	Description
Cross-region Standby can be an Elastic Pool Member or an Elastic Pool Leader	A cross-region standby can be added to an elastic pool or can be an elastic pool leader. Elastic pools do not operate across regions; when a cross-region standby is added to an elastic pool, the primary and the standby cannot be in the same elastic pool. See Use and Manage Elastic Pools on Autonomous Database for more information.
Google Drive Support for Database Links with Oracle-Managed Heterogeneous Connectivity	Use Autonomous Database Oracle-managed heterogeneous connectivity to create database links to Google Drive. When you use database links with Oracle-managed heterogeneous connectivity, Autonomous Database configures and sets up the connection. See Create Database Links to Non-Oracle Databases with Oracle-Managed Heterogeneous Connectivity for more information.
Built-in Tools Billing and Default Configuration Changes	In the ECPU compute model, the built-in database tools that use a VM, including: OML, Graph Studio, and Data Transforms have default configuration changes and billing changes. See About Configuring Built-in Tools Compute Resources and Limits and Oracle Autonomous Database Serverless Billing for Database Tools for more information.
Elastic Pool Cost and Usage Information in Cost Reports	Elastic pool cost and usage details are recorded in the Oracle Cloud Infrastructure cost reports and in the `OCI_COST_DATA` view. See About Elastic Pool Billing and Compute Models in Autonomous Database for more information.
Offload Queries from Elastic Pool Member to Member Refreshable Clones: Includes Queries from PL/SQL and from Oracle Scheduler Jobs	When you have heavy read workloads, where a relatively small amount of data is scanned by multiple queries, including from PL/SQL or from Oracle Scheduler Jobs, you can offload queries (reads) from either an elastic pool leader or from an elastic pool member to Refreshable Clones. See Offload Queries to Member Refreshable Clones for more information.

May 2025

Feature	Description
Remove OCPU Compute Model Option for Data Warehouse and Transaction Processing Workloads During Provisioning or Cloning	The OCPU compute model option is removed when you provision or clone a new Autonomous Database instance with Data Warehouse or Transaction Processing workloads. See Oracle Support Document 2998742.1 for more information.
Support Zstd Compression in `DBMS_CLOUD`	`DBMS_CLOUD` supports loading and querying data using external tables as well as exporting and importing of data from Autonomous Database to object storage. `DBMS_CLOUD` supports Zstd (Zstandard) format for compression or decompression. See DBMS_CLOUD Package Format Options and DBMS_CLOUD Package Format Options for EXPORT_DATA for more information.
Specify Default Bind Variables in Table Hyperlinks	You can specify default values for one or more bind variables in a SQL statement when the `default_bind_values` parameter is included with `DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL`. See Generate a Table Hyperlink with a Select Statement and Use Table Hyperlink to Access Data for more information.

Feature	Description
Select AI Enhancements	Select AI includes the following new features: • Select AI automatically detects relevant tables and sends metadata only for those specific tables, relevant to the query, in Oracle Database 23ai. You can achieve this by specifying `object_list_mode` as *automated*. See "*Automatically Detect Relevant Table Metadata*" in Example: Improve SQL Query Generation. • Select AI enables you to restrict table access by considering only the tables listed in the AI profile attributes for SQL generation. You can achieve this by specifying `enforce_object_list` to *true*. See Example: Restrict Table Access in AI Profile. • Select AI enables you to specify case sensitivity such that the LLM produces case-insensitive responses from the database and LLM. You can achieve this by specifying `case_sensitive_values` to *false*. See Example: Specify Case Sensitivity for Columns. • Administrators can control data access such that actual schema tables are not sent to the LLM using Select AI. See ENABLE_DATA_ACCESS Procedure, DISABLE_DATA_ACCESS Procedure, and Example: Enable or Disable Data Access.
Use In-database Transformer Model with Select AI RAG	You can use in-database transformers (also known as embedding models) that you imported into your Oracle Database 23ai instance with Select AI RAG to generate vector embeddings for document chunks and user prompts. See Use In-database Transformer Models and Example: Select AI with In-database Transformer Models to explore this feature.
Data Studio Enhancement: Use Table Hyperlinks to Access Data using Spreadsheet Add-ins	You can read and access data with a Table Hyperlink in a MS Excel Sheet using Oracle Autonomous Database Add-in for Excel. See Read and Access Data Using Table Hyperlinks in Excel Sheet.

Feature	Description
Data Transforms Updates	Data Transforms is updated with new features and changes in the following areas: • Run Data Studio Data Load using Data Transforms workflows You can now define a Data Studio Data Load in a Data Transforms workflow and then execute that data load as part of a Data Flow. You can then schedule the workflow to replay the Data Studio Data Load to capture any updates made to the files and load it into Autonomous Database. See Define a Data Studio Data Load in a Work Flow for more information about how you can run a Data Studio Data Load in Data Transforms. • Vector Embeddings in Data Flows Data Transforms now supports the integration of vector embeddings in a data flow using the Oracle Cloud Infrastructure (OCI) Generative AI service. Note that you need to have an Oracle Database 23ai connection to use this feature. For more information, see Use Text Embedding Vector in a Data Flow. • Use Machine Learning (ML) models in Data Flows Data Transforms now includes a new data entity type called ML Model. You can use ML Models in a data flow to apply prediction logic on source data and store the output in a target server. See Machine Learning (ML) Models for more information. • Set Custom Variable Value at Runtime For data flows and workflows that include variables as part of the flow, when you run the data flow or workflow you can choose to use the default value, the refresh value, or set a custom value at runtime to use for the current run. See Use Variables in a Data Flow and Use Variables in a Workflow for more information. • Show Queued/Waiting Jobs The Show Queued/Waiting Jobs option helps improve the overall performance when fetching the list of records to display on the Jobs page. See the Search for a Job section in Create and Manage Jobs for more information.

April 2025

Feature	Description
Data Studio Enhancements	The Data Studio is updated with new features and changes: • **The Autonomous Database Federated Database Catalog Enhancements**: You can connect multiple catalogs to make it easier to find and work with your data, including popular Iceberg Data Catalogs like Unity and Polaris. You can also link to remote Autonomous Database instances using Database Links, allowing you to query tables from those databases directly within your catalog — helping you get more from your data catalog without extra steps. One of the major improvements in the Federated Catalog is the ability to easily load a Delta Share directly from the Catalog tool itself — no need to switch to a separate data loading tool. See Inline Data Load of Delta Share from the Catalog Tool for more details. • **The Oracle Autonomous Database for Excel Add-in** can now be installed from the Microsoft Store. See Install the add-in from Microsoft AppSource for more details. • **UX Improvements for AI Table Assistant** - The Table AI Assist tool includes the ability to edit recipes, has auto-save of recipes, UX improvements such as showing new column header in italics. See The Table AI Assist Tool for more details. • **Live Feed Reporting enhancements**: Using the existing Simple Mail Transfer Protocol (SMTP) email registration setup in the Live Feed, you receive email when: – New rows are added: You receive email notifications when Live Feed runs and loads your table with new data. – Data Validation Issues Occur: You receive email notifications when inconsistencies arise, such as when the tool attempts to load non-numeric values into a numeric field. – Live Feed throws an Error: You receive email notifications if errors occur such as when the bucket containing your files is deleted or your credential for that bucket was dropped. See Send Email using Live Feed Tool for more details.
Long Pending Transaction Event	The `LongPendingTransaction` event is generated when a transaction is open for longer than 24 hours. See Information Events on Autonomous Database for more information.
Recovery Time Objective (RTO) for a Cross-region Autonomous Data Guard Standby	The Recovery Time Objective (RTO) for cross-region Autonomous Data Guard Standby databases is lowered to less than ten (10) minutes. See Autonomous Data Guard Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for more information.
Documentation Addition: Automatic Regression Detection	Automatic Regression Detection in Autonomous Database provides proactive handling of regressions and enables automated issue detection, diagnosis, and mitigation. See About Scheduled Maintenance and Patching for more information.
Average Active Sessions and Session Utilization Metrics	The `AverageActiveSessions` metric shows the average number of sessions that are actively executing or waiting for resources. The `SessionUtilization` metric shows the maximum session utilization. See Available Metrics: oci_autonomous_database for more information.

Feature	Description
Credential Comments in `DBMS_CLOUD.CREATE_CREDENTIAL`	The `DBMS_CLOUD.CREATE_CREDENTIAL` procedure supports comments when you use the `comments` parameter. See Create Credentials to Access Cloud Services and CREATE_CREDENTIAL Procedure for more information.
Migration from PostgreSQL, MySQL, or SQL Server to Oracle SQL	You can migrate SQL statements from PostgreSQL, MySQL, or SQL Server, to Oracle SQL and run the statements on Autonomous Database. See Migrate Applications from Non-Oracle SQL to Oracle SQL for more information.
Documentation Change: Pre-Authenticated Request (PAR) URLs Renamed to Table Hyperlinks	You can generate and manage Table Hyperlinks to access data on Autonomous Database. See Use Table Hyperlinks for Read Only Data Access on Autonomous Database for more information.
Additional File Types Supported for Select AI RAG Vector Index Creation	Select AI Retrieval Augmented Generation (RAG) supports multiple file formats, including PDF, DOCX, JSON, XML, and HTML, among others. For more information, see Select AI with Retrieval Augmented Generation (RAG) and, for the complete list of supported file formats, Build your Vector Store.

March 2025

Feature	Description
Elastic Pool with Any Workload Leader	When you create an elastic pool, the workload type of the leader is not restricted. The workload type of an elastic pool leader can be one of: Transaction Processing, Data Warehouse, JSON, or APEX. See About Elastic Pools and About Autonomous Database Workload Types for more information.
SQL Parse Error Warning Event: `SQLParseErrorWarning`	If a SQL parse error occurs several times in a certain time window, for example 200 times or more in an hour, this generates "`Too Many Parse Errors`" entries in the alert log. The `SQLParseErrorWarning` event is triggered when there is a new entry in the alert log, indicating that a SQL statement has too many parse errors. For example, when an invalid query runs 200 times in less than an hour, this triggers the `SQLParseErrorWarning` event. See Information Events on Autonomous Database for more information.
Create Refreshable Clone with 19c Source and 23ai Refreshable Clone	When you create a refreshable clone from an Oracle Database 19c source, you can select Oracle Database 23ai for the refreshable clone. You can use the refreshable clone to test your database on Oracle Database 23ai. After your testing is complete, you can disconnect the refreshable clone to upgrade to Oracle Database 23ai. See Upgrade an Autonomous Database Instance to Oracle Database 23ai and Create a Refreshable Clone for an Autonomous Database Instance for more information.
Offload Queries from Elastic Pool Leader to Member Refreshable Clones: Includes Queries from PL/SQL and from Oracle Scheduler Jobs	When you have heavy read workloads, where a relatively small amount of data is scanned by multiple queries, including from PL/SQL or from Oracle Scheduler Jobs, you can offload queries (reads) from an elastic pool leader to Refreshable Clones. See Offload Queries to Member Refreshable Clones for more information.

Feature	Description
Connections Dropped Metrics	The `ConnectionsDroppedByClient` metric shows the sum of the number of connections dropped by the client. The `ConnectionsDroppedByServer` metric shows the sum of the number of connections dropped by the server. See Available Metrics: oci_autonomous_database for more information.
Oracle APEX 24.2	Autonomous Database uses Oracle APEX Release 24.2. See Create Applications with Oracle APEX in Autonomous Database for more information.
Data Studio Enhancements	The Data Studio is updated with new features and changes: • The Table AI Assist Tool: The Data Studio suite of tools now features an AI-powered Table Assist Tool to effortlessly enhance your local data tables with the power of AI. You can use natural language prompts to augment and correct your information with this tool. See The Table AI Assist Tool for more information. • Catalog Tool Enhancements: The new multi-catalog capabilities of the Autonomous Database catalog mean you can add catalogs to enable data search, discovery and use across a wide range of external systems, including other Autonomous Databases, on-premises databases, cloud storage systems, shared data, and other data catalogs such as Oracle Cloud Infrastructure Data Catalog and AWS Glue. See The Catalog Tool for more information. • The Data Marketplace Tool: Using Data Studio Marketplace, Autonomous Database users can easily collaborate by sharing data with other users of the same or another Autonomous Database. Data providers publish data sets in just minutes, no coding required. Data consumers can browse available data sets, and query live data or copy it to their own Autonomous Database Instance. See The Data Marketplace Tool for more information. • Data Load Livefeed Enhancements: – With the new Basic mode, creating a Livefeed is as simple as selecting a folder from a drop-down menu. – New ability to use Merge semantics in livefeeds. The Livefeed feature automatically updates existing data in your tables with new information, eliminating redundant rows and ensuring data integrity. – With newly introduced **Show code** icon, you can view the underlying SQL code behind your livefeeds and even run the code manually in the SQL worksheet. See Feeding Data for more information.
Documentation Addition: Migrate Existing SQL Execution Plans to Autonomous Database	You can use Real-time SQL Plan Management when you migrate from a source Oracle Database to Autonomous Database. See Migrate Existing SQL Execution Plans to Autonomous Database to Mitigate the Risk of SQL Performance Regressions for more information.

February 2025

Feature	Description
Offload Queries from Elastic Pool Leader to Member Refreshable Clones	When you have heavy read workloads, where a relatively small amount of data is scanned by multiple queries, you can offload queries (reads) from an elastic pool leader to Refreshable Clones. See Offload Queries to Member Refreshable Clones for more information.
Create External Tables Over PAR URLs	You can create an external table over Pre-Authenticated Request (PAR) URLs. See Query Data that Uses Table Hyperlink Source for more information.
Update Cloud Links Registrations	With Cloud Links, a data owner registers a table or view for remote access for a selected audience as defined by the data owner, and the data is then accessible to those with access granted at registration time. After you register a data set you can update data set attributes. See Update Data Set Registration Attributes for more information.
Database User Password Warning Event: `DatabaseUserPasswordWarning`	The database user password warning event provides a message that the Autonomous Database user password is expiring within the specified grace period (default 30 days) or has already expired. See Critical Events on Autonomous Database for more information.

January 2025

Feature	Description
Documentation Addition: Create and Use Custom Metrics	You can create custom metrics to collect, publish, and analyze metrics. A custom metric collects specified data from your Autonomous Database instance and publishes the data to the OCI Monitoring service using a REST API provided by the OCI SDK. See Use Custom Metrics on Autonomous Database for more information.
Disconnect a Peer Database	You can disconnect an Autonomous Data Guard standby or a Backup-Based Disaster Recovery peer. When you disconnect a cross-region peer, the database is permanently disassociated from the Primary database. This converts the database from a peer database to a standalone database. See Disconnect a Peer Database for more information.
Change Patch Level After Provisioning	When you provision or clone an Autonomous Database instance you can select a patch level to apply to upcoming patches. You can also change the patch level after an Autonomous Database instance is provisioned. There are two patch level options: **Regular** and **Early**. See Set the Patch Level for more information.
Service Name Mapping for Cloud Links Consumers	You can set the service name mapping to use when Cloud Links consumers access data from a data set owner. See Set Database Service Name Mapping for Cloud Links Consumers for more information.

Feature	Description
External Tables with ORC, Parquet, and Avro Provide an Option to Set Size of String Columns	When you access ORC, Parquet, or Avro data stored in object store using external tables, you can either automatically or explicitly set the maximum size of text columns. By default the text column size is set based on the value of `MAX_STRING_SIZE`. See Query External Data with ORC, Parquet, or Avro Source Files and Explicitly Set the Text Column Size for more information.
Invoke Google Cloud Platform Cloud Run Functions	You can invoke Google Cloud Platform Cloud Run functions as SQL functions. See Steps to Invoke Google Cloud Run Functions as SQL Functions for more information.

About Autonomous Database

Provides information about Autonomous Database features, workload types, and regions.

- What is Oracle Autonomous Database?
  Autonomous Database provides an easy-to-use, fully autonomous database that scales elastically and delivers fast query performance. As a service, Autonomous Database does not require database administration.
- Key Features of Autonomous Database
  Provides information on key features of Autonomous Database, an affordable, feature-rich service in the cloud.
- About Autonomous Database Workload Types
  Autonomous Database supports different workload types, including: Data Warehouse, Transaction Processing, JSON Database, and APEX Service. Each of these workload types provides performance improvements and additional features that support operations for the specified workload.
- Autonomous Database Region Availability
  Describes availability information for Autonomous Database in commercial regions and government cloud regions.

What is Oracle Autonomous Database?

Autonomous Database provides an easy-to-use, fully autonomous database that scales elastically and delivers fast query performance. As a service, Autonomous Database does not require database administration.

With Autonomous Database you do not need to configure or manage any hardware or install any software. Autonomous Database handles provisioning the database, backing up the database, patching and upgrading the database, and growing or shrinking the database. Autonomous Database is a completely elastic service.

At any time you can scale, increase or decrease, either the compute or the storage capacity. When you make resource changes for your Autonomous Database instance, the resources automatically shrink or grow without requiring any downtime or service interruptions.

Autonomous Database is built upon Oracle Database, so that the applications and tools that support Oracle Database also support Autonomous Database. These tools and applications connect to Autonomous Database using standard SQL*Net connections. The tools and applications can either be in your data center or in a public cloud. Oracle Analytics Cloud and other Oracle Cloud services provide support for Autonomous Database connections.

Autonomous Database provides the following:

- **Increase your app development velocity**: Build and deploy quickly on unified platform supporting vector, JSON, graph, and more. Leverage scalable, cost-effective solutions with flexible cloud, multi-cloud, or on-premises deployment, advanced security, and self-service tools.
- **Fully automated database service**: Supports all modern Oracle data types, workloads including transaction processing, AI and analytics, reducing the need for multiple specialty databases.
- **Maximizes data security and availability**: With Oracle's automation services, you can keep your mission critical applications running by leveraging 99.995% availability.
- **Allows you to bring AI to your data**: Use LLMs with built-in AI Vector Search and your proprietary data to get more accurate answers. There’s no need to duplicate data to a separate vector database.

Autonomous Database also includes the following:

- **Oracle APEX**: a low-code development platform that enables you to build scalable, secure enterprise apps with world-class features.

  See About Oracle APEX for more information.
- **Oracle REST Data Services (ORDS)**: a Java Enterprise Edition based data service that makes it easy to develop modern REST interfaces for relational data and JSON Document Store.

  See Developing RESTful Services in Autonomous Database for more information.
- **Database Actions**: is a web-based interface that uses Oracle REST Data Services to provide development, data tools, administration, and monitoring features for Autonomous Database.

  See About Database Actions for more information.
- **Oracle Machine Learning**: is a set of components supporting data scientists, ML engineers, and data analysts, as well as SQL, R, and Python users. Components include SQL, R, and Python APIs, a built-in notebook interface, and no-code interfaces for AutoML, data and model monitoring, and model deployment, along with model management, deployment, and monitoring via REST endpoints.

  See Oracle Machine Learning for more information.

Key Features of Autonomous Database

Provides information on key features of Autonomous Database, an affordable, feature-rich service in the cloud.

Key Features

- **Managed**: Oracle simplifies end-to-end management of the database:
  - Provisioning new databases
  - Growing or shrinking storage and compute resources
  - Patching and upgrades
  - Backup and recovery
- **Fully elastic scaling**: Scale compute and storage independently to fit your database workload with no downtime:

  - Size the database to the exact compute and storage required
  - Scale the database on demand: Independently scale compute or storage
  - Shut off idle compute to save money
- **Auto scaling**: Allows your database to use more CPU and IO resources or to use additional storage automatically when the workload or storage demand requires additional resources:
  - Specify the number of CPUs for your Autonomous Database workload.
  - Use compute auto scaling to allow the database to use up to three times more CPU and IO resources, depending on workload requirements. Compute auto scaling is enabled by default when you create an Autonomous Database.
  - Use storage auto scaling to allow the database to expand to use up to three times the reserved base storage, depending on your storage requirements. Storage auto scaling is disabled by default when you create an Autonomous Database.
  - Manage auto scaling from the Oracle Cloud Infrastructure Console to enable or disable compute auto scaling or storage auto scaling for your Autonomous Database.
- **Autonomous Database supports**:
  - Existing applications, running in the cloud or on premises
  - Connectivity via SQL*Net, JDBC, ODBC

    See Connect to Autonomous Database for more information.
  - Third-party data-integration tools
  - Oracle cloud services: Oracle Analytics Cloud, Oracle GoldenGate Marketplace, and others
- **High-performance queries and concurrent workloads**: Optimized query performance with preconfigured resource profiles for different types of users.
- **Oracle SQL**: Autonomous Database is compatible with existing applications that support Oracle Database.
- **Built-in web-based data analysis tool**: Web-based notebook tool for designing and sharing SQL based data-driven, interactive documents.
- **Database migration utility**: Easily migrate from MySQL, Amazon AWS Redshift, PostgreSQL, SQL Server, and other databases.

Simple Cloud-based Data Loading

Autonomous Database provides:

- Fast, scalable data-loading from Oracle Cloud Infrastructure Object Storage, Azure Blob Storage or Azure Data Lake Storage, Amazon S3, Amazon S3-Compatible, GitHub Repository, Google Cloud Storage, or on-premises data sources.

Oracle Database Actions

Database Actions is a web-based interface that uses Oracle REST Data Services to provide development, data tools, and administration and monitoring features for Autonomous Database, including the following:

- **Development Tools**
  - SQL Navigator and Worksheet: view objects and enter and run SQL and PL/SQL statements, and create database objects

  - Data Modeler: provides an integrated version of Oracle SQL Developer Data Modeler with basic reporting features. You can create diagrams from existing schemas, retrieve data dictionary information, generate DDL statements, and export diagrams
  - REST: An IDE for your REST APIs that enables you to manage templates, handlers and OAuth clients, generate API documentation, and test APIs.
  - LIQUIBASE: View ChangeLogs applied to your schema.
  - JSON: Create collections, upload documents, query and filter your data, create diagrams for your JSON document structures, and create relational views and indexes.
  - Charts: Use SQL queries to build rich charts and dashboards containing multiple charts.
  - Scheduling: An interface for `DBMS_SCHEDULER` that enables you to monitor jobs, view execution history, forecast upcoming jobs, and visualize scheduler chains.
  - Oracle Machine Learning: provides several components accessible through a common user interface. OML Notebooks supports Python, R, SQL, PL/SQL, and Markdown interpreters, with access to in-database ML through OML4Py, OML4R, and OML4SQL. OML Models supports managing and deploying in-database models. OML AutoML UI provides a no-code user interface to build, evaluate, and deploy in-database models using automated machine learning.
  - APEX: Login to APEX, develop and run rich, low-code web applications.
  - Graph Studio: Oracle Graph Studio lets you create property graph databases and automates the creation of graph models and in-memory graphs from database tables.
- **Data Studio**
  - Data Load: load or access data from local files or remote databases.
  - Catalog: understand data dependencies and the impact of changes.
  - Data Insights: discover anomalies, outliers and hidden patterns in your data.
  - Data Analysis: analyze your data
  - Data Transforms: transform data for analysis and other applications.
- **Administration and Monitoring**
  - Manage users
  - Database Dashboard: Monitor database activity charts such as CPU usage, number of executing SQL statements, and wait events formerly found on your Autonomous Database Service Console.
  - Performance Hub: Access SQL Monitoring reports and Active Session History (ASH) Analytics.

Disaster Recovery Options

Autonomous Data Guard Autonomous Database provides Autonomous Data Guard to enable a standby (peer) database to provide data protection and disaster recovery for your Autonomous Database instance.When you add an Autonomous Data Guard standby database, the system creates a standby database that continuously gets updated with the changes from the primary database. You can use Autonomous Data Guard with a standby in the current region, a local standby, or with a standby in a different region, a cross-region standby. You can also use Autonomous Data Guard with both a local standby and a cross-region standby.

Backup-Based Disaster Recovery uses backups to instantiate a peer database at the time of switchover or failover. This enables you to have a lower cost and higher Recovery Time

Objective (RTO) disaster recovery option for your Autonomous Database, as compared with Autonomous Data Guard. For local backup-based disaster recovery, existing local backups are utilized. There are no additional costs for a local Backup-Based Disaster Recovery. Cross-Region Backup-Based Disaster Recovery incurs an additional cost.

SQL Developer Support

Using Autonomous Database with SQL Developer you can do the following:

- Connect to Autonomous Database
- Create tables, indexes, and materialized views in Autonomous Database
- Load data into an Autonomous Database
- Copy tables to Autonomous Database
- Transfer a schema to Autonomous Database

Business Intelligence Tools Support

Autonomous Database is compatible with a number of business intelligence and data visualization tools from Oracle and from trusted third parties:

- Oracle Analytics Cloud
- Oracle Analytics Desktop
- Third-party Business Intelligence tools

About Autonomous Database Workload Types

Autonomous Database supports different workload types, including: Data Warehouse, Transaction Processing, JSON Database, and APEX Service. Each of these workload types provides performance improvements and additional features that support operations for the specified workload.

- About Oracle Autonomous Data Warehouse
  Autonomous Database is designed as a "load and go" service: you start the service, define tables, load data, and then run queries.
- About Oracle Autonomous Transaction Processing
  Autonomous Database is designed to support all standard business applications and deliver scalable query performance.
- About Autonomous JSON Database
  Oracle Autonomous JSON Database is Oracle Autonomous Transaction Processing, but designed for developing NoSQL-style applications that use JavaScript Object Notation (JSON) documents. You can promote an Autonomous JSON Database service to an Autonomous Transaction Processing service.
- About Oracle APEX Application Development
  Oracle APEX Application Development (APEX Service) is a low cost, Oracle Cloud service offering convenient access to the Oracle APEX platform for rapidly building and deploying low-code applications.

About Oracle Autonomous Data Warehouse

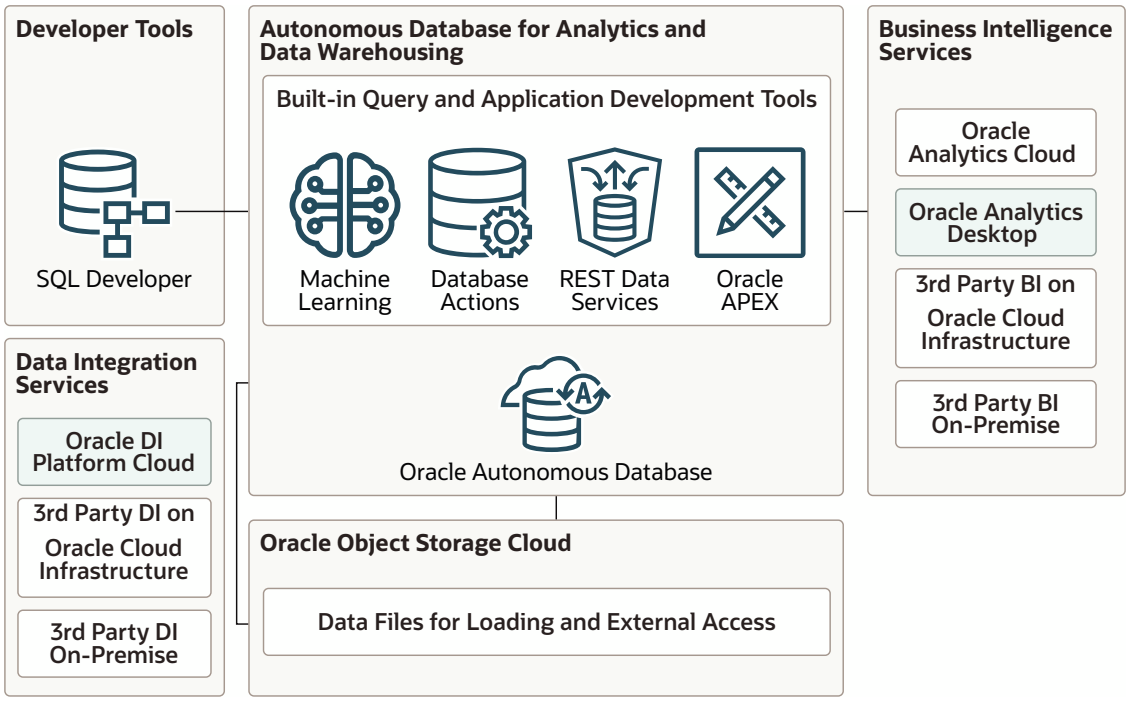
Autonomous Database is designed as a "load and go" service: you start the service, define tables, load data, and then run queries.

Autonomous Database is designed to support all standard SQL and business intelligence (BI) tools, and provides all of the performance of the market-leading Oracle Database in an environment that is tuned and optimized for data warehouse workloads.

To get started you create an Autonomous Database with workload type Data Warehouse and specify the number of ECPUs and the storage capacity in TB's for the Autonomous Database.

You can use Autonomous Database with Oracle Analytics Cloud or Oracle Analytics Desktop to easily create visualizations and projects that reveal trends in your company's data and help you answer questions and discover important insights about your business.

The following figure shows the Autonomous Database architecture with related components for analytics and data warehousing.



About Oracle Autonomous Transaction Processing

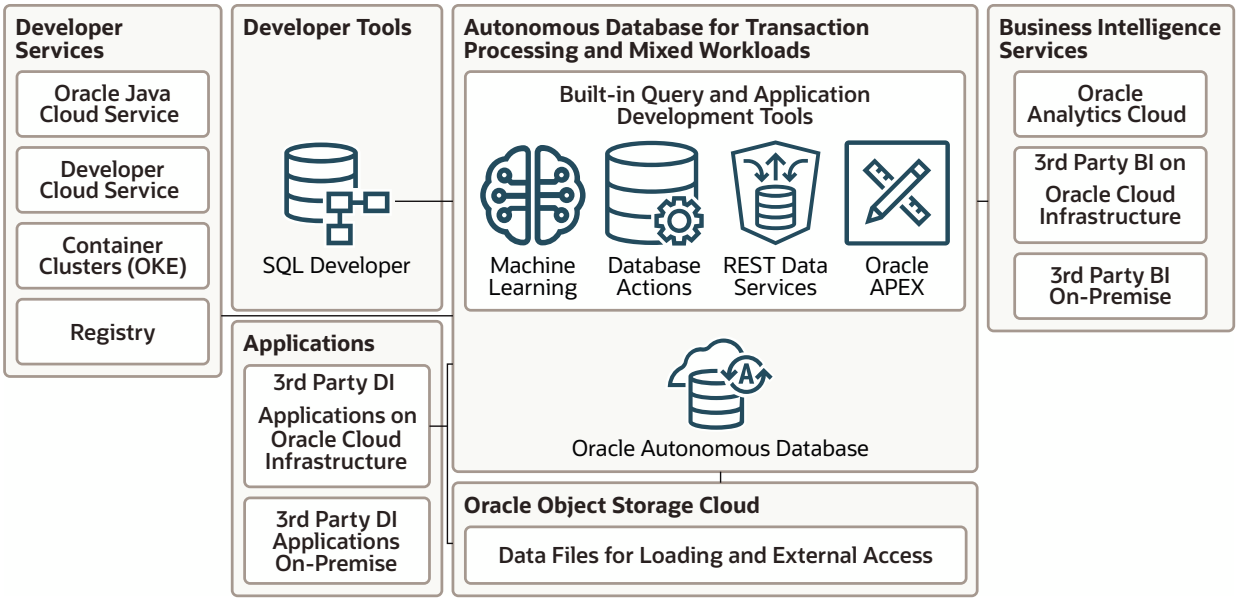
Autonomous Database is designed to support all standard business applications and deliver scalable query performance.

Autonomous Database provides all of the performance of the market-leading Oracle Database in an environment that is tuned and optimized to meet the demands of a variety of applications, including: mission-critical transaction processing, mixed transactions and analytics, IoT, and JSON document store.

To get started you create an Autonomous Database with the workload type Transaction Processing and specify the number of ECPUs and the storage capacity for the database.

You can use Autonomous Database with Oracle Analytics Cloud or Oracle Analytics Desktop to easily create visualizations and projects that reveal trends in your company's operational data and help you answer questions and discover important insights about your business.

The following figure shows the Autonomous Database architecture with related components for transaction processing and mixed workloads.



About Autonomous JSON Database

Oracle Autonomous JSON Database is Oracle Autonomous Transaction Processing, but designed for developing NoSQL-style applications that use JavaScript Object Notation (JSON) documents. You can promote an Autonomous JSON Database service to an Autonomous Transaction Processing service.

Oracle Autonomous JSON Database provides all of the same features as Autonomous Transaction Processing, with this *important limitation*: you can store only up to 20 GB of data other than JSON document collections. There is no storage limit for JSON collections.

Development of NoSQL-style, document-centric applications is particularly flexible because the applications use *schemaless* data. This lets you quickly react to changing application requirements. There's no need to normalize the data into relational tables, and no impediment to changing data structure or organization at any time, in any way. A JSON document has internal structure, but no relation is imposed on separate JSON documents.

With Oracle Autonomous JSON Database your JSON document-centric applications typically use Simple Oracle Document Access (SODA), which is a set of NoSQL-style APIs for various application-development languages and for the representational state transfer (REST) architectural style. You can use any SODA API to access any SODA collection.

SODA document collections are backed by ordinary database tables and views. To use other kinds of data, subject to the 20 GB limit, you typically need some knowledge of Structured Query Language (SQL) and how that data is stored in the database.

With Oracle Autonomous JSON Database, a SODA collection *can only contain JSON data*. For example, you cannot have a collection of image documents or a collection that contains both JSON documents and image documents. This is a limitation relative to Autonomous Transaction Processing, where you can define such heterogeneous collections.

No matter what kind of data your applications use, whether JSON or something else, you can take advantage of all Oracle Database features. This is true regardless of the kind of Oracle Autonomous Database you use.

JSON data is stored natively in the database. In a SODA collection on an Autonomous Database JSON data is stored in Oracle's native binary format, OSON.

About Oracle APEX Application Development

Oracle APEX Application Development (APEX Service) is a low cost, Oracle Cloud service offering convenient access to the Oracle APEX platform for rapidly building and deploying low-code applications.

APEX Service is designed to support all standard business applications and deliver scalable query performance. You can promote APEX Service to Autonomous Transaction Processing service.

See Oracle APEX Application Development and Upgrade APEX Service to Autonomous Transaction Processing for more information.

Autonomous Database Region Availability

Describes availability information for Autonomous Database in commercial regions and government cloud regions.

Commercial Regions and Availability Domains

Autonomous Database is available in all regions of the commercial realm.

Government Cloud

See Oracle Cloud Infrastructure Government Cloud for information about availability in Government Cloud regions.

Get Help, Search Forums, and Contact Support

When you use Oracle Autonomous Database, sometimes you need to get help from the community or to talk to someone in Oracle support. This topic provides more information about getting help by viewing and posting questions on forums and using Oracle Cloud Support to create a support request.

- Post Questions on Forums
  If you can't find an answer to a question through search, you can submit a question to one of the forums. This option is available to all customers.
- Open a Support Ticket to Create a Service Request
  If forums and other search options do not resolve your issue and you need to talk to someone, you can create a support request.
- Create a Service Request in My Oracle Cloud Support
- Register to My Oracle Cloud Support as a New User

Post Questions on Forums

If you can't find an answer to a question through search, you can submit a question to one of the forums. This option is available to all customers.

Cloud Customer Connect Forums

For any issue related to Autonomous Database or Oracle Cloud Infrastructure, you can post a question to Cloud Customer Connect:

- Autonomous Data Warehouse
- Autonomous Transaction Processing
- Oracle Cloud Infrastructure and Platform

Stack Overflow Knowledge Forum

If you are working with Autonomous Database and you have technical questions, you can use stackoverflow to post questions and to find answers or to help others answer their questions. When you post, tag your question with oracle-autonomous-db, as follows:

Questions tagged [oracle-autonomous-db]

Open a Support Ticket to Create a Service Request

If forums and other search options do not resolve your issue and you need to talk to someone, you can create a support request.

If you need to file a service request use Oracle Cloud Support or contact your support representative and provide the tenancy details. In addition to support for technical issues, you can open support requests if you need to:

Create a Support Ticket from Oracle Cloud Infrastructure Console

The first time you open a support ticket, you're automatically taken through a series of steps to provision your support account. If you want to make changes or if you run into problems, see Configuring Your Oracle Support Account.

To create a support request from the Oracle Cloud Infrastructure Console:

1. On the Oracle Cloud Infrastructure open the **Help** menu (?) and under **Request Help**, click **Create Support Request**.
2. Enter the following:
   - **Issue Summary**: Enter a title that summarizes your issue. Avoid entering confidential information.
   - **Describe Your Issue**: Provide a brief overview of your issue.
     - Include all the information that support needs to route and respond to your request.
       See Obtain Tenancy Details for details on obtaining Autonomous Database information.
     - Include troubleshooting steps taken and any available test results.
   - Select the severity level for this request.
3. Click **Create Support Request.**

Create a Service Request in My Oracle Cloud Support

If you are having a technical issue using Oracle Autonomous Database, you can contact Oracle Support by creating a service request in My Oracle Cloud Support.

If you are an Oracle Commercial Cloud customer, and you need the ability to request support for a tenancy that your cloud administrator created *after November 8th, 2024*, then you need to sign into My Oracle Cloud Support using **Sign in with your commercial cloud account**, and enter the new tenancy credentials. If you are unsure about your tenancy credentials, please contact your cloud administrator.

> **Tip:**
>
> Before following these steps, sign in to your Oracle Cloud account at cloud.oracle.com in a separate browser tab or window so that you can quickly and easily get any information you need to complete your service request.

- Oracle Commercial Cloud customer
- Oracle Non-Commercial Cloud Customer

Oracle Commercial Cloud customer

If you are an Oracle Commercial Cloud customer and you know your cloud account login information:

1. Go to support.oracle.com.
2. Enter your *email address* and click **Next**.
   - If you see a message that your email address is not recognized, follow the instructions outlined in Register to My Oracle Cloud Support as a New User.
   - If your email is registered with Oracle Support, you'll be directed to sign in to the applicable support portal for your products and services. Continue with step 3.
3. Click **Sign in with your commercial cloud account**.
4. Enter your *OCI credentials* (Tenancy, domain if applicable, username, and password).

   If this is your first time signing in to My Oracle Cloud Support, you will be prompted to associate your cloud account to My Oracle Cloud Support by verifying your email.
5. Accept the **Terms of Use**.
6. Complete your **personal Info**.
   - For **Security Pin**, enter a 6 digit number not starting with a 0. When you call Oracle support line, you will be asked to provide this pin.
   - If this is the first time you are filling the **personal Info**, you may have to verify your **Mobile Phone** using a one-time password sent to your mobile number.
7. Click **Save** an you will be routed to the **My Oracle Cloud Support** portal.
8. To create a new service request, click **Create Service Request**.

9. Provide the details prompted on the **Create a new service request** page, clicking **Next**, as needed.
10. When you see **Proceed with the next step to submit the SR.**, click **Next**.

Oracle Non-Commercial Cloud Customer

If you are an Oracle Cloud customer and you do not have cloud account login information OR you have a non-commercial cloud account (government realm or dedicated realm):

1. Go to support.oracle.com.
2. Enter your *Oracle account username* and *password*.
3. If this your first time signing in to My Oracle Cloud Support, you will be prompted to accept the **Terms of use**, and complete your *personal profile*. Refer to Register to My Oracle Cloud Support as a New User for more details.
4. Once you get access to My Oracle Cloud Support, sign in to support.oracle.com.
5. You are routed to the **My Oracle Cloud Support** portal.
6. To create a new service request, click **Create Service Request**.
7. Provide the details prompted on the **Create a new service request** page, clicking **Next**, as needed.
8. When you see **Proceed with the next step to submit the SR.**, click **Next**.

> **Note:**
>
> If you are the first user logging in using a specific Subscription ID and Account name, you will be prompted to accept the Administrator role.

Register to My Oracle Cloud Support as a New User

- Oracle Commercial Cloud customer
- Oracle Non-Commercial Cloud Customer

Oracle Commercial Cloud customer

If you are an Oracle Commercial Cloud customer:

1. Go to support.oracle.com.
2. Enter your *email address* and click **Next**.
3. Click **Sign in** with your commercial cloud account.
4. Enter your *Tenancy*, then click **Continue**.

   If you do not know your tenancy information, please contact your cloud administrator.
5. Choose an **Identity domain** if applicable, then click **Next**.
6. Enter your *cloud account user name* and *password*.

7. Associate your cloud account to My Oracle Cloud Support by verifying your email.
8. Accept **Terms of Use**.
9. Complete your **personal Info**.
    - For **Security Pin**, enter a 6 digit number not starting with a 0. When you call Oracle support line, you will be asked to provide this pin.
    - If this is the first time you are filling the **personal Info**, you may have to verify your **Mobile Phone** using a one-time password sent to your mobile number.
10. You will be directed to the My Oracle Cloud Support profile page as an **approved user**.

Oracle Non-Commercial Cloud Customer

If you are an Oracle Cloud customer and you do not have cloud account login information OR you have a non-commercial cloud account (government realm or dedicated realm):

1. Click **Create an Oracle Account**.
2. In **Create Your Oracle Account**, provide the required information and click **Create Account**.
3. You will receive a verification email to enable your new account. Once your email has been verified, you will be prompted to log in with your newly created Oracle account username and password.
4. Enter *subscription ID* and *Account or Company name* to complete identity verification.

    Contact your Cloud administrator for the subscription ID and Account name.
5. Complete your **personal Info**.
    - For **Security Pin**, enter a 6 digit number not starting with a 0. When you call Oracle support line, you will be asked to provide this pin.
    - If this is the first time you are filling the **personal Info**, you may have to verify your **Mobile Phone** using a one-time password sent to your mobile number.
6. You will be directed to the My Oracle Cloud Support profile page as a **pending user**. If approval is not granted within a reasonable amount of time, contact your administrator. Once you are approved, you will have access to My Oracle Cloud Support.

> **Note:**
>
> The first user logging in using a Subscription ID and Account name will be prompted to accept the Admin or CUA role.

Service Level Objectives (SLOs)

Describes the Service Level Objectives (SLOs) for Oracle Autonomous Database Serverless.

- Recovery Time Objective and Recovery Point Objective
  Oracle Autonomous Database Serverless is engineered to return an application online following an unplanned outage or a planned maintenance activity within single-digit seconds.

- Built-In Tool Availability
  Oracle will use commercially reasonable efforts to have the following tools with a Monthly Uptime Percentage (Availability) objective as defined below, during the service commitment of a calendar month.
- Zero-Regression Service Level Objective
  Oracle Autonomous Database Serverless automatically applies patches on your database. Oracle provides a service level objective of zero regressions in your production database due to these patches.

Recovery Time Objective and Recovery Point Objective

Oracle Autonomous Database Serverless is engineered to return an application online following an unplanned outage or a planned maintenance activity within single-digit seconds.

The following table outlines the target Recovery Time Objective (RTO) SLOs for different failure events.

Event	Recovery Time Objective (RTO) Service Level Objective	Recovery Point Objective (RPO) Service Level Objective, Maximum Possible Data Loss
Events requiring failover to a standby database when Autonomous Data Guard is enabled, such as: • Data corruptions • Full database failures • Complete storage failures • Availability Domain or Region failures	With an Autonomous Data Guard standby: **Local Standby Database**: two (2) minutes **Cross-Region Standby Database:** less than ten (10) minutes See Autonomous Data Guard Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for more information.	**Local Standby Database**: The maximum possible data loss objective is < 10 seconds **Cross-Region Standby Database:** The maximum possible data loss objective is 1 minute
Events requiring failover to a backup copy or restoring from a backup (when there is not an Autonomous Data Guard standby) such as: • Data corruptions • Full database failures • Complete storage failures • Availability Domain or Region failures	With Backup-Based Disaster Recovery, based on the size of the database: **Local backup copy**: 1 hour + (1 hour per 5 TB) **Cross-Region backup copy**: 1 hour + (1 hour per 5 TB) For example, a 5 TB database has a Recovery Time Objective (RTO) of up to 2 hours. See Backup-Based Disaster Recovery Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for more information.	**Local backup copy**: The maximum possible data loss objective is < 10 seconds **Cross-Region backup copy**: The maximum possible data loss objective is 1 minute

Built-In Tool Availability

Oracle will use commercially reasonable efforts to have the following tools with a Monthly Uptime Percentage (Availability) objective as defined below, during the service commitment of a calendar month.

Oracle Autonomous Database Built-in Tool	Availability Service Level Objective
Oracle REST Data Services (ORDS)	99.95%
Oracle APEX (on Autonomous Database)	99.95%
Database Actions	99.95%
Graph Studio	99.95%
Oracle Machine Learning: • Oracle Machine Learning Notebooks • Oracle Machine Learning Services • Oracle Machine Learning for Python/R • Oracle Machine Learning AutoML User Interface	99.95%
Data Transforms	99.95%

The following terms apply to the Availability Service Level Objective for the Oracle Autonomous Database Built-in Tools listed in this table:

- “HTTP Error Rate” applies separately to each Database Built-in Tool and means the percentage value corresponds to: (i) the total number of failed HTTP calls made to the applicable tool with a status of “Bad Gateway” or “Service Unavailable” in a minute period during a calendar month, divided by (ii) the total number of HTTP calls made to the tool in a minute period.
- “Monthly Uptime Percentage” is calculated by subtracting from 100%, the average of the HTTP Error Rate for each minute period during the applicable calendar month.

Oracle Autonomous Database Built-in Tool	Availability Service Level Objective
Oracle Database API for MongoDB	99.95%

The following terms apply to the Availability Service Level Objective for Oracle Database API for MongoDB:

- A "Connection through Oracle Database API for MongoDB" is a direct connection established from any tool or application to the Cloud Service using Oracle Database API for MongoDB.
- “Monthly Uptime Percentage” is calculated by subtracting from 100%, the percentage of minutes during the calendar month in which the applicable Cloud Service was "Unavailable".
- “Unavailable” means a minute period when: (i) no connection through Oracle Database API for MongoDB is or can be established and (ii) all continuous attempts, at least five, to establish such a connection fail.

Zero-Regression Service Level Objective

Oracle Autonomous Database Serverless automatically applies patches on your database. Oracle provides a service level objective of zero regressions in your production database due to these patches.

Oracle Autonomous Database Serverless automatically applies patches containing bug and security fixes during the maintenance windows announced on your database console. When you provision your test or pre-prod database using the Early patch level option, you can test the patches on these instances before the patches are applied to your production database. If you see issues in your test or pre-prod database, you can file service requests to get the issues addressed before the patch is applied to your production database.

Oracle provides a service level objective of zero regressions in your production database. "Regression" in this document is described as issues introduced by patches or updates to the Autonomous Database performed during the maintenance window announced on your database console.

After a patch is applied to your database on the Early patch level, if you find and report an issue on that database through a service request, Oracle will use commercially reasonable efforts to address the problem so that the same issue does not occur in your production database.

See Set the Patch Level for more information.

How Is Autonomous Database Billed?

Provides details for Autonomous Database billing.

- Autonomous Database Billing Summary
  You are billed for Autonomous Database usage according to the values of two parameters: **compute** and **storage**. You select values for these parameters when you provision or scale an Autonomous Database instance.
- Compute Models in Autonomous Database
  Autonomous Database provides the ECPU compute model when you create or clone an instance. The legacy OCPU billing metric has been retired and Oracle recommends using ECPUs for all new and existing Autonomous Database deployments.
- Oracle Autonomous Database Serverless Features Billing
- Oracle Autonomous Database Serverless Billing for Database Tools
  Shows billing information for Autonomous Database database tools for ECPU and OCPU compute models.
- Autonomous Database for Developers Billing and Tenancy Service Limit
  You are billed for Autonomous Database for Developers usage hourly, on a per instance basis.

Autonomous Database Billing Summary

You are billed for Autonomous Database usage according to the values of two parameters: **compute** and **storage**. You select values for these parameters when you provision or scale an Autonomous Database instance.

When you provision an instance you select a workload type based on the workloads you want to run on your database. The pricing for compute and storage resources is based on the workload type you select.

The workload types are:

- **Data Warehouse**: optimized for analytic workloads, including data marts, data warehouses, data lakes, and data lakehouses.
- **Transaction Processing**: optimized to run transactional, analytical, and batch workloads concurrently.
- **JSON**: a cloud document database service that makes it simple to develop JSON-centric applications.
- **APEX**: a fully managed, low-code application development platform for building and deploying modern cloud, mobile, and data-driven applications.

Depending on the workload type, you select an Autonomous Database billing model:

- **ECPU Billing Model**: (Recommended billing model) Based on the number of cores allocated from the pool of compute servers and storage servers.
- **OCPU Billing Model**: (Legacy billing model) Provides compute capacity equivalent of a physical core of a processor with hyper-threading enabled.

See the Autonomous Database Cloud Price List for up-to-date pricing for each workload type.

Use the Cloud Cost Estimator to estimate the monthly cost for an Autonomous Database.

For complete cost details, refer to the *Oracle PaaS and IaaS Universal Credits Service Descriptions* document available on the Oracle Cloud Services contracts page.

Compute and Storage Billing Summary

- ECPU Compute Model
- OCPU Compute Model

ECPU Compute Model

Compute Billing

Compute billing is based on the value you specify for **ECPU Count** and on the additional features you use. For example, if you use additional compute resources with **Compute auto scaling** selected, you are billed for the additional compute resources.

- Data Warehouse workload: See Billing Information: Autonomous Data Warehouse ECPU Compute Model
- Transaction Processing workload: See Billing Information: Autonomous Transaction Processing ECPU Compute Model
- JSON workload: See Billing Information: Autonomous JSON Database ECPU Compute Model
- APEX workload: See Billing Information: APEX Service ECPU Compute Model

Storage Billing

Billing for database storage is based on the value you specify for **Storage** and based on the features you use that require additional storage. For example, you are billed for additional storage usage if you select **Storage auto scaling** and you use additional storage resources.

The storage for backups is billed in addition to database storage.

- Data Warehouse workloads:
  - Database storage is provisioned in Terabytes (TBs) and charged in TB increments. The minimum database storage size is 1 TB.
  - Backup storage is separately charged in TB increments.

  See Billing Information: Autonomous Data Warehouse ECPU Compute Model.
- Transaction Processing, JSON, or APEX workloads:
  - Database storage is provisioned in Gigabytes or Terabytes (GBs or TBs) and charged in Gigabyte increments. The minimum database storage size is 20 GB.
  - Backup storage is separately charged in GB increments.

  See ECPU Compute Model Billing Information.

Also see Oracle Autonomous Database Serverless Features Billing

OCPU Compute Model

Compute Billing

Compute billing is based on the value you specify for **OCPU Count** and on the additional features you use. For example, if you use additional compute resources with **Compute auto scaling** selected, you are billed for the additional compute resources.

- Data Warehouse workload: See Billing Information: Autonomous Data Warehouse OCPU Compute Model
- Transaction Processing, JSON, or APEX workloads: See Billing Information: Autonomous Transaction Processing OCPU Compute Model

Storage Billing

Billing for database storage is based on the value you specify for **Storage** and based on the features you use that require additional storage. For example, you are billed for additional storage usage if you select **Storage auto scaling** and you use additional storage resources.

- For all workload types database storage is provisioned in Terabytes (TBs) and charged in TB increments.
- For all workload types, the storage for automatic backups is included in the cost of database storage at no additional cost.

See OCPU Compute Model Billing Information and Oracle Autonomous Database Serverless Features Billing

Compute Models in Autonomous Database

Autonomous Database provides the ECPU compute model when you create or clone an instance. The legacy OCPU billing metric has been retired and Oracle recommends using ECPUs for all new and existing Autonomous Database deployments.

- **ECPU**: An ECPU is an abstracted measure of compute resources. ECPUs are based on the number of cores elastically allocated from a pool of compute and storage servers.

  While provisioning a new database or cloning an existing database:

  – The CPU count defaults to 2 ECPUs.

  – For databases that need more than 2 ECPUs, you must specify the number of assigned ECPUs as an integer. For example, you cannot assign 3.5 ECPUs to a database. The next available number of ECPUs above 3 is 4.

- **OCPU**: An OCPU is a physical measure of compute resources. OCPUs are based on the physical core of a processor with hyper-threading enabled.

> **Note:**
>
> OCPU is a legacy billing metric and has been retired on Autonomous Database. ECPUs are the replacement billing metric for all new and existing Autonomous Database deployments. See Oracle Support Document 2998742.1 for more information.

See Provision an Autonomous Database Instance and Clone an Autonomous Database Instance for more information.

- ECPU Compute Model Billing Information
- OCPU Compute Model Billing Information
  Provides billing information for Autonomous Transaction Processing, Autonomous Data Warehouse, Autonomous JSON Database, and APEX Service with the OCPU compute model.

ECPU Compute Model Billing Information

Provides billing information for Autonomous Transaction Processing, Autonomous Data Warehouse, Autonomous JSON Database, and APEX Service with the ECPU compute model. Cost and usage reports are available from the Oracle Cloud Infrastructure Console, under **Billing and Cost Management**.

Cost reports look like this: `reports/cost-csv/0001000001010101-00001.csv.gz`. Download the zip file and extract to view the details.

Usage reports look like this: `reports/usage-csv/0001000001010101.csv.gz`. Download the zip file extract to view the details.

> **Note:**
>
> You can use OCI usage reports and `OCI_USAGE_DATA` view for information on usage *before* January 31st, 2025. The OCI usage reports are deprecated. You can continue to access your existing usage report CSV files until July 31, 2025. See Cost and Usage Reports for more information.

- Billing Information: Autonomous Transaction Processing ECPU Compute Model
  Shows billing information for Autonomous Transaction Processing with the ECPU compute model.
- Billing Information: Autonomous Data Warehouse ECPU Compute Model
- Billing Information: Autonomous JSON Database ECPU Compute Model
  Shows billing information for Autonomous JSON Database with the ECPU compute model.
- Billing Information: APEX Service ECPU Compute Model
  Shows billing information for APEX Service with the ECPU compute model.
- Billing Information: Autonomous Database for Developers

Billing Information: Autonomous Transaction Processing ECPU Compute Model

Shows billing information for Autonomous Transaction Processing with the ECPU compute model.

Resource Description	SKU Name	Part Number	Metric	Minimum and Minimum Increment
ECPU usage	Oracle Autonomous Transaction Processing – ECPU	B95702	ECPU per hour	Minimum ECPUs: 2 Minimum increment: 1

Resource Description	SKU Name	Part Number	Metric	Minimum and Minimum Increment
ECPU usage (Bring your own license)	Oracle Autonomous Transaction Processing – ECPU – BYOL	B95704	ECPU per hour	Minimum ECPUs: 2 Minimum increment: 1
ECPU usage for built-in-tools: Graph Studio, Oracle Machine Learning, Data Transforms	Oracle Autonomous Transaction Processing – ECPU – BYOL	B95704	ECPU per hour	Minimum ECPUs: 8 Minimum increment: 1
Database Storage usage	Oracle Autonomous Database Storage for Transaction Processing	B95706	Gigabyte Storage Capacity Per Month	1 GB Minimum storage size is 20 GB
Backup Storage usage	Oracle Autonomous Database Storage	B95754	Gigabyte Storage Capacity Per Month	1 GB Backups are billed per GB

The following tables show how to view cost and usage information for Autonomous Transaction Processing with the ECPU compute model.

See Cost and Usage Reports Overview for details on the cost and usage reports.

ECPU Cost and Usage Information for Oracle Autonomous Transaction Processing for Records Until January 31, 2025

Resource Description	Part Number	Cost Report	Usage Report
ECPU usage	B95702	Cost report: in the `product/Description` column, find, "Oracle Autonomous Transaction Processing – ECPU"	Usage report: in the `product/resource` column, find `PIC_ATP_ECPU_LI`
ECPU usage (Bring your own license)	B95704	Cost report: in the `product/Description` column, find, "Oracle Autonomous Transaction Processing – ECPU – BYOL"	Usage report: in the `product/resource` column, find `PIC_ATP_ECPU_BYOL`
ECPU usage for built-in-tools: Graph Studio, Oracle Machine Learning, Data Transforms	B95704	Cost report: in the `product/Description` column, find, "Oracle Autonomous Transaction Processing – ECPU – BYOL"	Usage report: in the `product/resource` column, find `PIC_ADBS_TOOLS_ECPU`
Database Storage usage	B95706	Cost report: in the `product/Description` column, find, "Oracle Autonomous Database Storage for Transaction Processing"	Usage report: in the `product/resource` column, find `PIC_ADB_STORAGE_ATP`
Backup Storage usage	B95754	Cost report: in the `product/Description` column, find, "Oracle Autonomous Database Backup Storage"	Usage report: in the `product/resource` column, find `PIC_ADB_STORAGE`

ECPU Cost and Usage Information for Oracle Autonomous Transaction Processing for Records After January 31, 2025

Resource Description	Part Number	Cost Report	Usage Information
ECPU usage	B95702	To find cost information, in the cost report find your OCID in the `product/resourceId` column and in the `product/Description` column, find, "Oracle Autonomous Transaction Processing – ECPU"	Usage is in the cost report in the `usage/billedQuantity` column.
ECPU usage (Bring your own license)	B95704	To find cost information, in the cost report find your OCID in the column `product/resourceId` and in the `product/Description` column, find, "Oracle Autonomous Transaction Processing – ECPU – BYOL"	Usage is in the cost report in the `usage/billedQuantity` column.
ECPU usage for built-in-tools: Graph Studio, Oracle Machine Learning, Data Transforms	B95704	To find cost information, in the cost report find your OCID in the `product/resourceId` column and in the `product/Description` column, find, "Oracle Autonomous Transaction Processing – ECPU – BYOL"	Usage is in the cost report in the `usage/billedQuantity` column.
Database Storage usage	B95706	To find cost information, in the cost report find your OCID in the `product/resourceId` and in the `product/Description` column, find, "Oracle Autonomous Database Storage for Transaction Processing"	Usage is in the cost report in the `usage/billedQuantity` column.
Backup Storage usage	B95754	To find cost information, in the cost report find your OCID in the `product/resourceId` column and in the `product/Description` column, find, "Oracle Autonomous Database Backup Storage"	Usage is in the cost report in the `usage/billedQuantity` column.

Billing Information: Autonomous Data Warehouse ECPU Compute Model

Shows billing information for Autonomous Data Warehouse with the ECPU compute model.

Resource Description	SKU Name	Part Number	Metric	Minimum and Minimum Increment
ECPU usage	Oracle Autonomous Data Warehouse – ECPU	B95701	ECPU per hour	Minimum ECPUs: 2 Minimum increment: 1
ECPU usage (Bring your own license)	Oracle Autonomous Data Warehouse – ECPU – BYOL	B95703	ECPU per hour	Minimum ECPUs: 2 Minimum increment: 1
ECPU usage for built-in-tools: Graph Studio, Oracle Machine Learning, Data Transforms	Oracle Autonomous Data Warehouse – ECPU – BYOL	B95703	ECPU per hour	Minimum ECPUs: 8 Minimum increment: 1

Resource Description	SKU Name	Part Number	Metric	Minimum and Minimum Increment
Database Storage usage	Oracle Autonomous Database Storage	B95754	Gigabyte Storage Capacity Per Month	1 TB (1024 GB)
Backup Storage usage	Oracle Autonomous Database Storage	B95754	Gigabyte Storage Capacity Per Month	1 GB Backups are billed per GB

The following table shows how to view cost and usage information for Autonomous Data Warehouse with the ECPU compute model.

See Cost and Usage Reports Overview for details on the cost and usage reports.

ECPU Cost and Usage Information for Autonomous Data Warehouse for Records Until January 31, 2025

Resource Description	Part Number	Cost Report	Usage Report
ECPU usage	B95701	Cost report: in the `product/Description` column, find, "Oracle Autonomous Data Warehouse – ECPU"	Usage report: in the `product/resource` column, find `PIC_ADW_ECPU_LI`
ECPU usage (Bring your own license)	B95703	Cost report: in the `product/Description` column, find, "Oracle Autonomous Data Warehouse – ECPU – BYOL"	Usage report: in the `product/resource` column, find `PIC_ADW_ECPU_BYOL`
ECPU usage for built-in-tools: Graph Studio, Oracle Machine Learning, Data Transforms	B95703	Cost report: in the `product/Description` column, find, "Oracle Autonomous Data Warehouse – ECPU – BYOL"	Usage report: in the `product/resource` column, find `PIC_ADBS_TOOLS_ECPU`
Database Storage usage	B95754	Cost report: in the `product/Description` column, find, "Oracle Autonomous Database Storage"	Usage report: in the `product/resource` column, find `PIC_ADB_STORAGE`
Backup Storage usage	B95754	Cost report: in the `product/Description` column, find, "Oracle Autonomous Database Backup Storage"	Usage report: in the `product/resource` column, find `PIC_ADB_STORAGE`

ECPU Cost and Usage Information for Autonomous Data Warehouse for Records After January 31, 2025

Resource Description	Part Number	Cost Report	Usage Information
ECPU usage	B95701	To find cost information, in the cost report find your OCID in the `product/resourceId` column and in the `product/Description` column, find, "Oracle Autonomous Data Warehouse – ECPU"	Usage is in the cost report in the `usage/billedQuantity` column.
ECPU usage (Bring your own license)	B95703	To find cost information, in the cost report find your OCID in the `product/resourceId` column and in the `product/Description` column, find, "Oracle Autonomous Data Warehouse – ECPU – BYOL"	Usage is in the cost report in the `usage/billedQuantity` column.

Resource Description	Part Number	Cost Report	Usage Information
ECPU usage for built-in-tools: Graph Studio, Oracle Machine Learning, Data Transforms	B95703	To find cost information, in the cost report find your OCID in the `product/resourceId` column and in the `product/Description` column, find, "Oracle Autonomous Data Warehouse – ECPU – BYOL"	Usage is in the cost report in the `usage/billedQuantity` column.
Database Storage usage	B95754	To find cost information, in the cost report find your OCID in the `product/resourceId` column and in the `product/Description` column, find, "Oracle Autonomous Database Storage"	Usage is in the cost report in the `usage/billedQuantity` column.
Backup Storage usage	B95754	To find cost information, in the cost report find your OCID in the `product/resourceId` column and in the `product/Description` column, find, "Oracle Autonomous Database Backup Storage"	Usage is in the cost report in the `usage/billedQuantity` column.

Billing Information: Autonomous JSON Database ECPU Compute Model

Shows billing information for Autonomous JSON Database with the ECPU compute model.

Resource Description	SKU Name	Part Number	Metric	Minimum and Minimum Increment
ECPU usage	Autonomous JSON Database – ECPU	B99708	ECPU per hour	Minimum ECPUs: 2 Minimum increment: 1
ECPU usage for built-in-tools: Graph Studio, Oracle Machine Learning, Data Transforms	Autonomous JSON Database – ECPU	B99708	ECPU per hour	Minimum ECPUs: 8 Minimum increment: 1
Database Storage usage	Oracle Autonomous Database Storage for Transaction Processing	B95706	Gigabyte Storage Capacity Per Month	1 GB Minimum storage size is 20 GB
Backup Storage usage	Oracle Autonomous Database Storage	B95754	Gigabyte Storage Capacity Per Month	1 GB Backups are billed per GB

The following table shows how to view cost and usage information for Autonomous JSON Database with the ECPU compute model.

See Cost and Usage Reports Overview for details on the cost and usage reports.

ECPU Cost and Usage Information for Autonomous JSON Database for Records Until January 31, 2025

Resource Description	Part Number	Cost Report	Usage Report
ECPU usage	B99708	Cost report: in the `product/Description` column, find, "Autonomous JSON Database – ECPU"	Usage report: in the `product/resource` column, find `PIC_AJD_ECPU_LI`

Resource Description	Part Number	Cost Report	Usage Report
ECPU usage for built-in-tools: Graph Studio, Oracle Machine Learning, Data Transforms	B99708	Cost report: in the `product/Description` column, find, "Autonomous JSON Database – ECPU"	Usage report: in the `product/resource` column, find `PIC_AJD_ECPU_LI`
Database Storage usage	B95706	Cost report: in the `product/Description` column, find, "Oracle Autonomous Database Storage for Transaction Processing"	Usage report: in the `product/resource` column, find `PIC_ADB_STORAGE_ATP`
Backup Storage usage	B95754	Cost report: in the `product/Description` column, find, "Oracle Autonomous Database Backup Storage"	Usage report: in the `product/resource` column, find `PIC_ADB_STORAGE`

ECPU Cost and Usage Information for Autonomous JSON Database for Records After January 31, 2025

Resource Description	Part Number	Cost Report	Usage Information
ECPU usage	B99708	To find cost information, in the cost report find your OCID in the `product/resourceId` column, and in the `product/Description` column, find, "Autonomous JSON Database – ECPU"	Usage is in the cost report in the `usage/billedQuantity` column.
ECPU usage for built-in-tools: Graph Studio, Oracle Machine Learning, Data Transforms	B99708	To find cost information, in the cost report find your OCID in the `product/resourceId` column and in the `product/Description` column, find, "Autonomous JSON Database – ECPU"	Usage is in the cost report in the `usage/billedQuantity` column.
Database Storage usage	B95706	To find cost information, in the cost report find your OCID in the `product/resourceId` and in the `product/Description` column, find, "Oracle Autonomous Database Storage for Transaction Processing"	Usage is in the cost report in the `usage/billedQuantity` column.
Backup Storage usage	B95754	To find cost information, in the cost report find your OCID in the `product/resourceId` and in the `product/Description` column, find, "Oracle Autonomous Database Backup Storage"	Usage is in the cost report in the `usage/billedQuantity` column.

Billing Information: APEX Service ECPU Compute Model

Shows billing information for APEX Service with the ECPU compute model.

Resource Description	SKU Name	Part Number	Metric	Minimum and Minimum Increment
ECPU usage	APEX Service – ECPU	B99709	ECPU per hour	Minimum ECPUs: 2 Minimum Increment: 1

Resource Description	SKU Name	Part Number	Metric	Minimum and Minimum Increment
Database Storage usage	Oracle Autonomous Database Storage for Transaction Processing	B95706	Gigabyte Storage Capacity Per Month	1 GB Minimum storage size is 20 GB
Backup Storage usage	Oracle Autonomous Database Storage	B95754	Gigabyte Storage Capacity Per Month	1 GB Backups are billed per GB

The following table shows how to view cost and usage information for APEX Service with the ECPU compute model.

See Cost and Usage Reports Overview for details on the cost and usage reports.

ECPU Cost and Usage Information for APEX Service for Records Until January 31, 2025

Resource Description	Part Number	Cost Report	Usage Report
ECPU usage	B99709	Cost report: in the `product/Description` column, find, "APEX Service – ECPU"	Usage report: in the `product/resource` column, find `PIC_APEX_ECPU_LI`
Database Storage usage	B95706	Cost report: in the `product/Description` column, find, "Oracle Autonomous Database Storage for Transaction Processing"	Usage report: in the `product/resource` column, find `PIC_ADB_STORAGE_ATP`
Backup Storage usage	B95754	Cost report: in the `product/Description` column, find, "Oracle Autonomous Database Backup Storage"	Usage report: in the `product/resource` column, find `PIC_ADB_STORAGE`

ECPU Cost and Usage Information for APEX Service for Records After January 31, 2025

Resource Description	Part Number	Cost Report	Usage Information
ECPU usage	B99709	To find cost information, in the cost report find your OCID in the `product/resourceId` column and in the `product/Description` column, find, "APEX Service – ECPU"	Usage is in the cost report in the `usage/billedQuantity` column.
Database Storage usage	B95706	To find cost information, in the cost report find your OCID in the `product/resourceId` column and in the `product/Description` column, find, "Oracle Autonomous Database Storage for Transaction Processing"	Usage is in the cost report in the `usage/billedQuantity` column.
Backup Storage usage	B95754	To find cost information, in the cost report find your OCID in the `product/resourceId` column and in the `product/Description` column, find, "Oracle Autonomous Database Backup Storage"	Usage is in the cost report in the `usage/billedQuantity` column.

Billing Information: Autonomous Database for Developers

Shows billing information for Autonomous Database for Developers with the ECPU compute model.

Resource Description	SKU Name	Part Number	Metric	Min. Increments Allowed
Instance usage (includes ECPU and storage)	Oracle Autonomous Database - Developer	B110316	Instance per hour	Fixed 4 ECPUs and 20 GB storage included

The following table shows cost and usage reports information for Autonomous Database for Developers with the ECPU compute model.

See Cost and Usage Reports Overview for details on the cost and usage reports.

Billing Information for Autonomous Database for Developers for Records Until January 31, 2025

Resource Description	Part Number	Cost Report	Usage Report
Instance usage (includes ECPU and storage)	B110316	Cost report: in the `product/Description` column, find, "Oracle Autonomous Database – Developer"	Usage report: in the`product/resource` column, find `PIC_ADB_DEVELOPER_LI` in the `product/Description` column.

Billing Information for Autonomous Database for Developers for Records After January 31, 2025

Resource Description	Part Number	Cost Reports	Usage Information
Instance usage (includes ECPU and storage)	B110316	Cost report: in the `product/Description` column, find, "Oracle Autonomous Database – Developer"	Usage is in the cost report in the `usage/billedQuantity` column.

OCPU Compute Model Billing Information

Provides billing information for Autonomous Transaction Processing, Autonomous Data Warehouse, Autonomous JSON Database, and APEX Service with the OCPU compute model.

OCPU is a legacy billing metric and has been retired on Autonomous Database. ECPUs are the replacement billing metric for all new and existing Autonomous Database deployments. See Oracle Support Document 2998742.1 for more information.

> **Note:**
>
> You can use OCI usage reports and `OCI_USAGE_DATA` view for information on usage *before* February 1, 2025. The OCI usage reports are deprecated. You can continue to access your existing usage report CSV files until July 31, 2025. See Cost and Usage Reports for more information.

- Billing Information: Autonomous Transaction Processing OCPU Compute Model
  Shows billing information for Autonomous Transaction Processing with the OCPU compute model.
- Billing Information: Autonomous Data Warehouse OCPU Compute Model
  Shows billing information for Autonomous Data Warehouse with the OCPU compute model.
- Billing Information: Autonomous JSON Database OCPU Compute Model
  Shows billing information for Autonomous JSON Database with the OCPU compute model.
- Billing Information: APEX Service OCPU Compute Model
  Shows billing information for APEX Service with the OCPU compute model.

Billing Information: Autonomous Transaction Processing OCPU Compute Model

Shows billing information for Autonomous Transaction Processing with the OCPU compute model.

Resource Description	SKU Name	Part Number	Metric	Min. Increments Allowed	Cost and Usage Reports
OCPU usage	Oracle Autonomous Transaction Processing	B90453	OCPU per hour	1 OCPU	Cost report: In the `product/Description` column, find "Oracle Autonomous Transaction Processing" Usage report: in the `product/resource` column, find `PIC_ATP_COMPUTE`
OCPU usage (Bring your own license)	Oracle Autonomous Transaction Processing – BYOL	B90454	OCPU per hour	1 OCPU	Cost report: Cost report: in the `product/Description` column, find, "Oracle Autonomous Transaction Processing – BYOL" Usage report: in the `product/resourceID` column, find `PIC_ATP_COMPUTE_BYOL`

Resource Description	SKU Name	Part Number	Metric	Min. Increments Allowed	Cost and Usage Reports
Database Storage usage	Oracle Autonomous Transaction Processing - Exadata Storage	B90455	Terabyte Storage Capacity Per Month	1 TB	Cost report: in the `product/Description` column, find, "Oracle Autonomous Transaction Processing - Exadata Storage" Usage report: in the `product/resource` column, find `PIC_ATP_EXADATA_STORAGE`

See Cost and Usage Reports Overview for details on the cost and usage reports.

Billing Information: Autonomous Data Warehouse OCPU Compute Model

Shows billing information for Autonomous Data Warehouse with the OCPU compute model.

Resource Description	SKU Name	Part Number	Metric	Min. Increments Allowed	Cost and Usage Reports
OCPU usage	Oracle Autonomous Data Warehouse	B89040	OCPU per hour	1 OCPU	Cost report: in the `product/Description` column, find "Oracle Autonomous Data Warehouse" Usage report: in the `product/resource` column, find `PIC_ADWC_COMPUTE`
OCPU usage (Bring your own license)	Oracle Autonomous Data Warehouse – BYOL	B89039	OCPU per hour	1 OCPU	Cost report: in the `product/Description` column, find "Oracle Autonomous Data Warehouse – BYOL" Usage report: in the `product/resource` column, find `PIC_ADWC_COMPUTE_BYOL`

Resource Description	SKU Name	Part Number	Metric	Min. Increments Allowed	Cost and Usage Reports
Database Storage usage	Oracle Autonomous Data Warehouse - Exadata Storage	B89041	Terabyte Storage Capacity Per Month	1 TB	Cost report: in the `product/Description` column, find "Oracle Autonomous Data Warehouse - Exadata Storage" Usage report: in the `product/resource` column, find `PIC_ADWC_EXADATA_STORAGE`

See Cost and Usage Reports Overview for details on the cost and usage reports.

Billing Information: Autonomous JSON Database OCPU Compute Model

Shows billing information for Autonomous JSON Database with the OCPU compute model.

Resource Description	SKU Name	Part Number	Metric	Min. Increments Allowed	Cost and Usage Reports
OCPU usage	Oracle Autonomous JSON Database	B92212	OCPU per hour	1 OCPU	Cost report: in the `product/Description` column, find "Oracle Autonomous JSON Database" Usage report: in the `product/resource` column, find `PIC_AJD_COMPUTE`
Database Storage usage	Oracle Autonomous Transaction Processing - Exadata Storage	B90453	Terabyte Storage Capacity Per Month	1 TB	Cost report: in the `product/Description` column, find "Oracle Autonomous JSON Database - Exadata Storage" Usage report: in the `product/resource` column, find `PIC_ATP_EXADATA_STORAGE`

> **Note:**
>
> In the OCPU model, 60 days of automatic backups are included in the cost of database storage and have no additional charge. Long-term backups are billed additionally, as database storage (Part Number: B90453).

See Cost and Usage Reports Overview for details on the cost and usage reports.

Billing Information: APEX Service OCPU Compute Model

Shows billing information for APEX Service with the OCPU compute model.

Resource Description	SKU Name	Part Number	Metric	Min. Increments Allowed	Cost and Usage Reports
OCPU usage	Oracle APEX Application Development	B92911	OCPU per hour	1 OCPU	Cost report: In the `product/Description` column, find "Oracle APEX Application Development" Usage report: in the `product/resource` column, find `PIC_ADBS_APEX_LI`
Database Storage usage	Oracle Autonomous Transaction Processing - Exadata Storage	B90453	Terabyte Storage Capacity Per Month	1 TB	Cost report: in the `product/Description` column, find, "Oracle Autonomous Transaction Processing - Exadata Storage" Usage report: in the `product/resource` column, find `PIC_ATP_EXADATA_STORAGE`

> **Note:**
>
> In the OCPU model, 60 days of automatic backups are included in the cost of database storage and have no additional charge. Long-term backups are billed additionally, as database storage (Part Number: B90453).

See Cost and Usage Reports Overview for details on the cost and usage reports.

Oracle Autonomous Database Serverless Features Billing

Shows billing information for Autonomous Database features for ECPU and OCPU compute models.

- ECPU Compute Model
- OCPU Compute Model

ECPU Compute Model

- **Automatic Backups**: The storage for backups is billed, per GB, in addition to your selected database storage.

  For example, if backups are occupying 200 GB of storage you are billed for 200 GB of backup storage (in addition to the usage billed for your selected number of ECPUs and database storage). See ECPU Compute Model Billing Information for more detail about which SKUs are billed for backups.

- **Long-Term Backups**: The storage for long-term backups is billed, per GB, as backup storage, in addition to your database storage.

  For example, if automatic backups are currently occupying 200 GB and long-term backups are occupying 600 GB of storage, you are billed for 800 GB of backup storage, in addition to the usage billed for your selected ECPUs and database storage. See ECPU Compute Model Billing Information for more detail about which SKUs are billed for each workload type and backups.

- **Compute Auto-Scaling**: When compute auto scaling is enabled, your database may use and you may be billed for additional ECPU consumption as needed by your workload, up to three times (3x) the number of base ECPUs, as shown in the **ECPU count** on the Oracle Cloud Infrastructure Console.

  – Your billed ECPU usage per hour while your database is running is based on the base number of ECPUs you selected for your database, plus any additional ECPU usage due to auto-scaling.

  – A stopped Autonomous Database instance has zero ECPU usage.

  – ECPU usage is measured each second, in units of whole ECPUs and averaged across an hour. If your database is running for less than an hour or auto-scales for only part of an hour, it is billed per second for the average ECPU consumption, over the base ECPUs, during that hour. The minimum ECPU consumption is one minute.

  For example, if your database has ECPU count 4 with **Compute auto scaling** enabled:

  – Assume in hour one your database is Available for the entire hour and ECPU utilization is below 4 ECPUs. Your database will be billed for 4 ECPUs.

  – Assume in hour two your database is Available for the entire hour and ECPU utilization is below 4 ECPUs for 30 minutes, 50% of the hour, and auto scales to 8 ECPUs for 30 minutes (the other 50% of the hour). The usage for this period, for billing, is 6 ECPUs (based on the average per-second ECPU usage for hour two).

- **Storage Auto-Scaling**:

  – For storage usage below your reserved base storage, you are billed based on your base storage.

- After your allocated storage exceeds your reserved base storage, storage usage is billed based on your allocated storage rounded up to the nearest TB, in a given hour.

For example, if your reserved base storage is 4 TB, until your allocated storage exceeds 4 TB of storage, you are billed based on your base storage (4 TB). After you exceed 4 TB, storage is billed based on the allocated storage rounded up to the nearest TB, in a given hour. In this example, if the allocated storage grows over 4 TB in a given hour, say to 4.9 TB, you are billed for 5 TB of storage from that hour onward.

If you then delete 1 TB of data, your allocated storage remains at 4.9 TB and you are billed for 5 TB until you perform a shrink operation. When you perform a shrink operation, may be able to shrink /reduce your allocated storage back to 3.9 TB. After the shrink operation completes and your allocated storage (3.9 TB) is once again below your reserved base storage (4 TB), you will once again be billed for your reserved base storage of 4 TB. See Shrink Storage for more information.

- **Autonomous Data Guard Standby – Local (Same-Region)**

  Local Autonomous Data Guard peer databases incur the additional cost of the base ECPUs and the storage of the Primary database, including any auto-scaled storage usage, billed on the Primary database itself. Auto-scaled ECPUs of the Primary database are not billed for additionally on the local Autonomous Data Guard peer database. The number of base ECPUs is specified by the number of ECPUs as shown in the **ECPU count** on the Oracle Cloud Infrastructure Console.

  For example, if you enable a local Autonomous Data Guard peer on a source database with:

  - 2 (base) ECPUs with **Compute auto-scaling** enabled and are consuming about 4 ECPUs per hour
  - 1 TB of (base) storage with storage auto-scaling, consuming a total of 2 TB of database storage

  For the local Autonomous Data Guard peer you are billed for an additional 2 ECPUs (your base ECPU selection), plus an additional 2 TB of storage (that is, the same amount of storage reserved for the source Primary with auto-scale, on the Primary database).

  When the Primary database is stopped, neither the primary nor the peer database is billed for ECPU.

- **Autonomous Data Guard Standby – Remote (Cross-Region)**

  Autonomous Data Guard cross-region peer databases incur the additional cost of the base ECPUs and twice (2x) the storage of the Primary database, including any auto-scaled storage usage, billed on the remote peer database. Auto-scaled ECPUs of the Primary are not billed for additionally on the remote peer database. The number of base ECPUs is specified by the number of ECPUs as shown in the **ECPU count** on the Oracle Cloud Infrastructure Console.

  For example, if you enable a cross-region Autonomous Data Guard peer for a source database with:

  - 2 (base) ECPUs with **Compute auto-scaling** enabled and are consuming about 4 ECPUs per hour
  - 1 TB of (base) storage with storage auto-scaling, consuming a total of 2 TB of database storage

  For the cross-region Autonomous Data Guard peer you are billed an additional 2 ECPUs (your base ECPU selection), plus 4 TB of storage (that is 2x the storage reserved for the source Primary with auto-scaling, billed on the remote peer database).

When the Primary database is stopped, neither the primary nor the peer database are billed for ECPU.

When the option **Enable cross-region backup replication to disaster recovery peer** is selected, you are billed for twice (2x) the backup storage size required for the replicated backups, billed to the remote standby.

When a cross-region peer operates as a Snapshot standby, the Snapshot standby CPU usage is billed based on the base CPU count and any additional CPU usage if compute auto scaling is enabled. The number of base CPUs is specified by the number of ECPUs as shown in the **ECPU count** field on the Oracle Cloud Infrastructure Console.

- **Backup-Based Disaster Recovery – Local (Same-Region) Backup Copies** There are no additional costs for local Backup-Based Disaster Recovery, other than the cost of storage for automatic backups.
- **Backup-Based Disaster Recovery – Remote (Cross-Region) Backup Copies** Billing for a cross-region Backup-Based Disaster Recovery is twice (2x) the amount of backup storage required for the replicated cross-region backups, billed to the remote peer.

  For example, if you enable a cross-region backup copy on a source database with:

  – 2 (base) ECPUs

  – 2 TB of database storage

  If the backups replicated to the remote region take up 1.9 TB of storage, you will be billed for 3.8 TB of backup storage on the remote backup copy peer database.

  When the option **Enable cross-region backup replication to disaster recovery peer** is selected, you are billed for twice (2x) the amount of backup storage size required for the additional replicated backups, billed to the remote peer. This billing is based on the number of days set for the backup retention on the Primary, as follows:

  – When automatic backup retention is set to 7 days or greater, billing is based on the storage size for 7 days of replicated backups.

  – When automatic backup retention is set to less than 7 days, billing is based on the storage size for the specified number of days of data that is replicated to the cross-region standby.
- **Refreshable Clone Local (Same-Region)** Local refreshable clones have their own configurable ECPU selection, and so they are billed for ECPU based on the user-selected number of ECPUs, with or without auto scaling; they do not get billed additionally over the ECPU selection. The number of ECPUs is specified by the number of ECPUs as shown in the **ECPU count** on the Oracle Cloud Infrastructure Console.

  Local refreshable clones are billed for the same amount of storage as their source database.

  For example, if you create a 2 ECPU local refreshable clone from a source database with:

  – 4 ECPUs

  – 1 TB of storage with storage auto-scaling and consuming 2 TB of storage

  For the local refreshable clone you will be billed for 2 ECPUs, that is, the refreshable clone's **ECPU count** value, and 2 TB of storage (that is, the storage reserved for the source database).

  When you start or stop the source database, your actions on the source database do not affect the refreshable clone. A refreshable clone is started or stopped independently of the source database.
- **Refreshable Clone Remote (Cross-Region)** Remote refreshable clones have their own configurable ECPU selection, and so they are billed for ECPU based on the user-selected

ECPU (with or without auto scaling); they do not get billed additionally over the ECPU selection. The number of ECPUs is specified by the number of ECPUs as shown in the **ECPU count** on the Oracle Cloud Infrastructure Console.

Remote refreshable clones are billed for twice (2x) the amount of storage as their source database.

For example, if you create a 2 ECPU remote refreshable clone from a source database with:

- 4 ECPUs
- 1 TB of storage with storage auto-scaling and consuming 2 TB of storage

For the remote refreshable clone you are billed for 2 ECPUs (that is, the refreshable clone's ECPU selection) and 4 TB of storage (that is, 2x the storage reserved for the source database)

Starting or stopping the source database does not affect the refreshable clone - The refreshable clone can be started or stopped independently.

- **Snapshot Standby for Remote (Cross-Region) Disaster Recovery**

  Snapshot standby ECPU usage is billed based on the base ECPU count and any additional ECPU usage if compute auto-scaling is enabled. The number of base ECPUs is specified by the number of ECPUs, as shown in the **ECPU count** on the Oracle Cloud Infrastructure Console.

  Snapshot standby storage usage is billed based on the storage of the snapshot standby plus (1x) the storage of the source primary database.

  For example, if you have a 2 ECPU and 3 TB snapshot standby from a source database with:

  - 4 ECPUs
  - 1 TB of storage with storage auto-scaling and consuming 2 TB of storage

  Your snapshot standby will be billed for 2 ECPUs (that is, the snapshot standby's ECPU selection) and 3 TB + 2 TB = 5 TB of database storage (that is, the storage reserved on the snapshot standby + the storage reserved on its source database)

- **Elastic Pools**: An elastic pool allows you to provision 4 times the ECPU count that you have as your pool size. For example, if you have a pool with pool size of 128 ECPUs, you can provision up to 512 ECPUs in this pool. In other words, when your pool size is 128 ECPUs, your pool capacity is four times the pool size (in this example, 512 ECPUs).

  Databases that belong to a pool are not billed individually for compute. The compute billing for all the pool members and the leader happens through the leader. In other words, individual members of an elastic pool do not get billed for compute as long as they are part of the pool. This applies no matter what the workload type is for the pool member. For example, when a pool member with workload type **Data Warehouse** is added to a pool, it's compute usage is billed to the pool leader at the **Transaction Processing** compute usage rate. The storage billing, on the other hand, continues to be billed to individual Autonomous Database instances irrespective of them being part of a pool or not.

  Let’s assume you have an elastic pool with pool size of 128 ECPUs. Given the pool size, the pool capacity is 512 ECPUs (pool capacity = 4x pool size). For this sample, here are some common billing questions and answers:

  - What is the maximum number of Autonomous Database instances allowed in this pool? A total of 512 Autonomous Database instances with 1 ECPU each (elastic pool members or the leader can have an individual ECPU allocation as low as 1 ECPU).

  - What happens if the aggregated ECPU utilization high watermark of the pool is greater than the pool size? If the aggregated ECPU utilization high watermark is less than or equal to the pool size in a given billing hour, the hourly charge is in the amount of the pool size. If the aggregated ECPU utilization high watermark is more than the pool size but less than or equal to the 2x pool size in a given billing hour, the hourly charge is in the amount of 2x pool size. If the aggregated ECPU utilization high watermark is more than 2x pool size in a given billing hour, the hourly charge is in the amount of 4x pool size.

    For example, let's assume 512 Autonomous Database instances with 1 ECPU each are in an elastic pool with a pool size of 128 ECPUs. If the aggregated ECPU utilization high watermark of these databases is 100 ECPUs between 1pm-2pm, and 250 ECPUs between 2pm-3pm; then the billing is 128 ECPU hours between 1pm-2pm and 256 ECPU hours between 2pm-3pm.

  See About Elastic Pool Billing for more information.

- **Undelete an Autonomous Database Instance**

  If you undelete an Autonomous Database instance, in the first hour after the undelete operation you are billed for the base CPU and Storage, including database storage and long term backups, for the total time when the database was deleted, as if the database was never deleted and was running.

  For example, if you terminate an Autonomous Database instance with:

  - 4 ECPU with **Compute auto scaling** enabled
  - 2 TB Storage, with 100 GB of automatic backup storage and 20 GB of long-term backup storage

  If after 5 hours and 30 minutes you undelete the terminated instance, in the first hour after the undelete operation the billing includes the additional costs as if the database was never deleted and was running, including:

  - 5 hours and 30 minutes for the base 4 ECPUs
  - 2 TB Storage
  - 100 GB of automatic backup storage
  - 20 GB of long-term backup storage

OCPU Compute Model

- **Automatic Backups**: The storage for automatic backups is included in the cost of database storage. See OCPU Compute Model Billing Information for more detail about which SKUs for database storage.
- **Long-Term Backups**: The storage for long-term backups is billed per TB as additional database storage, in addition to your selected database storage usage.

  For example, if automatic backups occupy 200 GB and long-term backups occupy 600 GB of storage, you will be billed for 1 TB (600 GB of long-term backup storage rounded up to the nearest TB) as database storage, in addition to the usage billed for your selected OCPU and database storage. See OCPU Compute Model Billing Information for more detail about which SKUs are billed for each workload type and backups.
- **Compute auto-scaling**: When compute auto scaling is enabled your database may use and you may be billed for additional OCPU consumption as needed by your workload, up to three times (3x) the number of base OCPUs (as shown in **OCPU count** on the Oracle Cloud Infrastructure Console)

  - Your billed OCPU usage per hour while your database is running is based on the base number of OCPUs you selected for your database, plus any additional OCPU usage due to auto-scaling.
  - A stopped Autonomous Database instance has zero OCPU usage.
  - OCPU usage is measured each second, in units of whole OCPUs and averaged across an hour. If your database is running for less than an hour or auto-scales for only part of an hour, it will be billed per second for the average OCPU consumption (over the base OCPU) during that hour. The minimum OCPU consumption is one minute.

- **Storage Auto-Scaling**:
  - For storage usage below your reserved base storage, you are billed based on your base storage.
  - After your allocated storage exceeds your reserved base storage, storage usage is billed based on your allocated storage rounded up to the nearest TB, in a given hour.

  For example, if your reserved base storage is 4 TB, until your allocated storage exceeds 4 TB of storage, you are billed based on your base storage (4 TB). After you exceed 4 TB, storage is billed based on the allocated storage rounded up to the nearest TB, in a given hour. In this example, if the allocated storage grows over 4 TB in a given hour, say to 4.9 TB, you are billed for 5 TB of storage from that hour onward.

  If you then delete 1 TB of data, your allocated storage remains at 4.9 TB and you are billed for 5 TB until you perform a shrink operation. When you perform a shrink operation, Autonomous Database may be able to shrink / reduce your allocated storage back to 3.9 TB. After the shrink operation completes and your allocated storage (3.9TB) is once again below your reserved base storage (4 TB), you will once again be billed for your reserved base storage of 4 TB. See Shrink Storage for more information.

- **Autonomous Data Guard Standby Local (same-region)**

  Local Autonomous Data Guard peer databases incur the additional cost of the base OCPUs and the storage of the Primary database, including any auto-scaled storage usage, billed on the Primary database itself. Auto-scaled OCPUs of the Primary database are not billed for additionally on the local Autonomous Data Guard peer database. The number of base OCPUs is specified by the number of OCPUs, as shown in the **OCPU count** on the Oracle Cloud Infrastructure Console.

  When the Primary database is stopped, neither the primary nor the peer database is billed for OCPU.

- **Autonomous Data Guard Standby Remote (cross-region)**

  Autonomous Data Guard cross-region standby databases incur the additional cost of the base OCPUs and twice (2x) the storage of the Primary database, including any auto-scaled storage usage, billed on the remote peer database. Auto-scaled OCPUs of the Primary are not billed for additionally on the remote peer database. The number of base OCPUs is specified by the number of OCPUs, as shown in **OCPU count** on the Oracle Cloud Infrastructure Console.

  When the Primary database is stopped, neither the primary nor the peer database is billed for OCPU.

  When the option **Enable cross-region backup replication to disaster recovery peer** is selected, you are billed to the remote standby database's OCPU database storage, for twice (2x) the 7 day replicated backup storage size, rounded up to the nearest TB.

- **Backup-based Disaster Recovery Local (same-region) Backup Copies** There are no additional costs for local Backup-Based Disaster Recovery, other than the cost of keeping automatic backups.

- **Backup-Based Disaster Recovery Remote (cross-region) Backup Copies**

  Billing for a cross-region Backup-Based Disaster Recovery with OCPUs is twice (2x) the amount of storage required for backups replicated to the remote region, billed as database storage to the remote peer, rounded up to the nearest TB.

  When the option **Enable cross-region backup replication to disaster recovery peer** is selected, you are billed to the remote peer database's OCPU database storage, for twice (2x) the replicated backup storage size, rounded up to the nearest TB.

- **Refreshable Clone Local (same-region)**

  Local refreshable clones have their own configurable OCPU selection, and so they are billed for OCPU based on the user-selected OCPU (with or without auto scaling); they do not get billed additionally over the OCPU selection. The number of OCPUs is specified by the number OCPUs as shown in the **OCPU count** on the Oracle Cloud Infrastructure Console.

  Local refreshable clones are billed for the same amount of storage as their source database.

  Starting or stopping the source database does not affect the refreshable clone. The refreshable clone can be started or stopped independently.

- **Refreshable Clone Remote (cross-region)**

  Remote refreshable clones have their own configurable OCPU selection, and so they are billed for OCPU based on the user-selected OCPU (with or without auto scaling); they do not get billed additionally over the OCPU selection. The number of OCPUs is specified by the number of OCPUs as shown in the **OCPU count** on the Oracle Cloud Infrastructure Console.

  Remote refreshable clones are billed for twice (2x) the amount of storage as their source database.

  Starting or stopping the source database does not affect the refreshable clone. The refreshable clone can be started or stopped independently.

- **Snapshot Standby for Remote (cross-region) disaster recovery**

  Snapshot standby OCPU usage is billed based on the base OCPU count and any additional OCPU usage if compute auto-scaling is enabled. The number of base OCPUs is specified by the OCPUs as shown in the **OCPU count** on the Oracle Cloud Infrastructure Console.

  Snapshot standby storage usage is billed based on the storage of the snapshot standby plus (1x) the storage of the source primary database.

Oracle Autonomous Database Serverless Billing for Database Tools

Shows billing information for Autonomous Database database tools for ECPU and OCPU compute models.

- ECPU Compute Model
- OCPU Compute Model

ECPU Compute Model

The following built-in tools have separate, configurable compute resource limits:

- Graph Studio
- Oracle Machine Learning
- Data Transforms

For these tools compute resource usage does not use the database's base compute, as specified with the **ECPU count**, or any auto-scaled ECPUs. These tools have separate ECPU allocation, independent from the database compute **ECPU count** (and independent of **Compute auto scaling**).

For Graph Studio, Oracle Machine Learning, and Data Transforms, the ECPU count and Maximum idle time configuration options let you specify compute resource allocation for the VMs that run the associated built-in tool. When you configure these built-in tools, the values for these configuration options mean the following in terms of resource usage and billing:

- You do not pay for a tool's ECPU allocation if you do not use the tool.
- The VMs associated with a tool are provisioned when you start using a tool. For example, if Graph Studio is disabled, billing does not begin when you enable the tool. Billing begins when you start using Graph Studio.
- The ECPU count specifies the compute resource allocation that is dedicated to the tool. The value you enter for ECPU count applies in addition to the database ECPU count you specify for the instance.

  After you start using a tool you are billed per ECPU hour reserved from the time the built-in Tool is launched. ECPU usage is measured each second, in units of whole ECPUs and averaged across an hour. If your database is running for less than an hour or you change the number of ECPUs for only part of an hour, your built-in Tool ECPU usage is billed per second for the average ECPU consumption during that hour. The minimum ECPU consumption is one minute.
- Billing stops for a built-in tool's allocated ECPUs when the built-in tool is disabled, the instance stops or is terminated, or when the built-in tool is idle for more than the specified Maximum idle time. The default maximum idle times depend on the tool:

For example, assume you allocate 4 ECPUs as the base ECPU count for the Autonomous Database instance, with compute auto-scaling disabled, and you use 8 ECPUs for Graph Studio with the maximum idle time set to 30 minutes. If you use Graph Studio for 30 minutes starting at 2 pm, your bill for the specified time period (2-3pm) is a total of 12 ECPUs. This sum includes the 4 base ECPUs plus the additional tools usage of 8 ECPUs for Graph Studio (including 30 minutes of actual usage and 30 minutes of idle time).

As another example, assume you allocate 6 ECPUs as the base ECPU count for the Autonomous Database instance, with compute auto-scaling disabled, and you use the default 8 ECPUs for Oracle Machine Learning (OML) with the maximum idle time set to 15 minutes. When the database is open and available for the hour between 2 to 3 pm, and you use OML for 15 min starting at 2 pm, with the specified 15 minute maximum idle time, you are charged for 8 ECPUs for OML compute usage for 30 minutes plus 6 ECPUs for the database compute resources for the full hour, for a total of 14 ECPUs assuming no other tools are used in this period.

Note: After you start using a tool with a specified **ECPU count** you are billed per ECPU hour reserved from the time the built-in Tool is launched. ECPU usage is measured each second, in units of whole ECPUs and averaged across an hour. If your database is running for less than an hour or you change the number of ECPUs for only part of an hour, it is billed per second for

the average ECPU consumption during that hour. The minimum ECPU consumption is one minute.

Depending on the workload type of the Autonomous Database instance, built-in tool usage is charged to the following SKUs:

- Oracle Autonomous Transaction Processing – ECPU – BYOL
- Oracle Autonomous Data Warehouse – ECPU – BYOL
- Autonomous JSON Database – ECPU

See ECPU Compute Model Billing Information for details on built-in tools billing.

OCPU Compute Model

Built-in tools do not have separate, configurable compute resource limits. The built-in tools use compute from the database's selected OCPU count. Tool usage may also use auto-scaled OCPUs, if necessary.

Autonomous Database for Developers Billing and Tenancy Service Limit

You are billed for Autonomous Database for Developers usage hourly, on a per instance basis.

Autonomous Database for Developers databases are billed as follows:

- Each instance is billed at a flat rate, per hour.
- An instance cannot be stopped for billing. After an Autonomous Database for Developers instance is created, it is billed until the instance is terminated.
- An Autonomous Database for Developers instance may be stopped and started as required, but this does not affect billing.
- Terminating an Autonomous Database for Developers instance stops the billing for the instance.

See Billing Information: Autonomous Database for Developers for compute and storage billing details.

When you provision an instance you select a workload type based on the workloads you want to run on your database. The pricing is the same for each workload type.

See the Autonomous Database Cloud Price List for up-to-date pricing.

Use the Cloud Cost Estimator to estimate the monthly cost for an Autonomous Database for Developers instance.

Autonomous Database for Developers Service Limit

The number of Autonomous Database for Developers instances in a tenancy is limited based on the `adb-developer-count` service limit. The default limit is 100, and you may request a limit increase for more. See Subscribed Region Limits for more information.

Always Free Autonomous Database

You have the option to create a limited number of Always Free Autonomous Databases that do not consume cloud credits. Always Free databases can be created in Oracle Cloud Infrastructure accounts that are in a trial period, have paying status, or are always free.

Sign Up with Oracle Cloud Free Tier

These are the services you can use for an unlimited time:

- Two Oracle Autonomous Databases with powerful tools like Oracle APEX and Oracle SQL Developer
- Two Oracle Cloud Infrastructure Compute VMs; Block, Object, and Archive Storage; Load Balancer and data egress; Monitoring and Notifications

See Oracle Cloud Free Tier to start for free.

Resource Restrictions for Always Free Autonomous Database

- The Always Free option provides databases that have CPU and storage included and you are never billed for an Always Free instance until the instance is upgraded to a paid Autonomous Database.
- Maximum of approximately 20 GB Exadata storage per database (you may see more than this)
- Maximum of 30 simultaneous database sessions
- Maximum of 2 Always Free Autonomous Database instances per Oracle Cloud Infrastructure tenancy. The Always Free Autonomous Database workload types are: Data Warehouse, Transaction Processing, JSON Database, and APEX Service. If you create 2 Always Free instances, they can be the same or different Autonomous Database workload types.
- The HTTP interface for Always Free Autonomous Databases is rate limited to restrict the number of simultaneous service users. Approximately 3-6 simultaneous users can be supported across all of the APEX, Oracle REST Data Services, and Database Actions running on your Always Free Autonomous Databases. Additional simultaneous users beyond that may result in users encountering HTTP errors such as HTTP status code 429.

  This HTTP interface rate limit applies only for Always Free Autonomous Databases.

Notes:

- For details on Always Free Oracle APEX Application Development (APEX Service), see Always Free Oracle APEX Application Development.
- Always Free Autonomous Databases cannot be scaled manually or automatically beyond the fixed resource restrictions described above.
- The Maximum of 30 simultaneous database sessions limit for Always Free allows you to work with Autonomous Database; however, if your usage includes many simultaneous users and/or many concurrent database client connections then you can exceed these limits, resulting in errors. To avoid such errors, obtain more resources for your Autonomous Database by upgrading to paid service.
- Always Free Autonomous Databases cannot be provisioned as a private endpoint and cannot reside within a Virtual Cloud Network (VCN). See Configure Private Endpoints When You Provision or Clone an Instance for further information on private endpoints.

Oracle Database Versions for Always Free Autonomous Database

The available Database versions for Always Free Autonomous Database are:

- **Oracle Database 19c**

  If you have Oracle Database 19c, then many database concepts and features of this service are further documented here:

  – Oracle Database 19c

- **Oracle Database 23ai**

  If you are using Always Free Autonomous Database with Oracle Database 23ai, then many of the concepts and features of this service are further documented here:

  – Oracle Database 23ai

  – *Oracle Database New Features*

  Notes for Always Free Autonomous Database with Oracle Database 23ai:

  – Always Free Autonomous Database with Oracle Database 23ai is available in the following regions:

    * US West: Phoenix (PHX)
    * US East: Ashburn (IAD)
    * UK South: London (LHR)
    * France Central: Paris (CDG)
    * Australia East: Sydney (SYD)
    * India West: Mumbai (BOM)
    * Singapore (SIN)
    * Japan East: Tokyo (NRT)

  – To take advantage of new features available with Oracle Database 23ai, your tenancy must have a home region of one of these regions and you must provision your Always Free Autonomous Database in your home region.

See Autonomous Database with Oracle Database 23ai Notes for additional information and limitations on Always Free Autonomous Database with Oracle Database 23ai.

Regional Availability for Always Free Autonomous Database

- Always Free Autonomous Databases are available worldwide in a subset of Oracle Cloud Infrastructure data regions. See Data Regions for more details on where Always Free databases are supported.
- When you sign up for Oracle Cloud Infrastructure, Oracle creates a tenancy and designates a home data region for the tenancy that you specify. You can create Always Free Autonomous Databases only in this home data region. You cannot create an Always Free Autonomous Database in other data regions that you subsequently subscribe to. See The Home Region for more information.

Backup Functionality Not Available in Always Free Autonomous Database

- Always Free Autonomous Databases do not support long-term backups or manual backups to your Oracle Cloud Infrastructure object storage.
- Always Free Autonomous Databases do not support restore (you cannot restore the current database to the selected past backup or timestamp).
- To use backup and restore functionality on Autonomous Database you must upgrade an Always Free Autonomous Database to a paid instance.

See Backup and Restore Autonomous Database Instances and Upgrade an Always Free Autonomous Database to a Paid Instance for more information.

Autonomous Data Guard Not Available for Always Free Autonomous Database

Autonomous Data Guard is not available with Always Free Autonomous Databases. See Use Standby Databases with Autonomous Data Guard for Disaster Recovery for more information.

Tools Configuration Options Not Available for Always Free Autonomous Database

Always Free Autonomous Database does not provide configuration options for Autonomous Database tools and does not allow you to disable Autonomous Database tools. For example, you cannot specifically disable HTTP access to Oracle APEX and REST Data Services on Always Free databases.

Supplemental Logging and Oracle GoldenGate Extract are Not Available for Always Free Autonomous Database

The following are not available for Always Free Autonomous Databases:

- Supplemental logging
- Oracle GoldenGate Extract

See Supplemental Logging and Configuring Extract to Capture from an Autonomous Database for more information.

Inactivity Monitoring and Database Stoppage

Persistently inactive Always Free Autonomous Databases are detected and handled as follows:

- After being inactive for 7 days, the database will be stopped automatically, preserving its stored data. Inactivity measurements leading up to 7 days are based on database connections and CPU usage. Successfully making a SQL*Net or HTTPS connection and running SQL commands on your database resets these measurements to zero.

- A database that is automatically or manually stopped and stays inactive for 90 days, cumulative, may be reclaimed and permanently deleted. Inactivity measurements leading up to 90 days are based on the database being inactive or in the stopped state. Starting a stopped database resets these measurements to zero.

  Start an Always Free Autonomous Database by clicking the **Start** button on the Oracle Cloud Infrastructure console. Start a stopped Always Free Autonomous Database before 90 days to avoid losing access to its data.

  When you start an Always Free Autonomous Database from the stopped state, you need to wait about 5 minutes before attempting to connect to an APEX application or to an Oracle REST Data Services (ORDS) endpoint. If you attempt to connect before the background APEX and ORDS startup completes, then you may see HTTP error messages.
- On an Always Free database the Oracle Cloud Infrastructure console shows banner alerts prior to automatic stop and permanent delete operations occurring. If you subscribe to Oracle Cloud Infrastructure Alerts and Notifications, you also will receive email notifications.

> **Note:**
>
> After an Always Free Autonomous Database has been stopped and is later started, you may need to reconnect to the database from SQL*Net database clients. You can use the same Oracle Wallet and database user credentials to reconnect.

Topics

- Upgrade an Always Free Autonomous Database to a Paid Instance
  You can upgrade an Always Free Autonomous Database to a paid instance to provide the instance with additional compute and storage resources.

Upgrade an Always Free Autonomous Database to a Paid Instance

You can upgrade an Always Free Autonomous Database to a paid instance to provide the instance with additional compute and storage resources.

If your Oracle Cloud Infrastructure account is in a trial period or has paying status, you can upgrade the Always Free database to a paid instance as follows:

1. From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
2. On the Autonomous Databases page select an Always Free Autonomous Database from the links under the **Display name** column.
3. On the Autonomous Database Details page, from the **More actions** drop-down list select **Upgrade instance to paid**.
4. Click **Upgrade instance to paid**.

When you upgrade from an Always Free instance to a paid instance you get a paid instance of the same workload type with the minimum CPU and minimum storage available for that workload type. After you upgrade you can scale up CPU and storage resources to fit your needs.

For example:

- **Data Warehouse** workload type: When you upgrade to paid from an Always Free instance, the upgrade process provides a paid Autonomous Database instance with workload type **Data Warehouse**, with 2 ECPUs and 1 TB of database storage.
- **Transaction Processing** workload type: When you upgrade to paid from an Always Free instance, the upgrade process provides a paid Autonomous Database instance with workload type **Transaction Processing**, with 2 ECPUs and 1 TB of database storage.

Notes for upgrading an Always Free Autonomous Database to a paid instance:

- If have an Always Free Autonomous Database and you want an Autonomous Database for Developers instance you can clone the Always Free instance. To create the paid Autonomous Database for Developers instance, perform the clone operation and in the configure the database step, select the **Developer** option. See Clone an Autonomous Database Instance for more information.
- If your account has finished a trial without upgrading to paying status, you can continue using Always Free databases but you cannot upgrade Always Free instances to paid instances until the account is first upgraded to paying status. See Upgrade Your Free Oracle Cloud Promotion for more information.
- For details on upgrading your Always Free APEX Service to an Oracle APEX Application Development paid instance, see Upgrading Always Free APEX Service to a Paid Version.

Autonomous Database for Developers

Autonomous Database for Developers provides you with low cost instances for developers and others to build and test new Autonomous Database applications.

Billing for Autonomous Database for Developers databases is hourly per instance. See Autonomous Database for Developers Billing and Tenancy Service Limit for details.

Autonomous Database for Developers are low-cost, fixed shape databases intended for development and testing uses cases, and are not recommended for production use cases. When you need more compute or storage resources or you want to take advantage of additional Autonomous Database features, you can upgrade your Autonomous Database for Developers instance to a full paid service instance.

Topics

- Resource Restrictions with Autonomous Database for Developers
  Autonomous Database for Developers instances come with compute and storage resources included.
- Oracle Database Versions with Autonomous Database for Developers
  Shows the available Oracle Database versions for Autonomous Database for Developers.
- Supported Features with Autonomous Database for Developers
  Autonomous Database for Developers supports most but not all of the Autonomous Database features.
- Provisioning and Cloning Workflow with Autonomous Database for Developers
  You can provision or clone Autonomous Database for Developers from the Oracle Cloud Infrastructure(OCI) console or using the API.
- Upgrade Autonomous Database for Developers Instance to a Full Paid Instance
  You can upgrade an Autonomous Database for Developers instance to a full paid service instance.

- Features Not Available in Autonomous Database for Developers
  There are some features that are not available with an Autonomous Database for Developers instance, compared to a full paid service instance.

Resource Restrictions with Autonomous Database for Developers

Autonomous Database for Developers instances come with compute and storage resources included.

- 4 ECPUs per instance. The number of ECPUs cannot be scaled manually or automatically unless you upgrade to a full paid service instance.
- Maximum of approximately 20 GB storage per database (you may see more than this). The storage size cannot be scaled manually or automatically unless you upgrade to a full paid service instance.
- Maximum of 30 simultaneous database sessions
- The HTTP interface for Autonomous Database for Developers is rate limited to restrict the number of simultaneous service users. Approximately 3-6 simultaneous users can be supported across all of the APEX, Oracle REST Data Services, and Database Actions running on an Autonomous Database for Developers instance.

  Additional simultaneous users beyond that number may result in users encountering HTTP errors such as HTTP status code 429.
- The Maximum of 30 simultaneous database sessions limit for Autonomous Database for Developers allows you to work with Autonomous Database; however, if your usage includes many simultaneous users and/or many concurrent database client connections then you can exceed these limits, resulting in errors. To avoid such errors, obtain more resources for your Autonomous Database by upgrading to a full paid service instance. See Upgrade Autonomous Database for Developers Instance to a Full Paid Instance for more information.

Oracle Database Versions with Autonomous Database for Developers

Shows the available Oracle Database versions for Autonomous Database for Developers.

- **Oracle Database 19c**

  If you are using Autonomous Database for Developers with Oracle Database 19c, then many of the concepts and features of this service are further documented here:

  – Oracle Database 19c
- **Oracle Database 23ai**

  If you are using Autonomous Database for Developers with Oracle Database 23ai, then many of the concepts and features of this service are further documented here:

  – Oracle Database 23ai

  – *Oracle Database New Features*

Supported Features with Autonomous Database for Developers

Autonomous Database for Developers supports most but not all of the Autonomous Database features.

The following features help you use your Autonomous Database instance for development tasks:

- **Cloning**: Autonomous Database for Developers instances are low-cost, fixed shape databases intended for development and testing uses cases, and are not recommended for production use cases. For uses such as full scale load testing and production deployments or to access all Autonomous Database features, you can clone an Autonomous Database for Developers instance to a full paid service instance. See Clone an Autonomous Database Instance for instructions.
- **Service Maintenance:** Autonomous Database for Developers instances have the same patching schedule as full paid service instances.

  See About Scheduled Maintenance and Patching for more information.
- **Database Application Development and Developer Tools:** Autonomous Database for Developers provides all of the developer-related features and built-in tools that a full paid service Autonomous Database offers.

  See About Autonomous Database Built-in Tools for more information.

Autonomous Database for Developers comes with a service level objective (SLO) of 99.5%, and you can log service requests (SRs) to Oracle Support for assistance. However, there is no Severity 1 SR support for Autonomous Database for Developers instances. See Get Help, Search Forums, and Contact Support to learn how to contact Oracle Support for assistance.

Provisioning and Cloning Workflow with Autonomous Database for Developers

You can provision or clone Autonomous Database for Developers from the Oracle Cloud Infrastructure(OCI) console or using the API.

Provisioning or cloning Autonomous Database for Developers instances from the OCI console follows the same workflow as provisioning or cloning a full paid service Autonomous Database. In the **Configure the database** step, select **Autonomous Database for Developers**.

See Provision an Autonomous Database Instance and Clone an Autonomous Database Instance for more information.

Upgrade Autonomous Database for Developers Instance to a Full Paid Instance

You can upgrade an Autonomous Database for Developers instance to a full paid service instance.

If your Oracle Cloud Infrastructure account is in a trial period or has paying status you can upgrade your Autonomous Database for Developers database to a full paid service instance as follows:

1. From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
2. On the Autonomous Databases page select an Autonomous Database for Developers (**Developer**) instance from the links under the **Display name** column.
3. On the Autonomous Database Details page, from the **More actions** drop-down list select **Upgrade instance to full paid service**.
4. Click **Upgrade instance to full paid service**.

When you upgrade from an Autonomous Database for Developers instance to a full paid service instance, you get an instance with the same workload type. The ECPU and storage in the upgraded paid instance are initially set to the minimum ECPU and storage values based on your workload type.

> **Note:**
>
> After you upgrade you can scale up your compute and storage resources to fit your needs.

Autonomous Database for Developers Workload Type	Workload Type, ECPU and Storage Values
Data Warehouse	When you upgrade to full paid service from an Autonomous Database for Developers instance, the upgrade process provides a full paid service instance with workload type Data Warehouse, with 2 ECPUs and 1 TB of database storage.
Transaction Processing, JSON, or APEX	When you upgrade to full paid service from an Autonomous Database for Developers instance, the upgrade process provides a full paid service instance with the same workload type as you had, with 2 ECPUs and 20 GB of database storage.

Features Not Available in Autonomous Database for Developers

There are some features that are not available with an Autonomous Database for Developers instance, compared to a full paid service instance.

The limitations for Autonomous Database for Developers are in place to ensure that instances are specifically used as a development sandbox and to assist you in working with Autonomous Database. When you want to use a feature that is not available, you can upgrade to a full paid service Autonomous Database.

Autonomous Database for Developers restrictions and exclusions:

- Do not support disaster recovery options, including: Autonomous Data Guard, Backup-Based Disaster Recovery, or OCI Full Stack Disaster Recovery
- Supports ECPUs only
- Comes with fixed compute and storage sizing and does not support manual or auto-scaling and storage scaling.
- Does not support elastic pools
- Does not provide an SLA Guarantee
- Use of refreshable clones is not allowed
- Does not provide auto start-stop schedules
- Does not support Bring Your Own License (BYOL) licensing
- Autonomous Database for Developers databases cannot be provisioned as a private endpoint and cannot reside within a Virtual Cloud Network (VCN). See Configure Private Endpoints When You Provision or Clone an Instance for further information on private endpoints.
- Does not provide backup or restore operations:
  - No backups can be used (backups are shown)

- Does not support long-term backups
- Does not support restore (you cannot restore the current database to a selected past backup or timestamp).

To use backup and restore functionality on Autonomous Database you must upgrade an Autonomous Database for Developers to a full paid service instance.

See Upgrade Autonomous Database for Developers Instance to a Full Paid Instance for more information.

2
Quick Start

- Before You Begin with Oracle Autonomous Database
  Before you begin using Oracle Autonomous Database, you should be familiar with Oracle Cloud.
- Autonomous Database 15 Minute Quick Start
  Learn about Autonomous Database Serverless and learn how to create an Autonomous Database in just a few clicks.

Before You Begin with Oracle Autonomous Database

Before you begin using Oracle Autonomous Database, you should be familiar with Oracle Cloud.

Before you create an Autonomous Database:

- On Oracle Cloud, sign up for Oracle Cloud Free Tier or sign up for a paid Cloud Account. You cannot create an Autonomous Database deployment until you do so.

  See Welcome to Oracle Cloud Infrastructure for more information.
- (Optional) if you want to leverage an object store for data loading you need your object store credentials to use with Oracle Autonomous Database, including a *username* and a *password*. For details on the required credentials, depending on the object store you want to use, see the following:
  - **Oracle Cloud Infrastructure Object Storage**, the `username` is your Oracle Cloud Infrastructure user name. The `password` is your auth token. See Working with Auth Tokens.
  - **Oracle Cloud Infrastructure Object Storage Classic**, the `username` is your Oracle Cloud Infrastructure Classic user name and the `password` is your Oracle Cloud Infrastructure Classic password.
  - **Amazon S3**, the `username` is your AWS access key ID and the `password` is your AWS secret access key. See AWS Identity and Access Management.
  - **Azure Blob Storage or Azure Data Lake Storage**, the `username` is your Azure storage account name and the `password` is an Azure storage account access key. See About Azure storage accounts.
  - **Amazon S3-Compatible**, such as Oracle Cloud Infrastructure Object Storage, Google Cloud Storage, and Wasabi Hot Cloud Storage. See CREATE_CREDENTIAL Procedure for more information.
  - **GitHub Repository**, the `username` is your GitHub email and the `password` is your GitHub personal access token. See Creating a personal access token for more information.

Autonomous Database 15 Minute Quick Start

Learn about Autonomous Database Serverless and learn how to create an Autonomous Database in just a few clicks.

 Autonomous Database 15 Minute Quick Start Workshop

- Deploy an Autonomous Database instance that is optimized for data warehousing workloads.
- Use Autonomous Database tools to load object storage sources.
- Use advanced SQL to uncover issues and possibilities

3 Tasks

- Provision or Clone an Autonomous Database
  Provisioning a new Autonomous Database is easy. You can also provision a new Autonomous Database by cloning an existing instance.
- Connect to Autonomous Database
  Describes methods to securely connect to Autonomous Database.
- Manage Users
  Describes administration tasks for managing database users on Autonomous Database.
- Manage Credentials or Configure Policies and Roles to Access Resources
  Describes the methods in which Autonomous Database connects to other resources, either with credentials, or by configuring policies and roles for access.
- Load Data into Autonomous Database
  Describes packages and tools to load data with Autonomous Database.
- Link Data
  You can link to data in remote databases, in cloud storage, or in connected file systems or the local file system. You can also use Database Links to access objects on another database.
- Call Web Services from Autonomous Database
  Describes options for calling Web Services from Autonomous Database.
- Move Files
  Describes how to create directories and copy files to Autonomous Database.
- Replicate Data
  Describes using Oracle GoldenGate to replicate data between an Autonomous Database instance and to any target database or platform that Oracle GoldenGate supports.
- Exporting Data from Autonomous Database to Object Store or to Other Oracle Databases
  To export data from an Autonomous Database, use one of the following methods:
- Data Sharing
  Share data on Autonomous Database
- Develop
  Describes developer tasks and procedures on Autonomous Database.
- Analyze
  Provides information on the built-in tools that allow you to explore your data with reports, analytics, and visualization.
- Manage the Service
  Describes basic tasks for managing an Autonomous Database instance.
- Managing and Viewing Maintenance Windows, Patching Options, Work Requests, and Customer Contacts
- Monitor and Manage Performance
  Describes the Database Dashboard card in Database Actions where you can find information about the performance of an Autonomous Database instance.

- Use Sample Data Sets in Autonomous Database
  For users who want to start using the service without creating their own tables, Autonomous Database provides the read-only Sales History and Star Schema Benchmark data sets.
- Use the Autonomous Database Free Container Image

Provision or Clone an Autonomous Database

Provisioning a new Autonomous Database is easy. You can also provision a new Autonomous Database by cloning an existing instance.

See Provision an Autonomous Database Instance for details on how to create an Autonomous Database for your workload type using the Create Autonomous Database dialog.

See Clone an Autonomous Database Instance for details on cloning.

To get started, you can also try the LiveLabs Autonomous Database 15 Minute Quick Start.

After you create an Autonomous Database:

- (Optional) Reset the administrator password:

  When you create an Autonomous Database you are required to set the administrator password. If you want to change the administrator password or if you need to unlock the administrator account, see Manage the Administrator Account on Autonomous Database.
- (Optional) Add Oracle Machine Learning component users. See Create and Update User Accounts for Oracle Machine Learning Components on Autonomous Database for more information.
- (Optional) Download Oracle Analytics Desktop and add a connection to Autonomous Database. See Create Dashboards and Reports to Analyze and Visualize Your Data.

Connect to Autonomous Database

Describes methods to securely connect to Autonomous Database.

- About Connecting to an Autonomous Database Instance
  After you create database users, applications and tools connect to Autonomous Databases using Oracle Net Services (also known as SQL*Net). Oracle Net Services enables a network session from a client application to an Oracle Database server.
- Connect to Autonomous Database Using a Client Application
  Autonomous Database is preconfigured to support Oracle Net Services (a TNS listener is installed and configured to use secure TCPS).
- Download Database Connection Information
- Connect to Autonomous Database Using Oracle Database Tools
  Oracle Database Tools such as SQL Developer, SQL*Plus, and SQLcl can be used with Autonomous Database.
- Connect with Built-In Oracle Database Actions
  You can access Database Actions from Autonomous Database. Database Actions provides development tools, data tools, administration, and monitoring features for Autonomous Database. Using Database Actions you can run SQL statements, queries, and scripts in a worksheet.

- Connect with JDBC Thin Driver
  Autonomous Database mandates a secure connection that uses Transport Layer Security (TLSv1.2).
- Preparing for Oracle Call Interface Connections
  Using TLS with Oracle Call Interface connections with newer Oracle Instant Client/Oracle Database Client versions, a wallet is not required. For TLS connections with Oracle Call Interface running on older Oracle Instant Client/Oracle Database Client versions, a wallet is required for both mTLS authenticated connections and TLS authenticated connections.
- Predefined Database Service Names for Autonomous Database
  The `tnsnames.ora` file provided with the credentials zip file contains the database service names that allow you to connect to your database.
- Connect with Oracle Analytics Desktop
  Oracle Analytics Desktop makes it easy to visualize your data so you can focus on exploring interesting data patterns. Just connect to Autonomous Database, select the elements that you're interested in, and let Oracle Analytics Desktop find the best way to visualize it. Choose from a variety of visualizations to look at data in a specific way.
- Connect with Oracle Analytics Cloud
  Oracle Analytics Cloud is a scalable and secure public cloud service that provides a full set of capabilities to explore and perform collaborative analytics for you, your workgroup, and your enterprise.
- Connection and Networking Options and Features
  Autonomous Database provides a number of different connection and networking options and features for connecting to a database.

About Connecting to an Autonomous Database Instance

After you create database users, applications and tools connect to Autonomous Databases using Oracle Net Services (also known as SQL*Net). Oracle Net Services enables a network session from a client application to an Oracle Database server.

When a network session is established, Oracle Net Services acts as the data courier for both the client application and the database. It is responsible for establishing and maintaining the connection between the client application and the database, as well as exchanging messages between them.

Oracle Net Services supports a variety of connection types to connect to an Autonomous Database instance, including:

- **JDBC Thin Driver**: for Java applications, the JDBC Thin Driver is a pure Java driver. Many applications, including Oracle SQL Developer support JDBC Thin Driver connections.
- **JDBC OCI**: which is used by Java language applications. JDBC OCI adds a layer over Oracle Call Interface for Java applications. The Oracle SQLcl command-line interface uses JDBC OCI.
- **Oracle Call Interface (OCI)**: used by many applications written in C language. Examples that use Oracle Call Interface include Oracle utilities such as Oracle SQL*Plus, SQL*Loader, and Oracle Data Pump.
- **ODBC Drivers**: used by applications running on Microsoft Windows, that are layered over Oracle Call Interface (OCI).

Third-party products and custom applications may use any of these connection types.

Topics

- Secure Connections to Autonomous Database with mTLS or with TLS
  Connections to Autonomous Database are made either over the public Internet, optionally with Access Control Rules (ACLs) defined, or using a private endpoint inside a Virtual Cloud Network (VCN) in your tenancy.
- Connect to Autonomous Database Through a Firewall
  Most organizations protect networks and devices on a network using a firewall. A firewall controls incoming and outgoing network traffic using rules which allow the use of certain ports and access to certain computers (or, more specifically IP addresses or host names). An important function of a firewall is to provide separation between internal networks and the public internet.
- Using Application Continuity
  Application Continuity masks outages from end users and applications by recovering the in-flight work for impacted database sessions following outages. Application Continuity performs this recovery beneath the application so that the outage appears to the application as a slightly delayed execution.

Secure Connections to Autonomous Database with mTLS or with TLS

Connections to Autonomous Database are made either over the public Internet, optionally with Access Control Rules (ACLs) defined, or using a private endpoint inside a Virtual Cloud Network (VCN) in your tenancy.

When you specify a private endpoint configuration, this only allows traffic from the virtual cloud network you specify and blocks access to the database from all public IPs or VCNs. Configuring a private endpoint allows you to keep all traffic to and from your database off of the public internet. With a private endpoint, when public access is enabled with **Allow public access**, the instance has both a private endpoint and a public endpoint:

- The private hostname, endpoint URL, and private IP address allow you to connect to the database from the VCN where the database resides.
- The public hostname allows you to connect to the database from specific public IP addresses or from specific VCNs if those VCNs are configured to privately connect to Autonomous Database using a Service Gateway.

Many applications provide support for more than one connection type, but each type of connection to Autonomous Database uses certificate authentication and TCPS (Secure TCP) database connection using standard TLS 1.2. This ensures that there is no unauthorized access to the Autonomous Database and that communications between the client and server are fully encrypted and cannot be intercepted or altered.

Autonomous Database by default supports Mutual TLS (mTLS) connections (use port 1522 to connect with mTLS). You have the option to configure an Autonomous Database instance to support both mTLS and TLS connections (use either port 1521 or 1522 to connect with TLS).

There are advantages for clients using TLS authentication with Autonomous Database, including the following:

- TLS connections do not require that you download a wallet. For TLS connections using JDBC Thin Driver with JDK8 or higher, a wallet is not required. This includes connections coming from the clients such as SQL Developer and SQL Command Line (SQLcl).
- Clients connecting with TLS do not need to worry about wallet rotation. Wallet rotation is a regular procedure for mTLS connections.

- TLS connections can be faster (providing less connection latency). TLS authentication can provide reduced connection latency compared to mTLS.
- TLS and mTLS connections are not mutually exclusive. Mutual TLS (mTLS) authentication is enabled by default and always available. When you enable TLS authentication, you can use either mTLS or TLS authentication.
- Using TLS authentication does not compromise the fully encrypted end-to-end communication between a client and Autonomous Database.

See View TNS Names and Connection Strings for an Autonomous Database Instance for details on obtaining the mTLS connection strings for your Autonomous Database Instance.

About Mutual TLS (mTLS) Authentication

Using Mutual Transport Layer Security (mTLS), clients connect through a TCPS (Secure TCP) database connection using standard TLS 1.2 with a trusted client certificate authority (CA) certificate.

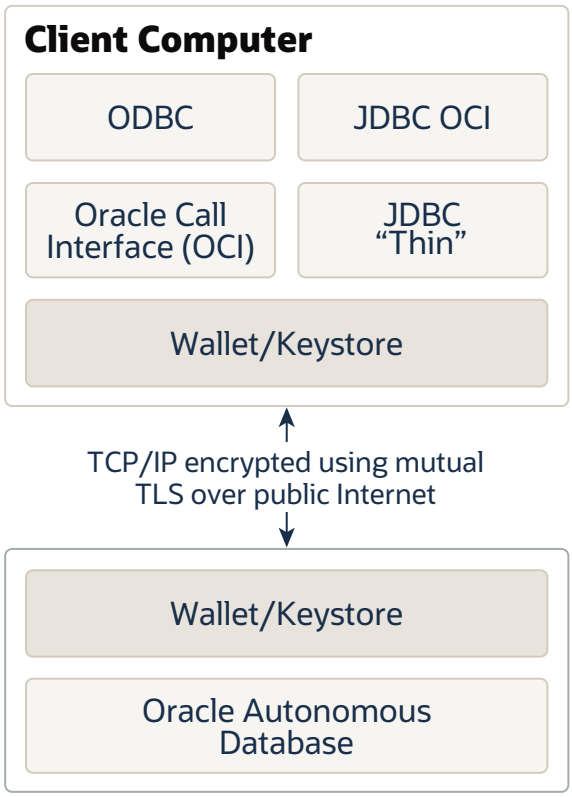
With mutual authentication both the client application and Autonomous Database authenticate each other. Autonomous Database uses mTLS authentication by default. Use port 1522 to connect to an Autonomous Database instance with mTLS (the 1522 port assignment cannot be changed).

Mutual TLS authentication requires that the client downloads or obtains a trusted client CA certificate for connecting to an Autonomous Database instance. Autonomous Database then uses the certificate to authenticate the client. This provides increased security and specifies the clients that can communicate with an Autonomous Database instance.

Certification authentication with Mutual TLS uses an encrypted key stored in a *wallet* on both the client (where the application is running) and the server (where your database service on the Autonomous Database is running). The key on the client must match the key on the server to make a connection. A wallet contains a collection of files, including the key and other information needed to connect to your Autonomous Database instance. All communications between the client and the server are encrypted.

To secure the connection to your Autonomous Database instance a service administrator downloads the client credentials (wallet files) from Autonomous Database. If you are not an Autonomous Database service administrator, your administrator provides you with the client credentials. See Download Client Credentials (Wallets) for more information.

The following figure shows client secure connections to Oracle Autonomous Database over the public Internet using Mutual TLS connections. If you configure your database to use private endpoints, then the public internet is not used and the connection uses a private endpoint inside a Virtual Cloud Network (VCN) in your tenancy.



See View TNS Names and Connection Strings for an Autonomous Database Instance for details on obtaining the mTLS connection strings for your Autonomous Database Instance.

About TLS Authentication

Using Transport Layer Security (TLS), clients connect through a TCPS (Secure TCP) database connection using standard TLS 1.2. A client uses its list of trusted Certificate Authorities (CA)s to validate the server's CA root certificate. If the issuing CA is trusted, the client verifies that the certificate is authentic.

This allows the client and Autonomous Database to establish the encrypted connection before exchanging any messages. Use either port 1521 or 1522 to connect to an Autonomous Database instance with TLS (these port assignments cannot be changed).

When you connect with TLS authentication using JDBC Thin Driver clients, including Oracle SQL Developer and Oracle SQLcl, you do not need to download a wallet to secure the connection to your Autonomous Database instance. TLS authentication enables the client to verify the identity of the Autonomous Database service to provide secure communication.

Depending on the type of client, a TLS connection has the following support with Autonomous Database:

- For connections with JDBC Thin Driver using JDK8u162 or higher, including connections with Oracle SQL Developer and Oracle SQLcl, a wallet is not required.
- Oracle Call Interface (OCI) clients support TLS authentication without a wallet if you are using the following client versions:
  - Oracle Instant Client/Oracle Database Client 19.13 - only on Linux x64
  - Oracle Instant Client/Oracle Database Client 19.14 (or later) and 21.5 (or later) - only on Linux x64 and Windows
- If the client is connecting with managed ODP.NET or ODP.NET Core versions 19.13 or 21.4 (or above) using TLS authentication, the client can connect without providing a wallet.

There are network access prerequisites for TLS connections. See Network Access Prerequisites for TLS Connections for more information.

See View TNS Names and Connection Strings for an Autonomous Database Instance for details on obtaining the TLS connection strings for your Autonomous Database Instance.

Connect to Autonomous Database Through a Firewall

Most organizations protect networks and devices on a network using a firewall. A firewall controls incoming and outgoing network traffic using rules which allow the use of certain ports and access to certain computers (or, more specifically IP addresses or host names). An important function of a firewall is to provide separation between internal networks and the public internet.

When Autonomous Database is configured for access using the public internet, you must configure the firewall to allow access to Autonomous Database servers.

To access the Autonomous Database from behind a firewall, the firewall must permit the use of the port specified in the database connection when connecting to the servers in the connection. Use port 1522 for Autonomous Database mTLS connections (you can see the port number in the connection string from the `tnsnames.ora` file in your `credentials ZIP` file). For example, see the `port` value in the following `tnsnames.ora` file:

```
db2022adb_high = (description = (
                address=(protocol=tcps)
                (port=1522)
                (host=adb.example.oraclecloud.com))
                (connect_data=(service_name=example_high.adb.oraclecloud.com))
                (security=(ssl_server_dn_match=yes)))
```

Your firewall must allow access to servers within the `.oraclecloud.com` domain using port 1522. To connect to Autonomous Database, depending upon your organization's network configuration, you may need to use a proxy server to access this port or you may need to request that your network administrator open this port.

See IP Address Ranges for information about the public IP address ranges in Oracle Cloud Infrastructure. You must allow traffic to these CIDR blocks to ensure access to an Autonomous Database instance on a public endpoint.

Using Application Continuity

Application Continuity masks outages from end users and applications by recovering the in-flight work for impacted database sessions following outages. Application Continuity performs this recovery beneath the application so that the outage appears to the application as a slightly delayed execution.

> **Note:**
>
> By default Application Continuity is disabled.

See Use Application Continuity on Autonomous Database for more information on Application Continuity.

Connect to Autonomous Database Using a Client Application

Autonomous Database is preconfigured to support Oracle Net Services (a TNS listener is installed and configured to use secure TCPS).

The client computer must be prepared to use Oracle Net Services to connect to Autonomous Database.

- About Connecting to Autonomous Database Using a Client Application
  Applications can connect to Autonomous Database using any of the connection types supported by Oracle Net Services.
- Prepare for Oracle Call Interface, ODBC, and JDBC OCI Connections with Wallets (mTLS)
  Preparing for any type of Oracle Call Interface (OCI) connection with mTLS authentication requires the installation of client software, downloading client credentials, and configuring certain files and environment variables.
- Prepare for Oracle Call Interface, ODBC, and JDBC OCI Connections Using TLS Authentication
  Preparing for any type of Oracle Call Interface (OCI) connection with TLS authentication requires the installation of client software and configuring certain files and environment variables.
- Prepare for JDBC Thin Connections
  Applications that use JDBC Thin connections include the software necessary to make an Oracle Net Services connection. It is not necessary to download and install Oracle Client software.
- Connect Microsoft .NET, Visual Studio Code, and Visual Studio with a Wallet (mTLS)
  Oracle Autonomous Database supports connectivity to the Microsoft .NET Framework, .NET Core, Visual Studio, and Visual Studio Code.
- Connect Microsoft .NET, Visual Studio Code, and Visual Studio Without a Wallet
  Oracle Autonomous Database supports connectivity to the Microsoft .NET Framework, .NET Core, Visual Studio, and Visual Studio Code using TLS authentication without a wallet.
- Connect Node.js and other Scripting Languages (mTLS)
  You can use programs in different languages, including Python, Node.js, PHP, Ruby, R, Go, and Perl to connect to an Autonomous Database instance using mTLS (with wallets). Security is enforced using client credentials.
- Connect Node.js, and Other Scripting Languages Without a Wallet
  You can use programs in different languages, including Python, Node.js, PHP, Ruby, R, Go, and Perl to connect to an Autonomous Database instance using TLS authentication without a wallet.
- Connect Power BI and Microsoft Data Tools to Autonomous Database
  Oracle Client for Microsoft Tools (OCMT) is a graphical user interface (GUI) native Microsoft Software installer (MSI) that simplifies ODP.NET setup and Oracle database connectivity to multiple Microsoft data tools.
- Connect Python Applications to Autonomous Database
  You can connect Python applications to your Autonomous Database instance either with a wallet (mTLS) or without a wallet (TLS).
- Connect Go Applications to Autonomous Database
  You can connect Go applications to your Autonomous Database instance either with a wallet (mTLS) or without a wallet (TLS).

About Connecting to Autonomous Database Using a Client Application

Applications can connect to Autonomous Database using any of the connection types supported by Oracle Net Services.

Consult your application documentation for details about how your application connects to Oracle.

The following steps describe the process of connecting to Autonomous Database using a client application:

1. Determine what connection type your application uses, (for example OCI, ODBC, JDBC Thin, and so on).
2. The steps required to prepare the client computer depend on the type of connection the client application uses. Determine what authentication type your application uses, either mTLS or TLS and prepare your client computer:
   - **Mutual TLS authentication connections**: In all cases, with mTLS connections, client credentials in the form of a wallet file must be downloaded to the client.
   - **TLS authentication connections**: For connections with JDBC Thin Driver using JDK8u162 or higher, including connections with Oracle SQL Developer and Oracle SQLcl, a wallet is not required.

     Oracle Call Interface (OCI) clients support TLS authentication without a wallet if you are using the following client versions:

     – Oracle Instant Client/Oracle Database Client 19.13 - only on Linux x64

     – Oracle Instant Client/Oracle Database Client 19.14 (or later) and 21.5 (or later) - only on Linux x64 and Windows

   See the details for each driver type for information on the steps required to prepare your application to connect to Autonomous Database.
3. Within your application, set up the connection.

Prepare for Oracle Call Interface, ODBC, and JDBC OCI Connections with Wallets (mTLS)

Preparing for any type of Oracle Call Interface (OCI) connection with mTLS authentication requires the installation of client software, downloading client credentials, and configuring certain files and environment variables.

This topic covers the steps to prepare an application to connect using mTLS authentication with a wallet that you download from an Autonomous Database instance. See Prepare for Oracle Call Interface, ODBC, and JDBC OCI Connections Using TLS Authentication for information on the steps to prepare for TLS authentication with these connection types.

New Oracle Client Installation

The following steps assume Oracle client software has not already been installed on the client computer. If Oracle client software has already been installed and there are working copies of `sqlnet.ora` and `tnsnames.ora`, see Updating an Existing Oracle Client Installation.

Before making an Oracle Call Interface (OCI), ODBC, or JDBC OCI connection, do the following:

1. Install Oracle Client software on your computer.

- **Oracle Instant Client/Oracle Database Client**: 18.19 (or later), 19.2 (or later), or 21 (base release or later). The Instant Client contains the minimal software needed to make an Oracle Call Interface connection.

2. Download client credentials and store the file in a secure folder on your client computer. See Download Client Credentials (Wallets).
3. Unzip/uncompress the credentials file into a secure folder on your client computer.
4. Edit the `sqlnet.ora` file in the folder where you unzip the credentials file, replacing `"?/network/admin"` with the name of the folder containing the client credentials.

   For example, edit `sqlnet.ora` as follows:

   ```
   WALLET_LOCATION = (SOURCE = (METHOD = file) (METHOD_DATA = (DIRECTORY="?/
   network/admin")))
   SSL_SERVER_DN_MATCH=yes
   ```

   The changed value on UNIX/Linux is:

   ```
   WALLET_LOCATION = (SOURCE = (METHOD = file) (METHOD_DATA = (DIRECTORY="/
   home/adb_credentials")))
   SSL_SERVER_DN_MATCH=yes
   ```

   The changed value for Windows is:

   ```
   WALLET_LOCATION = (SOURCE = (METHOD = file) (METHOD_DATA = (DIRECTORY="D:\
   \myapp\\adb_credentials")))
   SSL_SERVER_DN_MATCH=yes
   ```

5. Create the `TNS_ADMIN` environment variable and set it to the location of the credentials file.

   Use this environment variable to change the directory path of Oracle Net Services configuration files from the default location of `ORACLE_HOME\network\admin` to the location of the secure folder containing the credentials file you saved in Step 2. Set the `TNS_ADMIN` environment variable to the directory where the unzipped credentials files are, not to the credentials file itself.

   For example, on UNIX/Linux set `TNS_ADMIN` to the full path of the directory where you unzipped the client credentials:

   ```
   export TNS_ADMIN=/home/adb_credentials
   ```

   For example on Windows:

   ```
   set TNS_ADMIN=d:\myapp\adb_credentials
   ```

Connections with an HTTP Proxy

If the client is behind a firewall and your network configuration requires an HTTP proxy to connect to the internet, then perform the following steps to update the `sqlnet.ora` and `tnsnames.ora` files. Connections through an HTTP proxy are only available with Oracle Client software version 12.2.0.1 or later.

> **Note:**
>
> To avoid manual updates in `sqlnet.ora` and `tnsnames.ora` files, you can use SQLcl and specify the HTTP proxy on the command line. See Connect Oracle SQLcl Cloud with a Wallet (mTLS) for more information.

1. Add the following line to the `sqlnet.ora` file to enable connections through an HTTP proxy:

   ```
   SQLNET.USE_HTTPS_PROXY=on
   ```

2. Add the HTTP proxy hostname and port to the connection definitions in `tnsnames.ora`. You need to add the `https_proxy` and `https_proxy_port` parameters in the address section of connection definitions. For example, the following sets the HTTP proxy to `proxyhostname` and the HTTP proxy port to `80`; replace these values with your HTTP proxy information:

   ```
   ADB1_high =
          (description=
                (address=
                      (https_proxy=proxyhostname)(https_proxy_port=80)
   (protocol=tcps)(port=1522)(host=adb.example.oraclecloud.com)
                )
                (connect_data=(service_name=adb1_high.adb.oraclecloud.com)
                )
                (security=(ssl_server_dn_match=yes))
                )
          )
   ```

> **Note:**
>
> Configuring `sqlnet.ora` and `tnsnames.ora` for the HTTP proxy may not be enough depending on your organization's network configuration and security policies. For example, some networks require a `username` and `password` for the HTTP proxy. In such cases contact your network administrator to open outbound connections to hosts in the oraclecloud.com domain using port 1522 without going through an HTTP proxy.

For more information on `SQLNET.USE_HTTPS_PROXY`, see *Net Services Reference*.

For information on `HTTPS_PROXY` and `HTTPS_PROXY_PORT`, see Protocol Address Section.

Updating an Existing Oracle Client Installation

If you have an existing Oracle Client installation, you already have `sqlnet.ora` and `tnsnames.ora` files and the `TNS_ADMIN` environment variable. In this case, do the following:

1. Update your `sqlnet.ora` file by adding the following:

   ```
   WALLET_LOCATION = (SOURCE = (METHOD = file) (METHOD_DATA = (DIRECTORY="/
   home/adb_credentials")))
   ```

2. Copy the entries in the `tnsnames.ora` file provided in the Autonomous Database wallet to your existing `tnsnames.ora` file.

Prepare for Oracle Call Interface, ODBC, and JDBC OCI Connections Using TLS Authentication

Preparing for any type of Oracle Call Interface (OCI) connection with TLS authentication requires the installation of client software and configuring certain files and environment variables.

Oracle Call Interface (OCI) clients support TLS authentication without a wallet if you are using the following client versions:

- Oracle Instant Client/Oracle Database Client 19.13 - only on Linux x64
- Oracle Instant Client/Oracle Database Client 19.14 (or later) and 21.5 (or later) - only on Linux x64 and Windows

See Update your Autonomous Database Instance to Allow both TLS and mTLS Authentication for information on allowing TLS connections.

1. Install Oracle Instant Client.
   a. Got to the Oracle Instant Client page and click **Download Now**: Oracle Instant Client
   b. On the Oracle Instant Client Downloads page, select your platform.

      For example, under **Instant Client for Linux**, select the Instant Client for Linux x86-64 architecture (for this example, to download the Linux x86-64 version).
   c. Under **Version 19.14.0.0.0 (Requires glibc 2.14)**, select an Instant Client package to download.
   d. If you are building a language API or driver from source code, you may also need to download the Instant Client SDK Package version 19.14: Oracle Instant Client
   e. Unzip the base package you selected. If you also download the SDK, unzip it in the same directory.
   f. On Linux, create a symbolic link if it does not exist. For example:

      ```
      cd /home/myuser/instantclient_19_14
      ln -s libclntsh.so.19.1 libclntsh.so
      ```

      If there is no other Oracle software on your system that will be impacted, add Instant Client to the runtime link path. For example:

      ```
      sudo sh -c "echo /home/myuser/instantclient_19_14 > /etc/ld.so.conf.d/
      oic.conf"
      sudo ldconfig
      ```

      Alternatively set the library path in each shell that runs your application. For example:

      ```
      export LD_LIBRARY_PATH=/home/myuser/instantclient_19_14:$LD_LIBRARY_PATH
      ```

> **Note:**
>
> The Linux Instant Client download files are available as `.zip` files or `.rpm` files. You can use either version.

2. If you have not already done so, enable TLS connections on your Autonomous Database instance.

   See Update your Autonomous Database Instance to Allow both TLS and mTLS Authentication for details.

3. Run Your Application

   a. Update your application to connect using your database username, your password, and the Oracle Net connect name given in the unzipped `tnsnames.ora` file. For example, user, `adb_user`, *password*, and `db2022adb_low` as the connect string.

   b. Alternatively, change the connect string in `tnsnames.ora` to match the string used by your application.

   c. Run your application.

Allowing TLS connections to Autonomous Database does not disallow mutual TLS (mTLS) connections. Both Mutual TLS (mTLS) and TLS connections are valid when an Autonomous Database instance is configured to allow TLS connections. See Connect Node.js and other Scripting Languages (mTLS) for information on connecting using mutual TLS (mTLS) with a wallet.

In this case, update your `sqlnet.ora` file by adding the following:

```
WALLET_LOCATION = (SOURCE = (METHOD = file) (METHOD_DATA = (DIRECTORY="/home/
wallet1")))
```

Prepare for JDBC Thin Connections

Applications that use JDBC Thin connections include the software necessary to make an Oracle Net Services connection. It is not necessary to download and install Oracle Client software.

Some applications use the JDK installed on your computer while others use a JDK that is embedded in the application installation. If your application uses the JDK installed on your computer and that JDK is version 8, 8u161 or later, no additional preparation is required. If your computer does not have JDK version 8, 8u161 or later, already installed then install the latest JDK first. You can download the latest JDK from https://www.java.com/.

If your application is using a JDK version 8, prior to 8u161, then the JCE Policy Files must be updated within your application.

See Connect with JDBC Thin Driver for the steps required to use JDBC Thin connections to connect to an Oracle Database server.

- Set JVM Networking Properties
  Autonomous Database uses DNS names that map to multiple IP addresses (multiple load balancers) for better availability and performance. Depending on your application, you may want to configure certain JVM networking properties.

Set JVM Networking Properties

Autonomous Database uses DNS names that map to multiple IP addresses (multiple load balancers) for better availability and performance. Depending on your application, you may want to configure certain JVM networking properties.

For the Java Virtual Machine (JVM) address cache, any address resolution attempt caches the result whether it was successful or not, so that subsequent identical requests do not have to access the naming service. The address cache properties allow you to tune how the cache operates. In particular, the `networkaddress.cache.ttl` value specifies the number of seconds a successful name lookup is kept in the cache. A value of -1, the default value, indicates a “cache forever” policy, while a value of `0` (zero) means no caching.

If your Java Virtual Machine (JVM) is configured to cache DNS address lookups, your application may be using only one IP address to connect to your Autonomous Database, resulting in lower throughput. To prevent this you can change your JVM's `networkaddress.cache.ttl` value to `0`, so that every connection request does a new DNS lookup. This ensures that different threads in your application are distributed over multiple load balancers.

To change the `networkaddress.cache.ttl` value for all applications, or in your application, do one of the following:

- Configure the security policy to set the value for all applications:
  Set `networkaddress.cache.ttl=0` in the file `$JAVA_HOME/jre/lib/security/java.security`
- Set the following property in your application code:

```
java.security.Security.setProperty("networkaddress.cache.ttl" , "0");
```

Connect Microsoft .NET, Visual Studio Code, and Visual Studio with a Wallet (mTLS)

Oracle Autonomous Database supports connectivity to the Microsoft .NET Framework, .NET Core, Visual Studio, and Visual Studio Code.

Oracle Data Provider for .NET (ODP.NET) provides run-time ADO.NET data access to a database.

ODP.NET has the following driver types:

- Unmanaged ODP.NET for .NET Framework applications
- Managed ODP.NET for .NET Framework applications
- ODP.NET Core for .NET Core applications

The following minimum versions are required:

- **ODP.NET**: 12.1 (Starting with April 2022 WINDBBP; ODP.NET Managed only), 18 (base release or later), 19.4 (or later, except for 19.10), 21 (base release or later), or 23 (base release or later).

Oracle Developer Tools for Visual Studio and Oracle Developer Tools for VS Code provide database application design-time support in the Microsoft development environment, including tools for managing Autonomous Databases in the Oracle Cloud.

Oracle Developer Tools for VS Code provides database application design-time support in Visual Studio Code.

These software components are available as a free download from the following sites:

- Managed ODP.NET and ODP.NET Core: NuGet Gallery
- Unmanaged ODP.NET: Oracle Data Access Components Downloads
- Oracle Developer Tools for Visual Studio Code: VS Code Marketplace
- Oracle Developer Tools for Visual Studio: Oracle Developer Tools for Visual Studio

Oracle recommends using the latest provider and tools version with Oracle Autonomous Database.

Click the following link for instructions to download, install, and configure these components for use with mutual TLS (mTLS) and wallets:

- Developing .NET Applications for Oracle Autonomous Database

To use mutual TLS (mTLS) authentication to connect to Autonomous Database, you must download a wallet. See Download Client Credentials (Wallets) for information downloading wallets.

ODP.NET connectivity to Oracle Autonomous Database supports mutual TLS (mTLS) with wallets as well as TLS. See Connect Microsoft .NET, Visual Studio Code, and Visual Studio Without a Wallet for information on using TLS connections.

To learn more about using Oracle Autonomous Database and .NET, try the free .NET Development with Oracle Autonomous Database Quick Start. This lab walks you through setting up a .NET web server on Oracle Cloud Infrastructure that connects to Oracle Autonomous Database. Next, the lab guides developing and deploying a simple ASP.NET Core web application that uses all these components. By the end, you will have a live, working website on the Internet.

Connect Microsoft .NET, Visual Studio Code, and Visual Studio Without a Wallet

Oracle Autonomous Database supports connectivity to the Microsoft .NET Framework, .NET Core, Visual Studio, and Visual Studio Code using TLS authentication without a wallet.

Oracle Data Provider for .NET (ODP.NET) provides run-time ADO.NET data access to Autonomous Database. ODP.NET has the following driver types:

- Unmanaged ODP.NET for .NET Framework applications
- Managed ODP.NET for .NET Framework applications
- ODP.NET Core for .NET Core applications

Oracle Developer Tools for Visual Studio and Oracle Developer Tools for VS Code provide database application design-time support in the Microsoft development environment, including tools for managing Autonomous Databases in the Oracle Cloud.

Oracle Developer Tools for VS Code provides database application design-time support in Visual Studio Code.

These software components are available as a free download from the following sites:

- Managed ODP.NET and ODP.NET Core: NuGet Gallery
- Unmanaged ODP.NET: Oracle Data Access Components Downloads
- Oracle Developer Tools for Visual Studio Code: VS Code Marketplace
- Oracle Developer Tools for Visual Studio: Oracle Developer Tools for Visual Studio

Oracle recommends using the latest provider and tools version with Oracle Autonomous Database.

When you connect using TLS authentication with managed ODP.NET and ODP.NET Core you do not need to deploy the Oracle wallet or the Oracle network configuration files `sqlnet.ora` or `tnsnames.ora` with your application. Instead, you supply the data source attribute, a TLS connection string, with the configuration information in the ODP.NET connection.

To use TLS connections with Managed ODP.NET and ODP.NET Core, do the following:

1. Obtain managed ODP.NET or ODP.NET Core versions 19.14 (or later), 21.5 (or later), or 23 (base release or later). Lower level versions do not support TLS connections with Oracle Autonomous Database.
2. Enable TLS connections on your Autonomous Database instance. See Update your Autonomous Database Instance to Allow both TLS and mTLS Authentication for details.
3. After you enable TLS connections, supply a TLS connection string in the ODP.NET data source to connect to an Autonomous Database instance. See View TNS Names and Connection Strings for an Autonomous Database Instance for details on viewing or copying TLS connection strings.

Allowing TLS connections to Autonomous Database does not disallow mutual TLS (mTLS) connections. Both Mutual TLS (mTLS) and TLS connections are valid when an Autonomous Database instance is configured to allow TLS connections. See Connect Microsoft .NET, Visual Studio Code, and Visual Studio with a Wallet (mTLS) for information on connecting using mutual TLS (mTLS) with a wallet.

To learn more about using Oracle Autonomous Database and .NET, try the free .NET Development with Oracle Autonomous Database Quick Start. This lab walks you through setting up a .NET web server on Oracle Cloud Infrastructure that connects to Oracle Autonomous Database. Next, the lab guides developing and deploying a simple ASP.NET Core web application that uses all these components. By the end, you will have a live, working website on the Internet.

Connect Node.js and other Scripting Languages (mTLS)

You can use programs in different languages, including Python, Node.js, PHP, Ruby, R, Go, and Perl to connect to an Autonomous Database instance using mTLS (with wallets). Security is enforced using client credentials.

These scripting languages have database access APIs or drivers that use the Oracle Call Interface libraries.

> **Note:**
>
> Oracle Call Interface (OCI) clients support mTLS authentication with a wallet if you are connecting using the following client versions:
>
> - Oracle Instant Client/Oracle Database Client: 18.19 (or later), 19.2 (or later), or 21 (base release or later).
>
> For details on connecting Node.js or other scripting languages without a wallet, see Connect Node.js, and Other Scripting Languages Without a Wallet.

For additional details, see the following video:

- Connect Node.js Apps to Autonomous Transaction Processing

Install the Language Driver and Client Libraries

To connect to Autonomous Database from your scripting language, first install the language driver and client libraries as follows:

1. Install Oracle Instant Client/Oracle Database Client: 18.19 (or later), 19.2 (or later), or 21 (base release or later).

   The Instant Client works well for most applications. To install the Instant Client do the following:

   a. Select your desired architecture from the Instant Client Downloads page and download a Basic Package (available on the download page): Oracle Instant Client Downloads.

      Alternatively download the Basic Light Package from the download page for your desired architecture if the Basic Light globalization limitations suit your use.

   b. If you are building a language API or driver from source code, you may also need to download the Instant Client SDK: Oracle Instant Client

   c. Unzip the base package you selected. For example unzip to `C:\instantclient_12_2` or `/home/myuser/instantclient_18_5`. If you also download the SDK, unzip it in the same directory.

   d. On Windows, add the path to the `PATH` variable in the "System variables" section of the Environment Variables pane (for example add `C:\instantclient_12_2`). On Windows 8 access the `PATH` variable setting area by navigating to Control Panel>System>Advanced System Settings>Environment Variables. If you have multiple versions of Oracle libraries installed make sure the new directory occurs first in the path.

   e. On non-Windows platforms, create a symbolic link if it does not exist. For example:

      ```
      cd /home/myuser/instantclient_18_5
      ln -s libclntsh.so.18.1 libclntsh.so
      ```

      If there is no other Oracle software on your system that will be impacted, add Instant Client to the runtime link path. For example:

      ```
      sudo sh -c "echo /home/myuser/instantclient_18_5 > /etc/ld.so.conf.d/
      oic.conf"
      sudo ldconfig
      ```

      Alternatively set the library path in each shell that runs your application. For example:

      ```
      export LD_LIBRARY_PATH=/home/myuser/instantclient_18_5:$LD_LIBRARY_PATH
      ```

      > **Note:**
      >
      > The Linux Instant Client download files are available as `.zip` files or `.rpm` files. You can use either version.

2. Install the relevant language driver for Oracle Database:

   - **Node.js** : To install node-oracledb for Node.js,, use the instructions on the following page: Installing node-oracledb.

- **ROracle**: To install ROracle for R, use the instructions on the following page: ROracle
- **PHP**: To install PHP OCI8 for PHP, use the instructions on the following page: Configuring PHP with OCI8.

  Windows DLLs are available on http://php.net/downloads.php and are also available from PECL oci8.
- **PHP PDO_OCI**: To install PHP PDO_OCI for PHP, use the instructions on the following page: Oracle Functions (PDO_OCI).

  Windows DLLs are available on http://php.net/downloads.php included in PHP.
- **Ruby**: To install ruby-oci8 for Ruby, use the instructions on the following page: Install for Oracle Instant Client
- **DBD for Perl**: To install DBD::Oracle for Perl, set `ORACLE_HOME` and your library search path such as `LD_LIBRARY_PATH` or `PATH` to the Instant Client directory and use the instructions on the following page: Installing DBD-Oracle.
- **Python**: For instructions for connecting Python applications, see Connect Python Applications with a Wallet (mTLS)

Enable Oracle Network Connectivity and Obtain the Security Credentials (Oracle Wallet)

1. Obtain client security credentials to connect to Autonomous Database. You obtain a zip file containing client security credentials and network configuration settings required to access your database. You must protect this file and its contents to prevent unauthorized database access. Obtain the client security credentials file as follows:
   - ADMIN user: Click **Database connection**. See Download Client Credentials (Wallets).
   - Other user (non-administrator): Obtain the Oracle Wallet from the administrator for your Autonomous Database.
2. Extract the client credentials (wallet) files:
   a. Unzip the client credentials zip file.
   b. If you are using Instant Client, make a `network/admin` subdirectory hierarchy under the Instant Client directory if necessary. Then move the files to this subdirectory. For example depending on the architecture or your client system and where you installed Instant Client, the files should be in the directory:

      ```
      C:\instantclient_12_2\network\admin
      ```

      or

      ```
      /home/myuser/instantclient_18_5/network/admin
      ```

      or

      ```
      /usr/lib/oracle/18.5/client64/lib/network/admin
      ```

      - If you are using a full Oracle Client move the file to `$ORACLE_HOME/network/admin`.

   c. Alternatively, put the unzipped wallet files in a secure directory and set the `TNS_ADMIN` environment variable to that directory name.

> **Note:**
>
> From the zip file, only these files are required: `tnsnames.ora`, `sqlnet.ora`, `cwallet.sso`, and `ewallet.p12`.

3. If you are behind a proxy follow the steps in “Connections with an HTTP Proxy”, in Prepare for Oracle Call Interface, ODBC, and JDBC OCI Connections with Wallets (mTLS).

Run Your Application

1. Update your application to connect using your database username, your password, and the Oracle Net connect name given in the unzipped `tnsnames.ora` file. For example, user, `adb_user`, *password*, and `db2022adb_low` as the connect string.
2. Alternatively, change the connect string in `tnsnames.ora` to match the string used by your application.
3. Run your application.

Connect Node.js, and Other Scripting Languages Without a Wallet

You can use programs in different languages, including Python, Node.js, PHP, Ruby, R, Go, and Perl to connect to an Autonomous Database instance using TLS authentication without a wallet.

These scripting languages have database access APIs or drivers that use the Oracle Call Interface libraries. The Oracle Call Interface libraries can be either from the full Oracle Client or from Oracle Instant Client.

> **Note:**
>
> Oracle Call Interface (OCI) clients support TLS authentication without a wallet if you are using the following client versions:
>
> - Oracle Instant Client/Oracle Database Client 19.13 - only on Linux x64
> - Oracle Instant Client/Oracle Database Client 19.14 (or later) and 21.5 (or later) - only on Linux x64 and Windows

1. Install Oracle Instant Client.
   a. Got to the Oracle Instant Client page and click **Download Now**: Oracle Instant Client
   b. On the Oracle Instant Client Downloads page, select your platform.

      For example, under **Instant Client for Linux**, select the Instant Client for Linux x86-64 architecture (for this example, to download the Linux x86-64 version).
   c. Under **Version 19.14.0.0.0 (Requires glibc 2.14)**, select an Instant Client package to download.
   d. If you are building a language API or driver from source code, you may also need to download the Instant Client SDK Package version 19.14: Oracle Instant Client
   e. Unzip the base package you selected. If you also download the SDK, unzip it in the same directory.

f. On Linux, create a symbolic link if it does not exist. For example:

```
cd /home/myuser/instantclient_19_14
ln -s libclntsh.so.19.1 libclntsh.so
```

If there is no other Oracle software on your system that will be impacted, add Instant Client to the runtime link path. For example:

```
sudo sh -c "echo /home/myuser/instantclient_19_14 > /etc/ld.so.conf.d/
oic.conf"
sudo ldconfig
```

Alternatively set the library path in each shell that runs your application. For example:

```
export LD_LIBRARY_PATH=/home/myuser/instantclient_19_14:$LD_LIBRARY_PATH
```

> **Note:**
>
> The Linux Instant Client download files are available as `.zip` files or `.rpm` files. You can use either version.

2. Install the relevant language driver for Oracle Database:
   - **Node.js** : To install node-oracledb for Node.js,, use the instructions on the following page: Installing node-oracledb.
   - **ROracle**: To install ROracle for R, use the instructions on the following page: ROracle
   - **PHP**: To install PHP OCI8 for PHP, use the instructions on the following page: Configuring PHP with OCI8.

     Windows DLLs are available on http://php.net/downloads.php and are also available from PECL oci8.
   - **PHP PDO_OCI**: To install PHP PDO_OCI for PHP, use the instructions on the following page: Oracle Functions (PDO_OCI).

     Windows DLLs are available on http://php.net/downloads.php included in PHP.
   - **Ruby**: To install ruby-oci8 for Ruby, use the instructions on the following page: Install for Oracle Instant Client
   - **DBD for Perl**: To install DBD::Oracle for Perl, set `ORACLE_HOME` and your library search path such as `LD_LIBRARY_PATH` or `PATH` to the Instant Client directory and use the instructions on the following page: Installing DBD-Oracle.
   - **Python**: For instructions for connecting Python applications, see Connect Python Applications Without a Wallet (TLS)
3. If you have not already done so, enable TLS connections on your Autonomous Database instance.

   See Update your Autonomous Database Instance to Allow both TLS and mTLS Authentication for details.
4. Run Your Application

a. Update your application to connect using your database username, your password, and the Oracle Net connect name given in the unzipped `tnsnames.ora` file. For example, user, `adb_user`, *password*, and `db2022adb_low` as the connect string.

b. Alternatively, change the connect string in `tnsnames.ora` to match the string used by your application.

c. Run your application.

Allowing TLS connections to Autonomous Database does not disallow mutual TLS (mTLS) connections. Both Mutual TLS (mTLS) and TLS connections are valid when an Autonomous Database instance is configured to allow TLS connections. See Connect Node.js and other Scripting Languages (mTLS) for information on connecting using mutual TLS (mTLS) with a wallet.

Connect Power BI and Microsoft Data Tools to Autonomous Database

Oracle Client for Microsoft Tools (OCMT) is a graphical user interface (GUI) native Microsoft Software installer (MSI) that simplifies ODP.NET setup and Oracle database connectivity to multiple Microsoft data tools.

OCMT supports the following Microsoft tools:

- Power BI Desktop (unmanaged ODP.NET)
- Power BI service (unmanaged ODP.NET)
- Excel (unmanaged ODP.NET)
- SQL Server Analysis Services (unmanaged ODP.NET)
- BizTalk Server (unmanaged ODP.NET)
- SQL Server Data Tools (managed ODP.NET)
- SQL Server Integration Services (managed ODP.NET)
- SQL Server Reporting Services (managed ODP.NET)

To begin with, you must download this tool from the Oracle Client for Microsoft Tools page by clicking the **64-bit Oracle Client for Microsoft Tools** link. Save the **Oracle-Client-for-Microsoft-Tools.exe** file to your Windows machine and double-click the file to run the GUI installer.

By default, OCMT installs with both the managed and unmanaged ODP.NET providers. However, you can choose to have only managed ODP.NET or unmanaged ODP.NET driver while installing the Oracle client. The ODP.NET driver type each Microsoft data tool uses is shown in parenthesis in the above list.

Managed ODP.NET works with both 64-bit and 32-bit MS tool runtimes. In the case of unmanaged ODP.NET provider, OCMT installs the 64-bit driver. Thus, it supports 64-bit tools, such as 64-bit Power BI Desktop and 64-bit Excel, but not their 32-bit counterparts. 32-bit MS tools can continue using 32-bit Oracle Data Access Client (ODAC) 19c xcopy installs to set up 32-bit unmanaged ODP.NET as before.

Irrespective of the Microsoft tool that you want to connect to your Autonomous Database, you must follow a sequence of steps to use OCMT. They are:

- Provision an Autonomous Database. Refer to Provision Autonomous Database for instructions.
- Download database credentials to the Windows client. Refer to Download Client Credentials (Wallets) for instructions.

- Depending on the Microsoft (MS) tool that you want to connect to the Autonomous Database, you can install any of the following MS tools to the Windows client:
  - Visual Studio and Microsoft Reporting Services extension
  - Visual Studio and Microsoft Analysis Services extension
  - Visual Studio with the SQL Server Integration Services extension, which also supports SQL Server Data Tools
  - Power BI On-premises Data gateway
  - Power BI Desktop
- Install and configure ODP.NET on the Windows client
- Validate the MS tool connection to Autonomous Database

Below are step-by-step connectivity tutorials for each of the MS tools listed above. These tutorials show how to connect these MS tools with either on-premises Oracle databases or Oracle Autonomous Database:

- Power BI Desktop: Connect to Oracle Database
- Power BI Service: Connect to Oracle Database
- SQL Server Data Tools and Integration Services: Connect to Oracle Database
- SQL Server Analysis Services: Connect to Oracle Database
- SQL Server Reporting Services: Connect to Oracle Database

Connect Python Applications to Autonomous Database

You can connect Python applications to your Autonomous Database instance either with a wallet (mTLS) or without a wallet (TLS).

- Install Python and the python-oracledb Driver
  To connect to Autonomous Database from your Python application, install Python and the python-oracledb driver.
- Connect Python Applications Without a Wallet (TLS)
- Connect Python Applications with a Wallet (mTLS)

Install Python and the python-oracledb Driver

To connect to Autonomous Database from your Python application, install Python and the python-oracledb driver.

1. Install Python 3, if it is not already available.

   The version of Python you use depends on your client-side OS and hardware. For example Windows, Linux, macOS, and others.

   > **Note:**
   >
   > Oracle recommends you keep up to date with Python and python-oracledb driver releases.

2. Install the python-oracledb driver from PyPI.

The python-oracledb driver is a Python programming language extension module allowing Python programs to connect to Oracle Database. The python-oracledb driver is the renamed, new major release of the popular cx_Oracle driver.

Supported python-oracledb driver versions: python-oracledb 1.0 (or later)

Run the following command to upgrade python:

```
python -m pip install oracledb --upgrade
```

You should see output similar to the following:

```
Collecting oracledb
  Downloading oracledb-1.0.3-cp310-cp310-win_amd64.whl (1.0 MB)
     ---------------------------------------- 1.0/1.0 MB 1.8 MB/s eta
0:00:00
Collecting cryptography>=3.4
  Downloading cryptography-37.0.4-cp36-abi3-win_amd64.whl (2.4 MB)
     ---------------------------------------- 2.4/2.4 MB 3.5 MB/s eta
0:00:00
Collecting cffi>=1.12
  Downloading cffi-1.15.1-cp310-cp310-win_amd64.whl (179 kB)
     ---------------------------------------- 179.1/179.1 kB 5.4 MB/s eta
0:00:00
Collecting pycparser
  Downloading pycparser-2.21-py2.py3-none-any.whl (118 kB)
     ---------------------------------------- 118.7/118.7 kB 7.2 MB/s eta
0:00:00
Installing collected packages: pycparser, cffi, cryptography, oracledb
Successfully installed cffi-1.15.1 cryptography-37.0.4 oracledb-1.0.3
pycparser-2.21
```

Notes for installing python-oracledb:

- If you are behind a proxy, use the `--proxy` option to add a proxy server to the command. For example:

  ```
  python -m pip install oracledb --upgrade --proxy=http://
  proxy.example.com:80
  ```

- In the case where you do not have permission to write to system directories, include the `--user` option. For example:

  ```
  python -m pip install oracledb --upgrade --user
  ```

- If a binary package is not available for your platform, running `pip` will download the source package instead. The source is compiled and the resulting binary is installed.

See Installing python-oracledb for additional options and tips.

3. If you want to use the python-oracledb driver in Thick mode, install Oracle Client software.

   By default, python-oracledb runs in Thin mode which connects directly to Oracle Database. Thin mode does not require Oracle Client libraries. However, some additional functionality is available when python-oracledb runs in Thick mode.

> **Note:**
>
> See Oracle Database Features Supported by python-oracledb for information on supported features in python-oracledb Thin and Thick modes. Not all of the features shown in this link are available with Autonomous Database.

Python-oracledb uses Thick mode when you use either the Oracle Instant client libraries or the Oracle Database Client libraries and you call `oracledb.init_oracle_client()` in your Python code.

When you install Oracle Client Software, there are differences in required minimum versions for mTLS and TLS connections, as follows:

- **Mutual TLS (mTLS) Connections**:
  - If your database is on a remote computer, then download the free Oracle Instant Client “Basic” or “Basic Light” package for your operating system architecture. Use a supported version: **Oracle Instant Client**: 18.19 (or later), 19.2 (or later), or 21 (base release or later).
  - Alternatively, you can use the Full Oracle Database client libraries when they are available on your system (including Full Oracle Database Client: **Oracle Database Client**: 18.19 (or later), 19.2 (or later), or 21 (base release or later).
- **TLS Connections**: Oracle Call Interface (OCI) clients support TLS authentication without a wallet if you are using the following client versions:
  - Oracle Instant Client/Oracle Database Client 19.14 (or later) and 21.5 (or later) - only on Linux x64 and Windows
  - Alternatively, you can use the Full Oracle Database client libraries when they are available on your system, including Full Oracle Database Client 19.14 (or later) and 21.5 (or later).

Connect Python Applications Without a Wallet (TLS)

You can connect Python applications to your Autonomous Database instance without a wallet. Connecting a Python application without a wallet (TLS) provides security for authentication and encryption, and security is enforced using client credentials (by providing a username and password).
The python-oracledb driver's default "Thin mode" connects directly to Oracle Database. It can optionally use Oracle Client libraries, "Thick mode", for some additional functionality. The Oracle Client libraries can be from Oracle Instant Client, the full Oracle Client, or an Oracle Database installation.

Follow these steps to connect your Python application to an Autonomous Database instance without a wallet (TLS):

1. Install Python and the python-oracledb Driver
2. Enable TLS on Autonomous Database and Obtain Connection String
3. Run Python Application Without a Wallet

See Enabling python-oracledb Thick mode for information on Thick mode.

Topics

- Enable TLS on Autonomous Database and Obtain Connection String

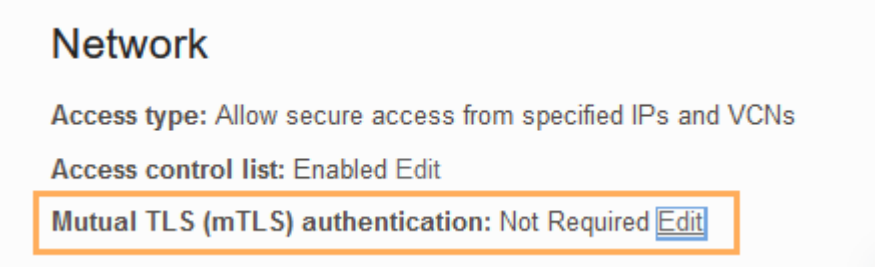
- Run Python Application Without a Wallet
  A Python application can connect to your Autonomous Database instance without a wallet (TLS) using the database credentials and a connect descriptor.

Enable TLS on Autonomous Database and Obtain Connection String

To run a Python application without a wallet, enable the Autonomous Database instance for TLS connections and obtain a connection string to contact the database from the Python application.

1. Determine if your Autonomous Database instance is enabled for TLS connections.

   If the instance is enabled for TLS connections, in the **Network** area on the Oracle Cloud Infrastructure Console the **Mutual TLS (mTLS) authentication** field shows: **Not Required**:

   If your instance requires Mutual TLS authentication, allow TLS connections on your Autonomous Database instance. See Update your Autonomous Database Instance to Allow both TLS and mTLS Authentication for details.

2. Obtain an Autonomous Database service connection string to access the database as follows:

   a. On the Oracle Cloud Infrastructure Console, click **Database connection**.

   b. Select **TLS** in the Database Connection dialog box, under **Connection Strings**, in the **TLS Authentication** drop-down list.

   > **Note:**
   >
   > You must select **TLS** in the **TLS Authentication** drop-down to obtain the TLS connection strings before you copy a connection string (when the value is **Mutual TLS** the connection strings have different values and do not work with TLS connections).

   c. Copy the Connection String for the database service you want to use with your application.

   See View TNS Names and Connection Strings for an Autonomous Database Instance for more information.

Run Python Application Without a Wallet

A Python application can connect to your Autonomous Database instance without a wallet (TLS) using the database credentials and a connect descriptor.

1. Obtain the connection string, as described in Enable TLS on Autonomous Database and Obtain Connection String.

2. In your Python application, set the following connection parameters to connect to an Autonomous Database instance:

- `dsn`: Use the connection string to specify the desired database service name.
- `password`: Specifies the database user password.
- `user`: Specifies the database user.

For example:

```
cs='''(description = (retry_count=20)(retry_delay=3)
(address=(protocol=tcps)
     (port=1521)(host=xxx.oraclecloud.com))
     (connect_data=(service_name=xxx.adb.oraclecloud.com))
     (security=(ssl_server_dn_match=yes)))'''

connection=oracledb.connect(
     user="admin",
     password=password,
     dsn=cs)
```

3. If you want to connect in Thick mode, include `oracledb.init_oracle_client()` in your Python application.

   For example:

```
cs='''(description = (retry_count=20)(retry_delay=3)
(address=(protocol=tcps)
     (port=1521)(host=xxx.oraclecloud.com))
     (connect_data=(service_name=xxx.adb.oraclecloud.com))
     (security=(ssl_server_dn_match=yes)))'''

oracledb.init_oracle_client()
connection=oracledb.connect(
     user="admin",
     password=password,
     dsn=cs)
```

Connect Python Applications with a Wallet (mTLS)

You can connect Python applications to your Autonomous Database instance with a wallet.

Connecting a Python application with a wallet (mTLS) provides enhanced security for authentication and encryption, and security is enforced using client credentials (by providing a username and password).

The python-oracledb driver's default "Thin mode" connects directly to Oracle Database. The driver can optionally use Oracle Client libraries, "Thick mode", for some additional functionality. The Oracle Client libraries can be from Oracle Instant Client, the full Oracle Client, or from an Oracle Database installation.

Follow these steps to connect your Python application to an Autonomous Database instance using a wallet (mTLS):

1. Install Python and the python-oracledb Driver
2. Obtain Security Credentials (Oracle Wallet) and Enable Network Connectivity
3. Perform this step if you only want to connect in Thin mode: Run Python Application with python-oracledb Thin Mode with a Wallet (mTLS)

4. Perform this step if you want to connect in Thick mode: Run Python Application with python-oracledb Thick Mode with a Wallet (mTLS)

Topics

- Obtain Security Credentials (Oracle Wallet) and Enable Network Connectivity
  Obtain client security credentials to connect to an Autonomous Database instance.
- Run Python Application with python-oracledb Thin Mode with a Wallet (mTLS)
  By default, python-oracledb uses Thin mode to connect directly to your Autonomous Database instance.
- Run Python Application with python-oracledb Thick Mode with a Wallet (mTLS)

Obtain Security Credentials (Oracle Wallet) and Enable Network Connectivity

Obtain client security credentials to connect to an Autonomous Database instance.

1. Download a wallet file from the Autonomous Database instance to obtain a zip file that contains the client security credentials and network configuration settings required to access an Autonomous Database instance.

   Obtain the client security credentials (`wallet.zip` file):

   - ADMIN user: On the Oracle Cloud Infrastructure Console, click **Database connection**. See Download Client Credentials (Wallets).
   - Other user (non-administrator): Obtain the Oracle Wallet from the administrator for your Autonomous Database instance.

   > **Note:**
   >
   > Protect the `wallet.zip` file and its contents to prevent unauthorized database access.

2. Unzip the client credentials file (`wallet.zip`).

Run Python Application with python-oracledb Thin Mode with a Wallet (mTLS)

By default, python-oracledb uses Thin mode to connect directly to your Autonomous Database instance.

In Thin mode only two files from the wallet zip are needed:

- `tnsnames.ora`: Maps net service names used for application connection strings to your database services.
- `ewallet.pem`: Enables SSL/TLS connections in Thin mode.

To connect in Thin mode:

1. Move `tnsnames.ora` and `ewallet.pem` files to a location on your system.

   - Linux
   - Windows

Linux

For example on Linux:

```
/opt/OracleCloud/MYDB
```

Windows

For example on Windows:

```
C:\opt\OracleCloud\MYDB
```

2. In your Python application, set the following connection parameters to connect to an Autonomous Database instance:
   - `config_dir`: Specifies the directory containing `tnsnames.ora`.
   - `dsn`: Use to specify the desired network alias from the `tnsnames.ora` file.
   - `password`: Specifies the database user password.
   - `user`: Specifies the database user.
   - `wallet_location`: Specifies the directory containing the PEM file (`ewallet.pem`).
   - `wallet_password`: Specifies the password for the PEM file (`ewallet.pem`). You set this password when you download the `wallet.zip` file.

- Linux
- Windows

Linux

For example, on Linux to connect as the ADMIN user using `oracledb.connect` with the `db2024_low` network service name (the service name is found in `tnsnames.ora`):

```
connection=oracledb.connect(
     config_dir="/opt/OracleCloud/MYDB",
     user="admin",
     password=password,
     dsn="db2024_low",
     wallet_location="/opt/OracleCloud/MYDB",
     wallet_password=wallet_pw)
```

Windows

For example, on Windows to connect as the ADMIN user using `oracledb.connect` with the `db2024_low` network service name (the service name is found in `tnsnames.ora`):

```
connection=oracledb.connect(
     config_dir=r"C:\opt\OracleCloud\MYDB",
     user="admin",
     password=password,
```

```
    dsn="db2024_low",
    wallet_location=r"C:\opt\OracleCloud\MYDB",
    wallet_password=wallet_pw)
```

The use of a 'raw' string `r"..."` means that backslashes are treated as directory separators.

As shown in this example, `wallet_location` and `config_dir` are set to the same directory (and the directory contains `tnsnames.ora` and `ewallet.pem`). Specifying the same directory for these files is not required.

If you are behind a firewall, you can tunnel TLS/SSL connections through a proxy using HTTPS_PROXY in the connect descriptor or by setting connection attributes. Successful connection depends on specific proxy configurations. Oracle does not recommend using a proxy in a production environment, due to the possible impact on performance.

In Thin mode you can specify a proxy by adding the `https_proxy` and `http_proxy_port` parameters.

For example, on Linux:

```
connection=oracledb.connect(
     config_dir="/opt/OracleCloud/MYDB",
     user="admin",
     password=password,
     dsn="db2024_low",
     wallet_location="/opt/OracleCloud/MYDB",
     wallet_password=wallet_pw,
     https_proxy='myproxy.example.com',
     https_proxy_port=80)
```

For example, on Windows:

```
connection=oracledb.connect(
     config_dir=r"C:\opt\OracleCloud\MYDB",
     user="admin",
     password=password,
     dsn="db2024_low",
     wallet_location=r"C:\opt\OracleCloud\MYDB",
     wallet_password=wallet_pw,
     https_proxy='myproxy.example.com',
     https_proxy_port=80)
```

Run Python Application with python-oracledb Thick Mode with a Wallet (mTLS)

By default, python-oracledb runs in Thin mode which connects directly to Oracle Database. Additional python-oracledb features are available when the driver runs in Thick mode.

> **Note:**
>
> Thick mode requires that the Oracle Client libraries are installed where you run Python. You must also call `oracledb.init_oracle_client()` in your Python code.

In Thick mode the following three files from the wallet zip file are required:

- `tnsnames.ora`: Contains the net service names used for application connection strings and maps the strings to your database services.
- `sqlnet.ora`: Specifies the SQL*Net client side configuration.
- `cwallet.sso`: Contains the auto-open SSO wallet.

To connect in Thick mode:

1. Place the files `tnsnames.ora`, `sqlnet.ora`, and `cwallet.sso` on your system.

   Use one of two options to place these files on your system:

   - If you are using Instant Client, move the files to a `network/admin` subdirectory hierarchy under the Instant Client directory. For example depending on the architecture or your client system and where you installed Instant Client, the files should placed be in a directory location such as:

     ```
     /home/myuser/instantclient_19_21/network/admin
     ```

     or

     ```
     /usr/lib/oracle/19.21/client64/lib/network/admin
     ```

     For example, on Linux if you are using the full Oracle Client move the files to `$ORACLE_HOME/network/admin`.

   - Alternatively, move the files to any accessible directory.

     For example, on Linux move the files to the directory `/opt/OracleCloud/MYDB` and edit `sqlnet.ora` to change the wallet location directory to the directory containing the `cwallet.sso` file.

     For example, on Linux edit `sqlnet.ora` as follows:

     ```
     WALLET_LOCATION = (SOURCE = (METHOD=file) (METHOD_DATA =
     (DIRECTORY="/opt/OracleCloud/MYDB")))
     SSL_SERVER_DN_MATCH=yes
     ```

     When the configuration files are not in the default location, your application needs to indicate where they are, either with the `config_dir` parameter in the call `oracledb.init_oracle_client()` or by setting the `TNS_ADMIN` environment variable.

     > **Note:**
     >
     > Neither of these settings are needed, and you do not need to edit `sqlnet.ora` if you put all the configuration files in the `network/admin` directory.

2. In your Python application set the following initialization and connection parameters to connect to the Autonomous Database instance:

   - `config_dir`: Specifies the configuration directory when you are putting the configuration files. This is only required when the configuration files are placed in a directory outside of the instant client configuration directory `network/admin`.

- `dsn`: Specifies the desired network alias from the `tnsnames.ora` file.
- `password`: Specifies the database user password.
- `user`: Specifies the database user.

In the first case for placement of the configuration files, connect to the Autonomous Database instance using your database credentials by setting the `dsn` parameter to the desired network alias from `tnsnames.ora`.

For example, to connect as the ADMIN user using `oracledb.init_oracle_client` and connect with the `db2024_low` network service name (where the service name is found in `tnsnames.ora`):

```
oracledb.init_oracle_client()
   connection=oracledb.connect(
       user="admin",
       password=password,
       dsn="db2024_low")
```

When configuration files are in a directory outside of the instant client configuration directory, set the `config_dir` parameter when you call `oracledb.init_oracle_client`.

- Linux
- Windows

Linux

For example, on Linux to connect as the ADMIN user using the `db2024_low` network service name:

```
oracledb.init_oracle_client(config_dir="/opt/OracleCloud/MYDB")
   connection=oracledb.connect(
      user="admin",
      password=password,
      dsn="db2024_low")
```

Windows

For example, on Windows to connect as the ADMIN user using the `db2024_low` network service name:

```
oracledb.init_oracle_client(config_dir=r"C:\opt\OracleCloud\MYDB")
   connection=oracledb.connect(
      user="admin",
      password=password,
      dsn="db2024_low")
```

The use of a ‘raw’ string `r"..."` means that backslashes are treated as directory separators.

If you are behind a firewall, you can tunnel TLS/SSL connections through a proxy using HTTPS_PROXY in the connect descriptor or by setting connection attributes. Successful connection depends on specific proxy configurations. Oracle does not recommend using a proxy in a production environment, due to the possible impact on performance.

In Thick mode you can specify a proxy by editing the `sqlnet.ora` file and adding a line:

```
SQLNET.USE_HTTPS_PROXY=on
```

In addition, edit `tnsnames.ora` and add an `HTTPS_PROXY` proxy name and `HTTPS_PROXY_PORT` port to the connect descriptor address list of any service name you plan to use.

For example:

```
mydb_high=(description=
(address=(https_proxy=myproxy.example.com)
(https_proxy_port=80)
(protocol=tcps)(port=1522)(host=...)
```

See Enabling python-oracledb Thick mode for information on Thick mode.

Connect Go Applications to Autonomous Database

You can connect Go applications to your Autonomous Database instance either with a wallet (mTLS) or without a wallet (TLS).

The Go language is an open source language and has support for SQL interactions. You can connect any Go program to your Autonomous Database using the Godror driver. The `godror` driver requires Oracle Client libraries such as from Oracle Instant Client.

> **Note:**
>
> The `godror` driver is a third party driver. It is not an Oracle driver.

- Install Go and the Godror driver
  To connect to Autonomous Database from your Go application, install Go and the `godror` driver.
- Connect Go Applications Without a Wallet (TLS) Using TLS
  You can connect Go applications to your Autonomous Database instance without a wallet using TLS.
- Connect Go Applications with a Wallet (mTLS)
  You can connect Go applications to your Autonomous Database instance with a wallet.

Install Go and the Godror driver

To connect to Autonomous Database from your Go application, install Go and the `godror` driver.

1. Download and install Go based on your operating system and the version of the operating system. This method installs from precompiled binary packages.

> **Note:**
>
> These steps describe installation of Go application for Microsoft Windows operating system. For Linux and Mac operating system, you can check out the link mentioned in this step.

2. Open the MSI file you downloaded and follow the prompts to download Go. The installer will install the `godror` driver.
3. Go to Program Files (x86).

   After installing, close and reopen any open command prompts so that changes to the environment made by the installer are reflected at the command prompt.
4. Verify that you've installed Go.
   - In **Windows**, click the **Start** menu.
   - In the menu's search box, type `cmd`, then press the **Enter** key.
   - In the Command Prompt window that appears, type the following command:

     ```
     go version
     ```

   - Confirm that the command prints the installed version of Go.
5. Install `godror`:

   To install `godror`, run the following command:

   ```
   go get github.com/godror/godror
   ```

   `Godror` is an open source database package in Go that allows a Go application user to connect and work with Autonomous Database using ODPI-C.

   See Go DRiver for ORacle User Guide (godror) for additional tips on installation.
6. Install an Oracle Client library to use `Godror`.

   `Godror` uses Oracle Database Programming Interface for C (ODPI-C) which supports connection to Oracle databases.
   Follow the ODPI-C installation instructions to install and configure the Oracle client libraries.

   You can use the optional `libDir` parameter to specify the Oracle Client library directory on macOS and Windows. For example,

   ```
   db, err := sql.Open("godror", `user="scott" password="tiger"
   connectString="dbhost:1521/orclpdb1"
                                 libDir="/Users/myname/instantclient_19_3"`)
   ```

   In this example, the `libDir` parameter value is set to the path to the location where you unzip the Oracle Instant Client zip files.

> **Note:**
>
> Oracle recommends you keep up to date with Go and `godror` releases.

Connect Go Applications Without a Wallet (TLS) Using TLS

You can connect Go applications to your Autonomous Database instance without a wallet using TLS.

Connecting a Go application without a wallet (TLS) provides security for authentication and encryption, and security is enforced using client credentials (by providing a username and password).

Topics

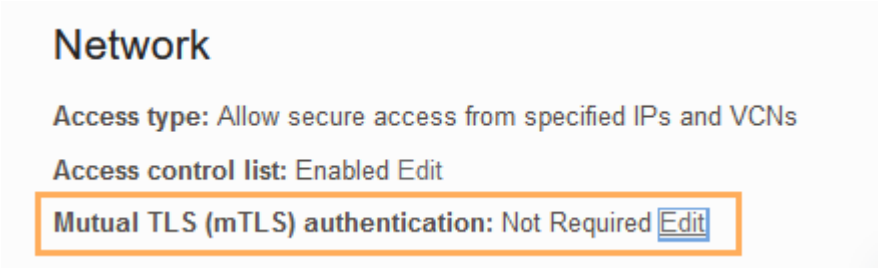
- Enable TLS on Autonomous Database and Obtain Connection String
  To run a Go application without a wallet using TLS, enable the Autonomous Database instance for TLS connections and obtain a connection string to contact the database from the Go application.
- Run Go Application Without a Wallet Using TLS
  The following section describes the steps to run a Go application on an Autonomous Database instance.

Enable TLS on Autonomous Database and Obtain Connection String

To run a Go application without a wallet using TLS, enable the Autonomous Database instance for TLS connections and obtain a connection string to contact the database from the Go application.

1. Determine if your Autonomous Database instance is enabled for TLS connections.

   If the instance is enabled for TLS connections, in the **Network** area on the Console the **Mutual TLS (mTLS) authentication** field shows: **Not Required**:

   If your instance requires Mutual TLS authentication, allow TLS connections on your Autonomous Database instance. See Update your Autonomous Database Instance to Allow both TLS and mTLS Authentication for details.

2. Obtain an Autonomous Database service connection string to access the database as follows:

   a. On the Oracle Cloud Infrastructure Console, click **Database connection**.

   b. Select **TLS** in the Database Connection dialog box, under **Connection Strings**, in the **TLS Authentication** drop-down list.

> **Note:**
>
> You must select **TLS** in the **TLS Authentication** drop-down to obtain the TLS connection strings before you copy a connection string (when the value is **Mutual TLS** the connection strings have different values and do not work with TLS connections).

c. Copy the connection string for the database service you want to use with your application.

   See View TNS Names and Connection Strings for an Autonomous Database Instance for more information.

Run Go Application Without a Wallet Using TLS

The following section describes the steps to run a Go application on an Autonomous Database instance.

1. Obtain the connection string, as described in Enable TLS on Autonomous Database and Obtain Connection String.
2. In your Go application, set the following parameters to connect to an Autonomous Database instance:
   - `password`: Specifies the database user password.
   - `user`: Specifies the database user.
   - `libDir`: Single directory where you unzip the Oracle Instant Client zip package.
   - `connectString`: Use the connection string you obtained from the previous step to connect using the desired database service name.

   To run a Go application using the `godror` driver, you must first import the `godror` package.

   Following is a sample code in the Go Application which enables you to connect to an Autonomous Database:

```
package main

import (
    "database/sql"
    "fmt"

    _ "github.com/godror/godror"
)

func main() {
    connectToADb()
}

func connectToADb() {
    fmt.Println("Connecting to Oracle Autonomous database !!!")
    db, err := sql.Open("godror", `user="admin" password="password"
    connectString="tcps://adb.ap-mumbai-1.oraclecloud.com:1522/
abcd_xxxx_medium.adb.oraclecloud.com?wallet_location=/Users/user/tool/
wallet_DB/"
    libDir="/home/user/tool/instantclient_19_8/"`)
```

```
    if err != nil {
        fmt.Println(err)
        return
    }
    defer db.Close()

    rows, err := db.Query("select 'hello' from dual")
    if err != nil {
        fmt.Println("Error running query")
        fmt.Println(err)
        return
    }
    defer rows.Close()

    var resData string
    for rows.Next() {

        rows.Scan(&resData)
    }
    fmt.Printf("The response is: %s\n", resData)
}
```

This code will return records retrieved from the Autonomous Database.

Connect Go Applications with a Wallet (mTLS)

You can connect Go applications to your Autonomous Database instance with a wallet.

Topics

- Download an Oracle Wallet File
  You must download an Oracle Wallet (mTLS) file to obtain client security credentials that enable you to connect to an Autonomous Database instance.
- Run Go Applications with a Wallet (mTLS)
  You can connect Go applications to your Autonomous Database instance with a wallet. You must download a zipped wallet file from the Autonomous Database instance that contains the client security credentials.

Download an Oracle Wallet File

You must download an Oracle Wallet (mTLS) file to obtain client security credentials that enable you to connect to an Autonomous Database instance.

1. Download a wallet file from the Autonomous Database instance to obtain a zip file that contains the client security credentials and network configuration settings required to access an Autonomous Database instance.

   Depending on whether you are an Admin user or a non-Admin user, obtain the client security credentials (wallet.zip file) as follows:

   - ADMIN user: On the Oracle Cloud Infrastructure Console, click **Database connection**. See Download Client Credentials (Wallets).
   - Other user (non-administrator): Obtain the Oracle Wallet from the administrator for your Autonomous Database instance.

> **Note:**
>
> Protect the `wallet.zip` file and its contents to prevent unauthorized database access.

2. Unzip the client credentials file (`wallet.zip`).

Run Go Applications with a Wallet (mTLS)

You can connect Go applications to your Autonomous Database instance with a wallet. You must download a zipped wallet file from the Autonomous Database instance that contains the client security credentials.

The `godror` driver requires Oracle Client Libraries. See Install Oracle Client Library to use Godror. The Oracle Client libraries are installed where you run the Go Application.

The wallet file is in a zip archive format. The database service names can be found in the `tnsnames.ora` file in the wallet zip file. In the example given below, you will use the `mydb_high` network service name found in the `tnsnames.ora` file.

The following three files from the wallet zip file are required:

- `tnsnames.ora`: Contains the net service names used for application connection strings and maps the strings to your database services.
- `sqlnet.ora`: Specifies the SQLNet client side configuration.
- `cwallet.sso`: Contains the auto-open SSO wallet.

Follow these steps to connect your Go application to an Autonomous Database instance using a wallet (mTLS):

1. Place the files `tnsnames.ora`, `sqlnet.ora`, and `cwallet.sso` on the optional (`opt`) directory of your Linux system. The /`opt` directory is used to install optional or add-on software packages that are not part of the core operating system. You can place the files anywhere on your Microsoft Windows operating system. For example, subdirectory on C drive (`C:`) of your Microsoft Windows operating system.

   You can move the files to any accessible directory.

   For example, on Linux move the files to the directory `/opt/OracleCloud/MYDB` and edit `sqlnet.ora` to change the wallet location directory to the directory containing the `cwallet.sso` file.

   For example, on Linux edit `sqlnet.ora` as follows:

   ```
   WALLET_LOCATION = (SOURCE = (METHOD=file) (METHOD_DATA = (DIRECTORY="/opt/
   OracleCloud/MYDB")))
   SSL_SERVER_DN_MATCH=yes
   ```

   When the configuration files are not in the default location, your Go application needs to indicate where they are with the `config_dir` parameter.

> **Note:**
>
> Neither of these settings are needed, and you do not need to edit `sqlnet.ora` if you put all the configuration files in the `network/admin` directory.

2. In your Go application set the following connection parameters to connect to the Autonomous Database instance:
   - `user`: Specifies the Cloud Database user.
   - `password`: Specifies the Cloud Database user's password.
   - `config_dir`: Specifies the configuration directory where you are placing the `tnsnames.ora` and `sqlnet.ora` files. This is required when the configuration files are placed in a directory outside of the instant client configuration directory `network/admin`.
   - `dsn`: Specifies the desired network alias from the `tnsnames.ora` file.

   For example, to connect as the ADMIN user and connect with the `mydb_high` network service name (where the service name is found in `/opt/OracleCloud/MYDB/tnsnames.ora`):

```
dsn := `user="admin"
      password="password"
      connectString="mydb_high"
      configDir="/opt/OracleCloud/MYDB"`
db, _ := sql.Open("godror", dsn)
```

   You can use the above code in your Go application to create a connection string and connect to the Autonomous Database.

Download Database Connection Information

You can download Oracle client credentials, wallet files, from an Autonomous Database instance.

If you are not an Autonomous Database administrator and your application requires a wallet to connect, then your administrator should provide you with the client credentials. You can also view TNS names and connection strings for your database.

- Download Client Credentials (Wallets)
  To download client credentials you can use the Oracle Cloud Infrastructure Console or Database Actions.
- Wallet README File
  The wallet `README` file contains the wallet expiration information and details for Autonomous Database tools and resources.
- View TNS Names and Connection Strings for an Autonomous Database Instance
  From the Database Connection page on the Oracle Cloud Infrastructure Console you can view Autonomous Database TNS names and connection strings.

Download Client Credentials (Wallets)

To download client credentials you can use the Oracle Cloud Infrastructure Console or Database Actions.

> **Note:**
>
> The password you provide when you download the wallet protects the downloaded Client Credentials wallet.

For commercial regions, the wallet password complexity for the password you supply requires the following:

- Minimum of 8 characters
- Minimum of 1 letter
- Minimum of 1 numeric character or 1 special character

For US Government regions, the wallet password complexity requires all of the following:

- Minimum of 15 characters
- Minimum of 1 lowercase letter
- Minimum of 1 uppercase letter
- Minimum of 1 numeric character
- Minimum 1 special character

To download client credentials from the Oracle Cloud Infrastructure Console:

1. Navigate to the Autonomous Database details page.
2. Click **Database connection**.
3. On the **Database connection** page select the **Wallet type**:
   - **Instance wallet**: Wallet for a single database only; this provides a database-specific wallet.
   - **Regional wallet**: Wallet for all Autonomous Databases for a given tenant and region (this includes all service instances that a cloud account owns).

   > **Note:**
   >
   > Oracle recommends you provide a database-specific wallet, using **Instance wallet**, to end users and for application use whenever possible. Regional wallets should only be used for administrative purposes that require potential access to all Autonomous Databases within a region.
4. Click **Download wallet**.
5. In the **Download wallet** dialog, enter a wallet password in the **Password** field and confirm the password in the **Confirm password** field.
6. Click **Download** to save the client security credentials zip file.

By default the filename is: `Wallet_`*databasename*`.zip`. You can save this file as any filename you want.

You must protect this file to prevent unauthorized database access.

To download client credentials from Database Actions:

First, access Database Actions as the ADMIN user. See Access Database Actions as ADMIN for more information.

1. Access Database Actions as the ADMIN user. See Access Database Actions as ADMIN for more information.
2. On the Database Actions Launchpad, under **Administration**, select **Download Client Credentials (Wallet)**.
3. On the **Download Client Credentials (Wallet)** page, enter a wallet password in the **Password** field and confirm the password in the **Confirm Password** field.
4. Click **Download** to save the client security credentials zip file. By default the filename is: `Wallet_`*databasename*`.zip`. You can save this file as any filename you want. You must protect this file to prevent unauthorized database access.

> **Note:**
>
> When you use Database Actions to download a wallet there is no **Wallet type** option on the **Download Client Credentials (Wallet)** page and you always download an instance wallet. If you need to download the regional wallet click **Database connection** on the Oracle Cloud Infrastructure Console.

The zip file includes the following:

File	Description
`cwallet.sso`	Auto-open SSO wallet
`ewallet.p12`	PKCS12 file. The PKCS12 file is protected by the wallet password provided while downloading the wallet.
`ewallet.pem`	Encoded certificate file used to authenticate with certificate authority (CA) server certificate.
`keystore.jks`	Java keystore file. This file is protected by the wallet password provided while downloading the wallet.
`ojdbc.properties`	Contains the wallet related connection property required for JDBC connection. This should be in the same path as `tnsnames.ora`.
`README`	Contains wallet expiration information and links for Autonomous Database tools and resources. See Wallet README File for information on the contents of the `README` file.
`sqlnet.ora`	SQL*Net client side configuration.
`tnsnames.ora`	Network configuration file storing connect descriptors.

File	Description
`truststore.jks`	Java truststore file. This file is protected by the wallet password provided while downloading the wallet.

Notes for wallet files and the wallet password:

- To invalidate database client certification keys associated with a wallet, see Rotate Wallets with Immediate Rotation.
- Wallet files, along with the Database user ID and password provide access to data in your database. Store wallet files in a secure location. Share wallet files only with authorized users. If wallet files are transmitted in a way that might be accessed by unauthorized users (for example, over public email), transmit the wallet password separately and securely.
- For better security, Oracle recommends using restricted permissions on wallet files. This means setting the file permissions to 600 on Linux/Unix. Similar restrictions can be achieved on Windows by letting the file owner have Read and Write permissions while all other users have no permissions.
- Autonomous Database uses strong password complexity rules for all users based on Oracle Cloud security standards. For more information on the password complexity rules see Create Users on Autonomous Database - Connecting with a Client Tool.
- The `README` file that contains wallet expiration information is not available in wallet zip files that were downloaded before April 2020.
- Starting six weeks before the wallet expiration date Autonomous Database sends notification emails each week, indicating the wallet expiration date. These emails provide notice before your wallet expires that you need to download a new wallet. You will receive these notification emails only if there is a connection that uses a wallet that is about to expire.

  You can also use the `WalletExpirationWarning` event to be notified when a wallet is due to expire. You will receive these notification events only if you are subscribed to Critical events and there is a connection that uses a wallet that is about to expire. See About Events Based Notification and Automation on Autonomous Database for more information.

Wallet README File

The wallet `README` file contains the wallet expiration information and details for Autonomous Database tools and resources.

The wallet expiration information at the top of the `README` file shows the following information:

- The date when the wallet was downloaded.
- The date when the wallet SSL certificate provided in the wallet expires. If your wallet is nearing expiration or is expired, then download a new wallet or obtain a new wallet from your Autonomous Database administrator. If you do not download a new wallet before the expiration date, you will no longer be able to connect to your database.

The Autonomous Database tools and resources area provides the following information:

Tool or Resource	Description
Database Actions	Load, explore, transform, model, and catalog your data. Use an SQL worksheet, build REST interfaces and low-code apps, manage users and connections, build and apply machine learning models. Access Link: provides the link to use Database Actions. See Connect with Built-In Oracle Database Actions for more information.
Graph Studio	Oracle Graph Studio lets you create scalable property graph databases. Graph Studio automates the creation of graph models and in-memory graphs from database tables. It includes notebooks and developer APIs that allow you to execute graph queries using PGQL (an SQL-like graph query language) and over 50 built-in graph algorithms. Graph Studio also offers dozens of visualization, including native graph visualization. Access Link provides the link to use Graph Studio. See About Oracle Graph Studio on Autonomous Database for more information.
Oracle APEX	Oracle APEX is a low-code development platform that enables you to build scalable, secure enterprise apps that can be deployed anywhere. Access Link: provides the link to use Oracle APEX. See Access Oracle APEX Administration Services for more information.
Oracle Machine Learning User Management	Create new Oracle Machine Learning user accounts and manage the credentials for existing Oracle Machine Learning users. Access Link: provides the link to use Oracle Machine Learning User Management. See Create and Update User Accounts for Oracle Machine Learning Components on Autonomous Database for more information.
Oracle Machine Learning User Notebooks	Oracle Machine Learning notebooks provide easy access to Oracle's parallelized, scalable in-database implementations of a library of Oracle Advanced Analytics' machine learning algorithms (classification, regression, anomaly detection, clustering, associations, attribute importance, feature extraction, times series, and so on), SQL, PL/SQL and Oracle's statistical and analytical SQL functions. Access Link: provides the link to use Oracle Machine Learning User Notebooks. See Work with Oracle Machine Learning User Interface for Data Access, Analysis, and Discovery for more information.
SODA Drivers	Simple Oracle Document Access (SODA) is a set of APIs that let you work with JSON documents managed by the Oracle Database without needing to use SQL. SODA drivers are available for REST, Java, Node.js, Python, PL/SQL, and C. Access Link: provides the link to download the SODA drivers. See Work with Simple Oracle Document Access (SODA) in Autonomous Database for more information.

Notes for wallet `README` file:

- If you rename your Autonomous Database instance, the tools links change and the old links no longer work. To obtain valid tools links you must download a new Wallet zip file with an updated `README` file. The SODA drivers link is a resource link and this link does not change when you rename an instance.

- The `README` in a regional wallet does not contain the Autonomous Database tools and resources links.
- If your Autonomous Database instance is on a private endpoint and you enable the **Allow public access** option, the `README` file includes two links for each tool: one link to use from the private endpoint when you are connecting from the same VCN where the database resides, and another link to use from specific public IP addresses or from specific VCNs if those VCNs are configured to privately connect to Autonomous Database using a Service Gateway. See Use a Private Endpoint with Public Access Allowed for more information.

View TNS Names and Connection Strings for an Autonomous Database Instance

From the Database Connection page on the Oracle Cloud Infrastructure Console you can view Autonomous Database TNS names and connection strings.

> **Note:**
>
> See Update your Autonomous Database Instance to Allow both TLS and mTLS Authentication for information on allowing TLS connections.

Perform the following steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To view the TNS names and connection strings, do the following:

1. On the Autonomous Database details page, click **Database connection**.

   By default this shows the Mutual TLS connection information in a table with the TNS names and connection strings for the Autonomous Database instance.

2. When both Mutual TLS (mTLS) and TLS connections are allowed, under **TLS authentication** select **TLS** to view the TNS names and connection strings for connections with TLS authentication.

   The TNS names are the same for mTLS and TLS authentication. The connection strings differ for mTLS and TLS connections, with different port definitions. Mutual TLS (mTLS) connections use port 1522. TLS connections use either port 1521 or 1522. The port assignments of 1522 for Mutual TLS (mTLS) and either 1521 or 1522 for TLS cannot be changed.

   In the Connection String column, click **Show** to display the full value of a connection string or click **Copy** to copy a connection string.

   For example, when you click **Show** you see the full connection string.

3. When your Autonomous Database instance is on a private endpoint and you enable the **Allow public access** option, also select an access type to show or copy the private or the public endpoint information.

   The private connection strings are the connections strings you use from within the VCN where the private endpoint database resides. Use the the public connection strings to

connect from public IPs or from other allowed VCNs. See Use a Private Endpoint with Public Access Allowed for more information.

Under **Access**, select either **Private endpoint** or **Public endpoint URL**.

Connect to Autonomous Database Using Oracle Database Tools

Oracle Database Tools such as SQL Developer, SQL*Plus, and SQLcl can be used with Autonomous Database.

The following sections provide step-by-step instructions for connecting to Autonomous Database using these tools.

- Connect Oracle SQL Developer with a Wallet (mTLS)
  Oracle SQL Developer is a free integrated development environment that simplifies the development and management of Autonomous Database.
- Connect Oracle SQL Developer Without a Wallet
  Oracle SQL Developer is a free integrated development environment that simplifies the development and management of Autonomous Database. Oracle SQL Developer provides support for connecting using TLS authentication without a wallet.
- Connect SQL*Plus with a Wallet (mTLS)
  SQL*Plus is a command-line interface used to enter SQL commands. SQL*Plus connects to an Oracle database.
- Connect SQL*Plus Without a Wallet
  SQL*Plus is a command-line interface used to enter SQL commands. SQL*Plus connects to an Oracle database.
- Connect Oracle SQLcl Cloud with a Wallet (mTLS)
  SQLcl is a command-line interface used to enter SQL commands. You can use SQLcl to connect to an Autonomous Database with client credentials configured (mTLS).
- Connect Oracle SQLcl Cloud Without a Wallet
  SQLcl is a command-line interface used to enter SQL commands. You can use SQLcl to connect to an Autonomous Database with TLS authentication without a wallet.

Connect Oracle SQL Developer with a Wallet (mTLS)

Oracle SQL Developer is a free integrated development environment that simplifies the development and management of Autonomous Database.

SQL Developer can connect to Autonomous Database and contains enhancements for key Autonomous Database features. You can download the latest version of Oracle SQL Developer for your platform from the **Download** link on this page: Oracle SQL Developer.

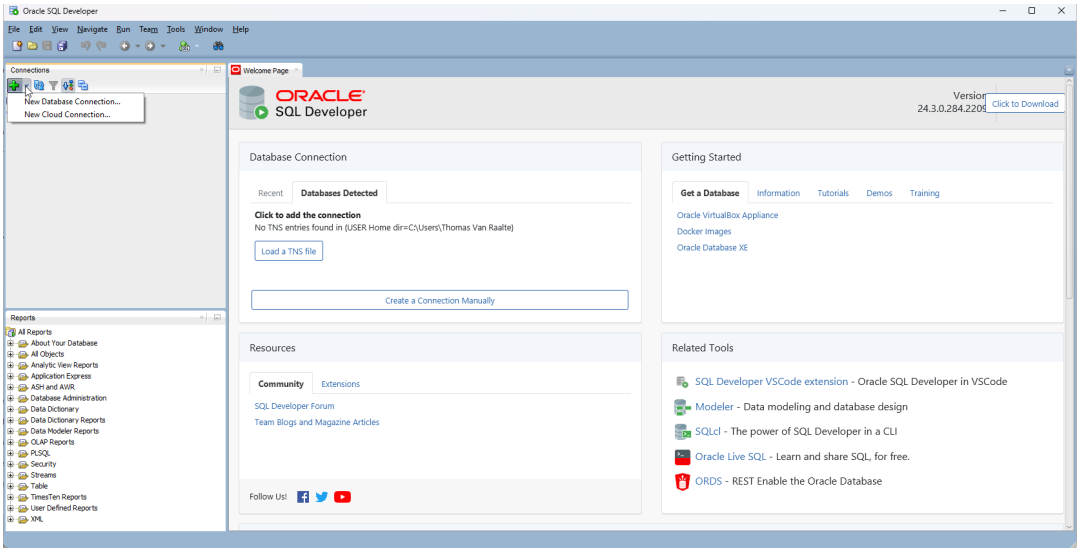
For connecting with mTLS authentication, Oracle SQL Developer provides support for wallet files using the **Cloud Wallet** Connection Type. Oracle recommends that you use version 18.2 (or later); however, earlier versions of SQL Developer will work with Autonomous Database using an Oracle Wallet.

For connecting with TLS authentication, Oracle SQL Developer provides support using the **Custom JDBC** Connection Type. See Connect with Oracle SQL Developer with TLS Authentication for details on connecting using TLS authentication.

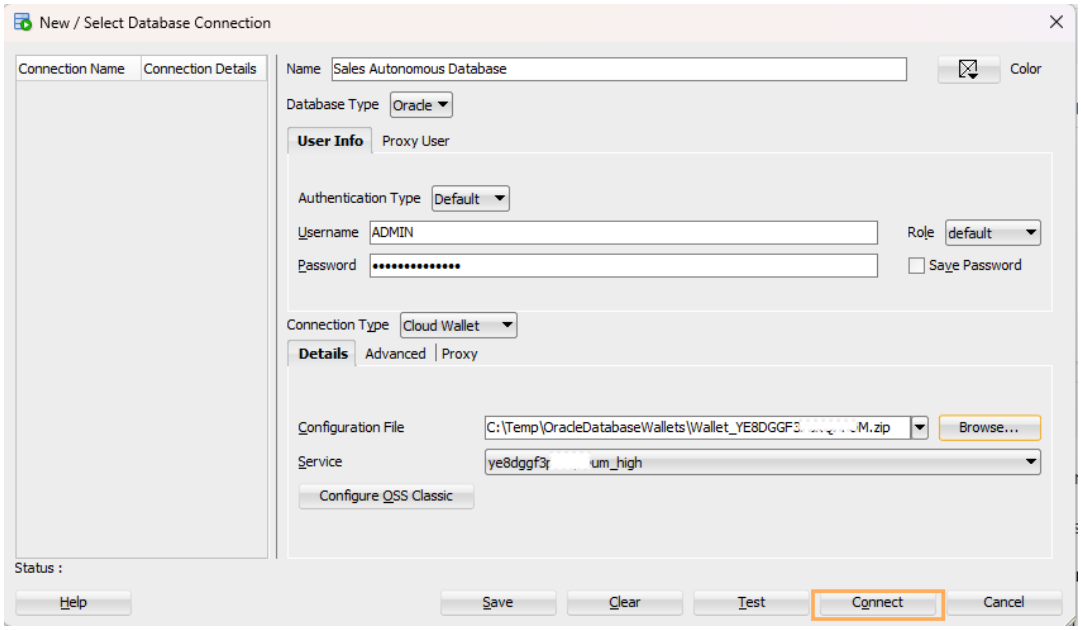
To create a new mTLS connection to Autonomous Database, do the following:

Obtain your credentials to access Autonomous Database. For more information, see Download Client Credentials (Wallets).

1. Start Oracle SQL Developer and in the connections panel, right-click **Connections** and select **New Database Connection...**.



2. Choose the Connection Type **Cloud Wallet**.
3. Enter the following information:
   - **Connection Name**: Enter the name for this connection.
   - **Username**: Enter the database username. You can either use the default administrator database account (`ADMIN`) provided as part of the service or create a new schema, and use it.
   - **Password**: Enter the password for the database user.
   - **Connection Type**: Select **Cloud Wallet**
   - **Configuration File** : Click **Browse**, and select the client credentials zip file.
   - **Service**: Enter the database TNS name. The client credentials file includes a `tnsnames.ora` file that provides database TNS names with corresponding services.



4. Click **Connect** to connect to the database.

> **Note:**
>
> If you are using Microsoft Active Directory, then for **Username** enter the Active Directory "*AD_domain\AD_username*" (you may include double quotes), and for the **Password**, enter the password for the Active Directory user. See Use Microsoft Active Directory with Autonomous Database for more information.

Connect Oracle SQL Developer Without a Wallet

Oracle SQL Developer is a free integrated development environment that simplifies the development and management of Autonomous Database. Oracle SQL Developer provides support for connecting using TLS authentication without a wallet.

> **Note:**
>
> See Update your Autonomous Database Instance to Allow both TLS and mTLS Authentication for information on allowing TLS connections.

To create a new TLS connection to Autonomous Database:

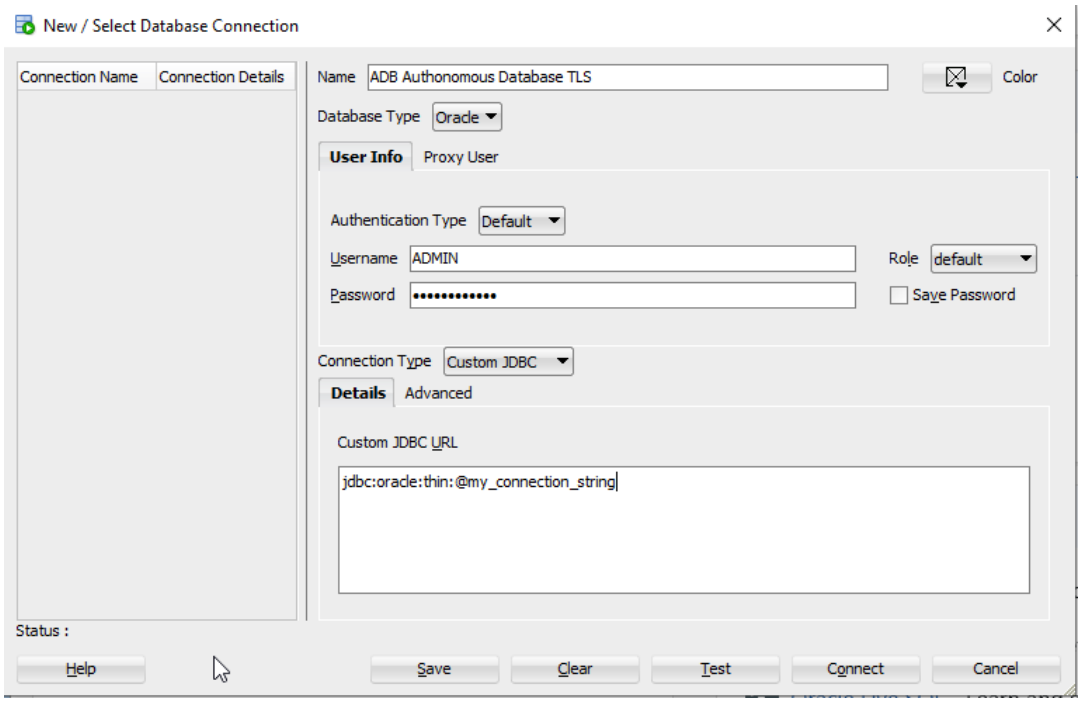
1. Copy a connection string for the Autonomous Database.

   To connect with TLS authentication copy a TLS connection string. On the Database Connection page, under **TLS Authentication**, select **TLS** to view the connection strings for connecting with TLS authentication.

   See View TNS Names and Connection Strings for an Autonomous Database Instance for information on viewing and copying connection stings.

2. Start Oracle SQL Developer and in the connections panel, right-click **Connections** and select **New Database Connection...**.
3. Choose the Connection Type **Custom JDBC**.
4. Enter the following information:
   - **Name**: Enter the name for this connection.
   - **Username**: Enter the database username. You can either use the default administrator database account `ADMIN` provided as part of the service or create a new schema, and use it.
   - **Password**: Enter the password for the database user.
   - **Connection Type**: Select **Custom JDBC**.
   - **Custom JDBC URL**: Enter the following:

     `jdbc:oracle:thin:@` followed by the connection string you copied in step one.



For example, a sample value for the Custom JDBC URL field is:

```
jdbc:oracle:thin:@(description= (retry_count=20)(retry_delay=3)
(address=(protocol=tcps)
(port=1521)(host=adb-preprod.us-phoenix-1.oraclecloud.com))
(connect_data=(service_name=u9adutfb2_fmexample1_medium.adb.oraclecloud.com
))
(security=(ssl_server_dn_match=yes)))
```

When you copy the connection string, the values for *`region`* and *`databasename`* are for your Autonomous Database instance.

5. Click **Connect** to connect to the database.

> **Note:**
>
> If you are using Microsoft Active Directory, then for **Username** enter the Active Directory "*AD_domain\AD_username*" (you may include double quotes), and for the **Password**, enter the password for the Active Directory user. See Use Microsoft Active Directory with Autonomous Database for more information.

Connect SQL*Plus with a Wallet (mTLS)

SQL*Plus is a command-line interface used to enter SQL commands. SQL*Plus connects to an Oracle database.

To install and configure the client and connect to the Autonomous Database using SQL*Plus with client credentials (mTLS), do the following:

1. Prepare for Oracle Call Interface (OCI), ODBC and JDBC OCI Connections. See Prepare for Oracle Call Interface (OCI), ODBC, and JDBC OCI Connections with Wallets (mTLS).
2. Connect using a database user, *password*, and database TNS name provided in the `tnsnames.ora` file.

   For example:

```
sqlplus adb_user@db2022adb_medium

SQL*Plus: Release 19.0.0.0.0 - Production on Mon Nov 23 15:08:48 2020
Version 19.8.0.0.0

Copyright (c) 1982, 2020, Oracle.  All rights reserved.

Enter password:
Last Successful login time: Wed Nov 18 2020 12:36:56 -08:00

Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
Version 19.5.0.0.0

SQL>
```

> **Notes:**
>
> - The Oracle Wallet is transparent to SQL*Plus because the wallet location is specified in the `sqlnet.ora` file. This is true for any Oracle Call Interface (OCI), ODBC, or JDBC OCI connection.
> - If you are connecting to a database using Microsoft Active Directory credentials, then connect using an Active Directory user name in the form of "*AD_domain\AD_username*" (double quotes must be included), and Active Directory user password. See Use Microsoft Active Directory with Autonomous Database for more information.

Connect SQL*Plus Without a Wallet

SQL*Plus is a command-line interface used to enter SQL commands. SQL*Plus connects to an Oracle database.

> **Note:**
>
> See Update your Autonomous Database Instance to Allow both TLS and mTLS Authentication for information on allowing TLS connections.

To install and configure the client and connect to the Autonomous Database using SQL*Plus with TLS authentication, do the following:

1. Prepare for Oracle Call Interface (OCI) connections.

   You can use TLS authentication without a wallet in SQL*Plus if you are using the following client versions:

   - Oracle Instant Client/Oracle Database Client 19.13 - only on Linux x64
   - Oracle Instant Client/Oracle Database Client 19.14 (or later) and 21.5 (or later) - only on Linux x64 and Windows

   See Prepare for Oracle Call Interface, ODBC, and JDBC OCI Connections Using TLS Authentication for more information.

2. Copy a connection string for the Autonomous Database.

   To connect with TLS authentication copy a TLS connection string. On the Database Connection page, under **TLS Authentication**, select **TLS** to view the connection strings for connecting with TLS authentication.

   See View TNS Names and Connection Strings for an Autonomous Database Instance for information on viewing and copying connection stings.

   See Database Service Names for Autonomous Database for information on the different databases services for each connection string.

3. Connect using a database user and *password*, and provide the connection string you copied in Step 2.

   On UNIX/Linux start `sqlplus` with the connection string, enclosed in quotes on the command line, as follows:

   ```
   sqlplus username/password@'my_connect_string'
   ```

   For example (for clarity line breaks added):

   ```
   sqlplus adb_user/password@'(description= (retry_count=20)(retry_delay=3)
   (address=(protocol=tcps)
   (port=1521)(host=adb.region.oraclecloud.com))
   (connect_data=(service_name=u9adutfb2ba8x4d_database_medium.adb.oraclecloud
   .com))
   (security=(ssl_server_dn_match=yes))'
   ```

On Windows, start `sqlplus` with the database user and password with the copied connection string, as follows (as compared to UNIX/Linux, on Windows do not surround the connection string with quotes):

```
sqlplus username/password@my_connect_string
```

For example (for clarity line breaks added in the connection string):

```
sqlplus adb_user/password@(description= (retry_count=20)(retry_delay=3)
(address=(protocol=tcps)
(port=1521)(host=adb.region.oraclecloud.com))
(connect_data=(service_name=u9adutfb2ba8x4d_database_medium.adb.oraclecloud
.com))
(security=(ssl_server_dn_match=yes))
```

> **Note:**
>
> If you are connecting to an Autonomous Database instance using Microsoft Active Directory credentials, then connect using an Active Directory user name in the form of "*AD_domain\AD_username*" (double quotes must be included), and Active Directory user password. See Use Microsoft Active Directory with Autonomous Database for more information.

Connect Oracle SQLcl Cloud with a Wallet (mTLS)

SQLcl is a command-line interface used to enter SQL commands. You can use SQLcl to connect to an Autonomous Database with client credentials configured (mTLS).

You can use SQLcl version 4.2 or later with Autonomous Database. Download SQLcl from oracle.com.

SQLcl can connect to an Autonomous Database instance using either an Oracle Call Interface (OCI) or a JDBC thin connection.

- If you use Oracle Call Interface (OCI), prepare for OCI, ODBC and JDBC OCI Connections. See Prepare for Oracle Call Interface (OCI), ODBC, and JDBC OCI Connections.
- If you use JDBC Thin, prepare for JDBC Thin Connections. See Prepare for JDBC Thin Connections.

SQLcl with Oracle Call Interface

To connect using Oracle Call Interface, use the `-oci` option, supply the database user name, a password, and the database service name provided in the `tnsnames.ora` file. For example:

```
sql -oci

SQLcl: Release 22.1 Production on Fri May 06 16:07:46 2022

Copyright (c) 1982, 2022, Oracle.  All rights reserved.

Username? (''?) adb_user@db2022adb_medium
Password? (**********?) ***************
```

```
Connected.
SQL>
```

When connecting using Oracle Call Interface, the Oracle Wallet is transparent to SQLcl.

SQLcl with a JDBC Thin Connection

To connect using a JDBC Thin connection, first configure the SQLcl cloud configuration and then connect to the database.

1. Start SQLcl with the /nolog option.

   ```
   sql /nolog
   ```

2. Configure the SQLcl session to use your Oracle Wallet:

   ```
   SQL> set cloudconfig directory/client_credentials.zip
   ```

3. Connect to the database:

   ```
   SQL> connect username@servicename
   password
   ```

   To avoid the prompt, connect and supply the password inline:

   ```
   SQL> connect username/password@servicename
   ```

   For example:

   ```
   sql /nolog

   SQLcl: Release 22.1 Production on Fri May 06 14:48:26 2022

   Copyright (c) 1982, 2022, Oracle.  All rights reserved.

   SQL> set cloudconfig /home/adb/Wallet_db2022ADB.zip

   SQL> connect adb_user@db2022adb_medium

   Password? (**********?) ***************
   Connected.
   SQL>
   ```

SQLcl with a JDBC Thin Connection with an HTTP Proxy

1. Start SQLcl with the /nolog option.

   ```
   sql /nolog
   ```

2. Configure the SQLcl session to use a proxy host and your Oracle Wallet:

   ```
   SQL> set cloudconfig -proxy=proxyhost:port directory/client_credentials.zip
   ```

3. Connect to the database.

   ```
   SQL> connect username@servicename
   password
   ```

   To avoid the prompt, connect and supply the password inline:

   ```
   SQL> connect username/password@servicename
   ```

For example:

```
sql /nolog

SQLcl: Release 22.1 Production on Fri May 06 11:59:38 2022

Copyright (c) 1982, 2022, Oracle.  All rights reserved.
SQL> set cloudconfig -proxy=http://myproxyhost.com:80 /home/adb/
Wallet_db2022.zip

SQL> connect adb_user@db2022adb_medium

Password? (**********?) ***************
Connected.
SQL>
```

> **Note:**
>
> If you are connecting to Autonomous Database using Microsoft Active Directory credentials, then connect using an Active Directory user name in the form of "*AD_domain\AD_username*" (double quotes must be included), and Active Directory user password. See Use Microsoft Active Directory with Autonomous Database for more information.

For more information, on the connection types specified in `tnsnames.ora`, see Manage Concurrency and Priorities on Autonomous Database.

For information on SQLcl, see Oracle SQLcl.

Connect Oracle SQLcl Cloud Without a Wallet

SQLcl is a command-line interface used to enter SQL commands. You can use SQLcl to connect to an Autonomous Database with TLS authentication without a wallet.

> **Note:**
>
> See Update your Autonomous Database Instance to Allow both TLS and mTLS Authentication for information on allowing TLS connections.

You can use SQLcl version 4.2 or later with Autonomous Database. Download SQLcl from oracle.com.

If you use JDBC Thin Driver, prepare for JDBC Thin connections. See Prepare for JDBC Thin Connections.

To connect using a JDBC Thin Driver with TLS authentication, do the following to connect to the database.

1. Copy a connection string for the Autonomous Database.

To connect with TLS authentication copy a TLS connection string. On the Database Connection page, under **TLS Authentication**, select **TLS** to view the connection strings for connecting with TLS authentication.

See View TNS Names and Connection Strings for an Autonomous Database Instance for information on viewing and copying connection stings.

See Database Service Names for Autonomous Database for information on the different databases services for each connection string.

2. Start SQLcl and connect to the database:

   On UNIX/Linux start `sql` with the connection string, enclosed in quotes on the command line, as follows:

   ```
   sql username/password@'my_connect_string'
   ```

   For example (for clarity line breaks added):

   ```
   $ sql admin/password@'(description= (retry_count=20)(retry_delay=3)
   (address=(protocol=tcps)(port=1521)(host=adb.region.oraclecloud.com))
   (connect_data=(service_name=u9adutfb2ba8x4d_database_medium.adb.oraclecloud
   .com))
   (security=(ssl_server_dn_match=yes)))'

   SQLcl: Release 21.2 Production on Thu Sep 16 10:43:00 2021
   Copyright (c) 1982, 2021, Oracle. All rights reserved.

   Last Successful login time: Thu Sep 16 2021 10:43:01 -07:00

   Connected to:
   Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
   Version 19.12.0.1.0

   SQL>
   ```

   On Windows, start `sql` with the `/nolog` option and then connect with the copied connection string, as follows (as compared to UNIX/Linux, on Windows do not surround the connection string with quotes):

   ```
   > sql /nolog

   SQLcl: Release 21.2 Production on Fri Sep 17 10:15:01 2021
   Copyright (c) 1982, 2021, Oracle.  All rights reserved.
   SQL> conn username/password@my_connect_string
   ```

   For example (for clarity line breaks are added):

   ```
   > sql admin/password@(description= (retry_count=20)(retry_delay=3)
   (address=(protocol=tcps)(port=1521)(host=adb.region.oraclecloud.com))
   (connect_data=(service_name=u9adutfb2ba8x4d_database_medium.adb.oraclecloud
   .com))
   (security=(ssl_server_dn_match=yes)))

   SQLcl: Release 21.2 Production on Thu Sep 16 10:43:00 2021
   Copyright (c) 1982, 2021, Oracle. All rights reserved.
   ```

```
Last Successful login time: Thu Sep 16 2021 10:43:01 -07:00

Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
Version 19.12.0.1.0

SQL>
```

> **Note:**
>
> If you are connecting to Autonomous Database using Microsoft Active Directory credentials, then connect using an Active Directory user name in the form of "*AD_domain\AD_username*" (double quotes must be included), and Active Directory user password. See Use Microsoft Active Directory with Autonomous Database for more information.

For information on SQLcl, see Oracle SQLcl.

Connect with Built-In Oracle Database Actions

You can access Database Actions from Autonomous Database. Database Actions provides development tools, data tools, administration, and monitoring features for Autonomous Database. Using Database Actions you can run SQL statements, queries, and scripts in a worksheet.

- About Database Actions (SQL Developer Web)
  Database Actions provides a web-based interface with development, data tools, administration, monitoring, and download features for Autonomous Database.
- Access Database Actions as ADMIN
  Database Actions (also known as SQL Developer Web) is bundled with each Autonomous Database instance.
- Provide Database Actions Access to Database Users
  The ADMIN user provides access to Database Actions for other database users.
- Required Roles to Access Tools from Database Actions
  Lists the database roles required to use the built-in Autonomous Database tools.
- IAM Policies to Access Database Actions and Tools from the OCI Console
  Lists the IAM policies required to access built-in Autonomous Database tools and Database Actions from the OCI Console.
- Access Database Actions, Oracle APEX, Oracle REST Data Services, and Developer Tools Using a Vanity URL
  By default you access Oracle APEX apps, REST endpoints, and developer tools on Autonomous Database using the `oraclecloudapps.com` domain name. You can optionally configure a vanity URL or custom domain name that is easy to remember to help promote your brand identity.

About Database Actions (SQL Developer Web)

Database Actions provides a web-based interface with development, data tools, administration, monitoring, and download features for Autonomous Database.

The following table lists the main features of Database Actions. Database Actions on your Autonomous Database instance may include additional cards if you download applications from Oracle Cloud Marketplace.

Feature Area	Database Actions Cards
Development	SQL, Data Modeler, REST, JSON, Charts, Scheduling, Oracle Machine Learning, Graph Studio, and Oracle APEX
Data Studio	Data Load, Catalog, Data Insights, Data Transforms, and Data Analysis
Administration	Database Users, APEX Workspaces, Data Pump, Download Client Credentials, and Set Resource Management Rules
Monitoring	Performance Hub and Database Dashboard
Downloads	Download Oracle Instant Client and Download SODA Drivers
Related Services	Restful Data Services (ORDS) and SODA and Access Oracle Machine Learning Restful Services

See About Database Actions in *Using Oracle Database Actions* for more information.

Access Database Actions as ADMIN

Database Actions (also known as SQL Developer Web) is bundled with each Autonomous Database instance.

Database Actions runs in Oracle REST Data Services and access is provided through schema-based authentication. To use Database Actions, you must sign in as a database user whose schema is enabled for Database Actions. By default the ADMIN user is enabled to access Database Actions.

See Provide Database Actions Access to Database Users to enable another database user's schema to access Database Actions.

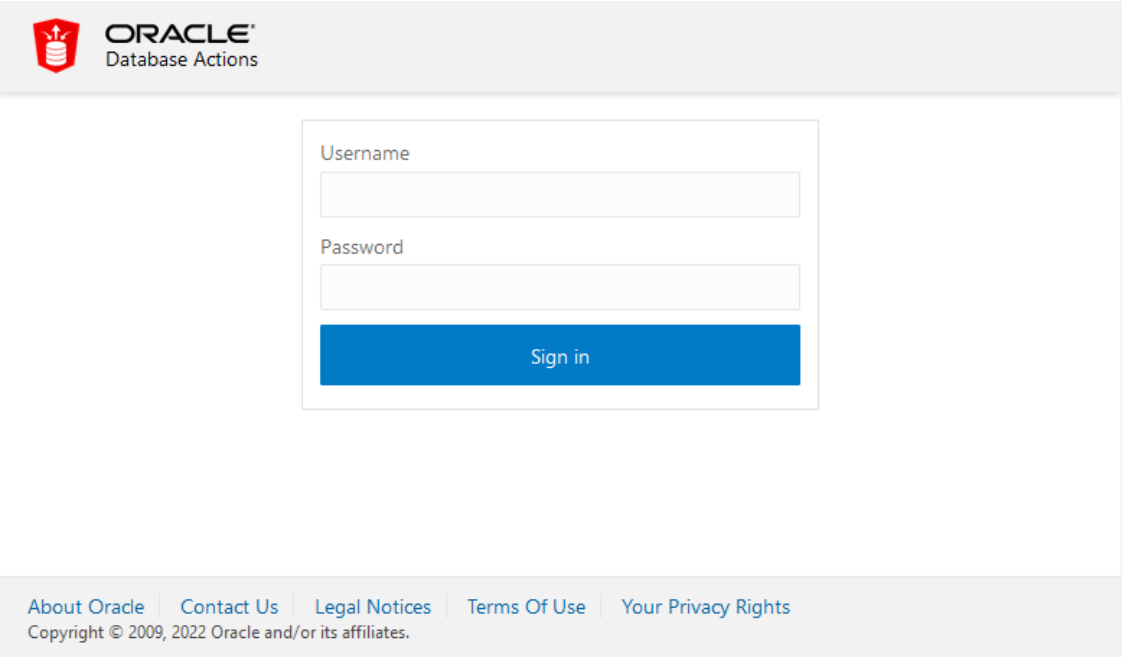
> **Note:**
>
> If your Autonomous Database is configured to use a Private Endpoint, then you can only access Database Actions from clients in the same Virtual Cloud Network (VCN). See Configure Network Access with Private Endpoints for more information.

To access Database Actions from the Oracle Cloud Infrastructure Console:

1. On the Autonomous Database Details page click the **Database actions** dropdown list.
2. Select a quick link to go directly to a quick link action or select **View all database actions** to access the full Database Actions Launchpad.

   For example, click **SQL** to use a SQL Worksheet. On the SQL Worksheet you can use the **Consumer Group** drop-down list to select the consumer group to run your SQL or PL/SQL code. See Executing SQL Statements in the Worksheet Editor for more information.

Depending on your browser, if the Console cannot access the database as ADMIN you will be prompted for your database ADMIN username and password.



Provide Database Actions Access to Database Users

The ADMIN user provides access to Database Actions for other database users.

Database users who are not service administrators do not have access to the Oracle Cloud Infrastructure Console. The ADMIN user provides access to Database Actions as follows:

- Use Database Actions to create a user and assign roles to the user. If the user already exists, check that Web Access is selected for the schema (with Web Access selected, the user's card shows **REST Enabled**).

  See Create Users on Autonomous Database for information on adding database users.

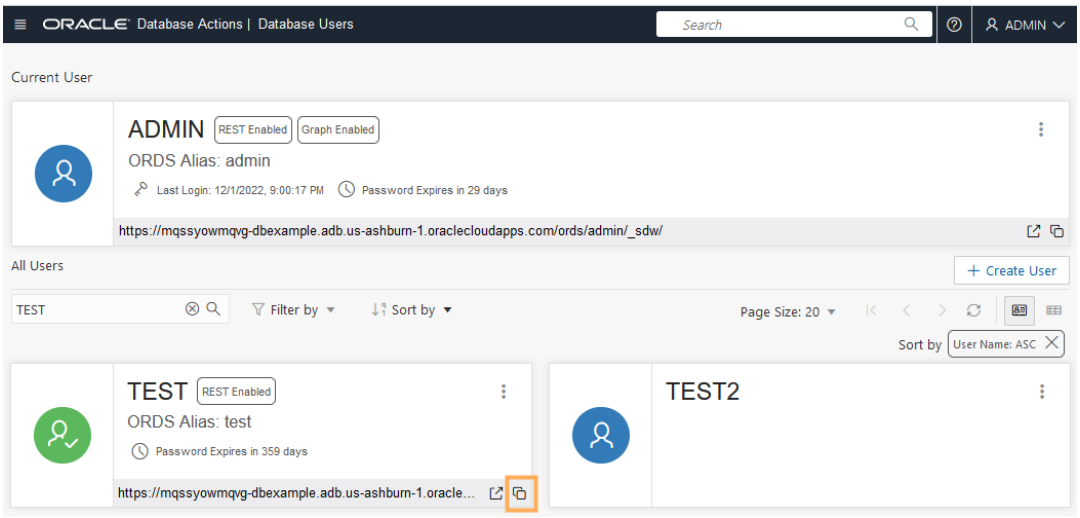
  See Required Roles to Access Tools from Database Actions for information on required roles for Database Actions.
- Provide the user with a URL to access Database Actions.

After adding a user and selecting Web Access, the ADMIN provides a user with the URL to access Database Actions:

1. In Database Actions, click ☰ to show the available actions.
2. Under **Administration** select **Database Users**.

   This displays information about users, such as user names, whether a user is REST Enabled, and the last login date and time. On a user's card, the icon on the left displays the user status with one of the following colors: green (Open), blue (Locked), or red (Expired).

   The default view is Card View. You can select either grid view or card view by clicking the **Card View** or **Grid View** icons.
3. A URL is displayed in the user's card only if the user is REST Enabled. It provides the URL to access Database Actions. Click ⧉ to copy the URL to the clipboard.



4. Provide the user with the URL you copied.

After you provide the URL to a user, to access Database Actions the user pastes the URL into their browser and then enters their Username and Password in the Sign-in dialog.

See Manage Users and User Roles on Autonomous Database - Connecting with Database Actions for more information.

Required Roles to Access Tools from Database Actions

Lists the database roles required to use the built-in Autonomous Database tools.

When you connect to Database Actions as the ADMIN user, many database roles are set to allow you to access the available Autonomous Database tools. When you connect to Database Actions as a non-ADMIN user, some of the Database Actions cards for Autonomous Database tools are hidden if you do not have required permissions.

As the ADMIN user, set the appropriate roles to allow user access to Autonomous Database tools. See Manage User Roles and Privileges on Autonomous Database for more information.

Autonomous Database Tool	Required Role	More Information
Oracle Machine Learning (OML)	`OML_DEVELOPER`	Oracle Machine Learning
Graph Studio	`GRAPH_DEVELOPER`	Using Oracle Graph with Autonomous Database
Data Studio (including: Catalog, Data Load, Data Analysis, Data Insights)	`DWROLE`	The Data Studio Overview Page
Data Transforms	`DATA_TRANSFORM_USER`	The Data Transforms Page

IAM Policies to Access Database Actions and Tools from the OCI Console

Lists the IAM policies required to access built-in Autonomous Database tools and Database Actions from the OCI Console.

In addition to having database roles to access the built-in tools and Database Actions, you must set OCI Identity and Access Management policies to either restrict access from the OCI Console or grant privileges to automatically login as ADMIN to tools or to Database Actions from the OCI Console.

IAM Policy	Description

`read autonomous-databases`
or
`inspect autonomous-databases`
	With these policies you must provide a *username* and *password* to access Database Actions from the OCI Console.
	For example, this policy allows members of the group `ADB-Admins` to review the details of an instance but not automatically login as ADMIN to Database Actions or Tools:
	`allow group ADB-Admins to inspect autonomous-databases in tenancy`

`use autonomous-databases`
or
`manage autonomous-databases`
	These policies provide automatic access to login to Database Actions from the OCI Console as ADMIN.
	For example, this policy allows members of the group `ADB-Admins2` to automatically login as ADMIN to Database Actions or Tools:
	`allow group ADB-Admins2 to manage autonomous-databases in tenancy`

See IAM Policies for Autonomous Database for more information.

Access Database Actions, Oracle APEX, Oracle REST Data Services, and Developer Tools Using a Vanity URL

By default you access Oracle APEX apps, REST endpoints, and developer tools on Autonomous Database using the `oraclecloudapps.com` domain name. You can optionally configure a vanity URL or custom domain name that is easy to remember to help promote your brand identity.

After you acquire a desired domain name and matching SSL certificate from a vendor of your choice, deploy an Oracle Cloud Infrastructure Load Balancer in your Virtual Cloud Network (VCN) using your Autonomous Database as the backend. Your Autonomous Database instance must be configured with a private endpoint in the same VCN. See Configuring Network Access with Private Endpoints for more information.

To learn more, see the following:

- Introducing Vanity URLs for APEX and ORDS on Oracle Autonomous Database
- Automate Vanity URL Configuration Using Terraform

Connect with JDBC Thin Driver

Autonomous Database mandates a secure connection that uses Transport Layer Security (TLSv1.2).

Java applications that use JDBC Thin driver connect with one of the following:

- **Mutual TLS (mTLS) Authentication**: requires either Oracle Wallet or Java KeyStore (JKS) where both the client and Autonomous Database authenticate each other.

  The wallet and keystore files are included in the client credentials `.zip` file that is available by clicking **Database connection** on the Oracle Cloud Infrastructure Console

- **TLS Authentication**: The client computer matches the server's CA root certificate against the client's list of trusted CAs. If the issuing CA is trusted, the client verifies that the certificate is authentic. This allows the client and Autonomous Database to establish the encrypted connection before exchanging any messages.

See Connecting to Oracle Database Manually for additional information on how to connect your Java applications to Autonomous Database using Oracle JDBC driver and Universal Connection Pool.

Topics

- JDBC Thin Connections with a Wallet (mTLS)
  Autonomous Database mandates a secure connection that uses Transport Layer Security (TLSv1.2). Depending on the network configuration options, Autonomous Database supports mTLS and TLS authentication.
- JDBC Thin Connections Without a Wallet (TLS)
  Autonomous Database mandates a secure connection that uses Transport Layer Security (TLSv1.2). Depending on the configuration options, Autonomous Database supports mTLS and TLS authentication. This section covers using JDBC Thin Connections with TLS authentication without a wallet.

JDBC Thin Connections with a Wallet (mTLS)

Autonomous Database mandates a secure connection that uses Transport Layer Security (TLSv1.2). Depending on the network configuration options, Autonomous Database supports mTLS and TLS authentication.

> **Note:**
>
> If you use TLS, instead of mTLS, for your connections using JDBC Thin Driver with JDK8u162 or higher, a wallet is not required.
>
> TLS connections are enabled for the following network configurations:
>
> - **Private endpoint access only**: network configuration with a private endpoint
> - **Secure access from allowed IPs and VCNs only**: configuration with an Access Control List (ACL)
>
>   If your Autonomous Database is on a public endpoint without any ACL, you can add `0.0.0.0/0` as your CIDR ACL and enable TLS authentication. Adding `0.0.0.0/0` as your CIDR ACL is identical to having your Autonomous Database on public endpoint with no ACL.
>
> See Secure Connections to Autonomous Database with mTLS or with TLS for more information.

- JDBC Thin Driver Connection Prerequisites Connections with Wallets (mTLS)
  Applications that use JDBC Thin driver support TLS and mutual TLS (mTLS) authentication. Using mTLS authentication requires that you supply Oracle database credentials including the Oracle wallets or Java KeyStore (JKS) files when connecting to the database.

- Using a JDBC URL Connection String with JDBC Thin Driver and Wallets
  The connection string is found in the file `tnsnames.ora` which is part of the client credentials download. The `tnsnames.ora` file contains the predefined service names. Each service has its own TNS alias and connection string.
- Using a JDBC Connection with 18.3 JDBC Driver
  Applications that use JDBC Thin driver can connect to Autonomous Databases using either Oracle Wallets or Java KeyStore (JKS).
- Connecting Using JDBC Thin Driver 12.2 or Older
  If you are using the JDBC driver 12.2.0.2 or older, set the Java properties prior to starting the application. Usually you set the properties in the application's startup script.
- JDBC Thin Connections with an HTTP Proxy
  If the client is behind a firewall and your network configuration requires an HTTP proxy to connect to the internet, you need to use the JDBC Thin Client 18.1 or higher which enables connections through HTTP proxies.

JDBC Thin Driver Connection Prerequisites Connections with Wallets (mTLS)

Applications that use JDBC Thin driver support TLS and mutual TLS (mTLS) authentication. Using mTLS authentication requires that you supply Oracle database credentials including the Oracle wallets or Java KeyStore (JKS) files when connecting to the database.

Perform the following steps before connecting to the database:

1. **Provision Autonomous Database**: Create a database and obtain your database credentials (username and password).

   See Provision an Autonomous Database Instance for more information.

2. For mutual TLS connections, **Download Client Credentials**: Unzip the `wallet_`*databasename*`.zip` to a secure location. Make sure that only authorized users have access to these files.

   See Download Client Credentials (Wallets) for information on downloading client credentials for Autonomous Database.

3. **Verify your JDK version for security**: If you are using JDK11, JDK10, or JDK9 then you don't need to do anything for this step. If your JDK version is less than JDK8u162 then you need to download the JCE Unlimited Strength Jurisdiction Policy Files. Refer to the `README` file for installation notes. Download the JCE files from Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files 8 Download.

4. **Check JDBC Driver Version**: Download a supported JDBC Thin driver (`ojdbc8.jar` and `ucp.jar`). You also need the additional jars: `oraclepki.jar`, `osdt_core.jar`, and `osdt_cert.jar` for use with Oracle wallets.

   Supported versions are:

   - **JDBC Thin**: 11.2.0.4 (or later with one-off patch for Bug 28492769), 12.2 (or later with one-off patch for Bug 28492769), 18 (base release or later with one-off patch for Bug 28492769), 19 (base release or later), or 21 (base release or later)

   For applications that use the Universal Connection Pool (UCP) feature of JDBC, it is highly recommended to use 19.13 or higher or 21.3 or higher versions of the JDBC driver. These versions include proper draining behavior to minimize impact to applications when planned maintenance is performed in Autonomous Database. UCP will replenish connections in the pool proactively so that active connections are not impacted by maintenance.

   For older versions of the driver, a patch for bug 31112088 can also be requested by filing a Service Request.

Download supported versions: Oracle Database JDBC driver and Companion Jars Downloads .

Using a JDBC URL Connection String with JDBC Thin Driver and Wallets

The connection string is found in the file `tnsnames.ora` which is part of the client credentials download. The `tnsnames.ora` file contains the predefined service names. Each service has its own TNS alias and connection string.

A sample entry, with `dbname_high` as the TNS alias and a connection string in `tnsnames.ora` follows:

```
dbname_high= (description=
      (address=(protocol=tcps)(port=1522)(host=adb.example.oraclecloud.com))
      (connect_data=(service_name=dbname_high.oraclecloud.com))
(security=(ssl_server_dn_match=yes)))
```

Set the location of `tnsnames.ora` with the property `TNS_ADMIN` in one of the following ways:

- As part of the connection string (only with the 18.3 or newer JDBC driver)
- As a system property, `-Doracle.net.tns_admin`
- As a connection property (`OracleConnection.CONNECTION_PROPERTY_TNS_ADMIN`)

Using the 18.3 JDBC driver, the connection string includes the TNS alias and the `TNS_ADMIN` connection property.

Sample connection string using 18.3 JDBC driver (Linux):

```
DB_URL="jdbc:oracle:thin:@dbname_high?TNS_ADMIN=/Users/test/wallet_dbname"
```

Sample connection string using 18.3 JDBC driver (Windows):

```
DB_URL="jdbc:oracle:thin:@dbname_high?TNS_ADMIN=C:\\Users\\test\
\wallet_dbname"
```

The `TNS_ADMIN` connection property specifies the following:

- The location of `tnsnames.ora`.
- The location of Oracle Wallet (`ewallet.sso`, `ewallet.p12`) or Java KeyStore (JKS) files (`truststore.jks`, `keystore.jks`).
- The location of `ojdbc.properties`. This file contains the connection properties required to use Oracle Wallets or Java KeyStore (JKS).

> **Note:**
>
> If you are using 12.2.0.1 or older JDBC drivers, then the connection string contains only the TNS alias. To connect using older JDBC drivers:
>
> - Set the location of the `tnsnames.ora`, either as a system property with `-Doracle.net.tns_admin` or as a connection property (`OracleConnection.CONNECTION_PROPERTY_TNS_ADMIN`).
> - Set the wallet or JKS related connection properties in addition to `TNS_ADMIN`.
>
> For example, in this case you set the TNS alias in the `DB_URL` without the `TNS_ADMIN` part as:
>
> ```
> DB_URL=”jdbc:oracle:thin:@dbname_high”
> ```

See Database Service Names for Autonomous Database for more details.

Using a JDBC Connection with 18.3 JDBC Driver

Applications that use JDBC Thin driver can connect to Autonomous Databases using either Oracle Wallets or Java KeyStore (JKS).

Using Oracle Wallet

To use Java and the 18.3 JDBC Thin Driver to connect to Autonomous Database with the Oracle Wallet, do the following:

1. **Make sure that the prerequisites are met:** See JDBC Thin Driver Connection Prerequisites Connections with Wallets (mTLS) for more information.
2. **Verify the connection**: You can either use a Java program, a servlet, or IDEs to verify the connection to the database. A simple test is to download DataSourceSample.java or UCPSample.java from JDBC code samples and update the connection URL to have the required TNS alias and pass `TNS_ADMIN`, providing the path for `tnsnames.ora` and the wallet files. Also, in the sample source code update the database username and password. For example:

   ```
   DB_URL="jdbc:oracle:thin:@dbname_high?TNS_ADMIN=/Users/test/wallet_dbname"
   ```

   > **Note:**
   >
   > If you are using Microsoft Active Directory with a database, then in the sample source code update the username with the Active Directory username and update the password with the Active Directory user password. See Use Microsoft Active Directory with Autonomous Database for more information.

3. **Set the wallet location**: The properties file `ojdbc.properties` is pre-loaded with the wallet related connection property.

   ```
   oracle.net.wallet_location=(SOURCE=(METHOD=FILE)(METHOD_DATA=(DIRECTORY=${TNS_ADMIN})))
   ```

> **Note:**
>
> You do not modify the file `ojdbc.properties`. The value of `TNS_ADMIN` determines the wallet location.

4. **Compile and Run:** Compile and run the sample to get a successful connection. Make sure you have `oraclepki.jar` , `osdt_core.jar`, and `osdt_cert.jar`, in the `classpath`. For example:

```
java -classpath
      ./lib/ojdbc8.jar:./lib/ucp.jar:./lib/oraclepki.jar:./lib/
osdt_core.jar:./lib/osdt_cert.jar:. UCPSample
```

> **Note:**
>
> The auto-login wallet part of Autonomous Database downloaded client credentials zip file removes the need for your application to use username/password authentication.

Using Java KeyStore

To use Java and the 18.3 JDBC Thin Driver to connect to Autonomous Database with Java KeyStore (JKS), do the following:

1. **Make sure that the prerequisites are met:** See JDBC Thin Driver Connection Prerequisites Connections with Wallets (mTLS) for more information.
2. **Ready the database details**: You can either use a Java program, a servlet, or IDEs to check the connection to your database. A simple test is to download DataSourceSample.java or UCPSample.java from JDBC code samples. In this sample, use the connection URL as shown. Note that the connection `DB_URL` contains the TNS alias, for example, `dbname_high` present in `tnsnames.ora`. You can provide the path for `tnsnames.ora` file through `TNS_ADMIN` property as shown in the URL. Make sure to use the database username and password related to your database.

```
DB_URL="jdbc:oracle:thin:@dbname_high?TNS_ADMIN=/Users/test/wallet_dbname"
```

> **Note:**
>
> If you are using Microsoft Active Directory with Autonomous Database, then make sure to change the sample source code to use the Active Directory username and the Active Directory user password. See Use Microsoft Active Directory with Autonomous Database for more information.

3. **Set JKS related connection properties**: Add the JKS related connection properties to `ojdbc.properties` file. The keyStore and truststore password are the password specified when you downloading the client credentials `.zip` file.

To use SSL connectivity instead of Oracle Wallet, specify the keystore and truststore files and their respective password in the `ojdbc.properties` file as follows:

```
# Properties for using Java KeyStore (JKS)
oracle.net.ssl_server_dn_match=true
javax.net.ssl.trustStore==${TNS_ADMIN}/truststore.jks
javax.net.ssl.trustStorePassword=password
javax.net.ssl.keyStore==${TNS_ADMIN}/keystore.jks
javax.net.ssl.keyStorePassword=password
```

> **Note:**
>
> Make sure to comment the wallet related property in `ojdbc.properties`. For example:
>
> ```
> # Property for using Oracle Wallets
> # oracle.net.wallet_location=(SOURCE=(METHOD=FILE)
> (METHOD_DATA=(DIRECTORY=${TNS_ADMIN})))
> ```

4. **Compile and Run:** Compile and run the sample to get a successful connection. For example:

```
java -classpath ./lib/ojdbc8.jar:./lib/ucp.jar UCPSample
```

Connecting Using JDBC Thin Driver 12.2 or Older

If you are using the JDBC driver 12.2.0.2 or older, set the Java properties prior to starting the application. Usually you set the properties in the application's startup script.

If you are not able to use the latest JDBC drivers, then you can connect to Autonomous Database using 12.2.0.2 or other older JDBC drivers. The 12.2 or older JDBC drivers do not support the `ojdbc.properties` file. With older JDBC driver versions, you need to pass wallets or JKS related properties either as system properties or as connection properties to establish a connection.

Using Oracle Wallet

To use Java and the 12.2 or older JDBC Drivers to connect to Autonomous Database with the Oracle Wallet, do the following:

1. **Make sure that the prerequisites are met:** See JDBC Thin Driver Connection Prerequisites Connections with Wallets (mTLS) for more information.
2. **Verify the connection**: You can either use a Java program, a servlet, or IDEs to verify the connection to the database. A simple test is to download DataSourceSample.java or UCPSample.java from JDBC code samples and update the connection URL to have the required TNS alias. Also, update the sample source code to use the database username and password. For example:

```
DB_URL="jdbc:oracle:thin:@dbname_high"
```

> **Note:**
>
> If you are using Microsoft Active Directory with Autonomous Database, then update the sample source code to use the Active Directory username and Active Directory user password. See Use Microsoft Active Directory with Autonomous Database for more information.

3. **Set the wallet location**: Add the `OraclePKIProvider` at the end of the provider list in the file `java.security` (this file is part of your JRE install located at `$JRE_HOME/jre/lib/security/java.security`) which typically looks like:

```
security.provider.14=oracle.security.pki.OraclePKIProvider
```

4. **Compile and Run:** Compile and run the sample to get a successful connection. Make sure to have `oraclepki.jar` , `osdt_core.jar`, and `osdt_cert.jar`, in the `classpath`. Also, you need to pass the connection properties. Update the properties with the location where `tnsnames.ora` and wallet files are located.

```
java -classpath
./lib/ojdbc8.jar:./lib/ucp.jar:./lib/oraclepki.jar:./lib/
osdt_core.jar:./lib/osdt_cert.jar:.
-Doracle.net.tns_admin=/users/test/wallet_dbname
-Doracle.net.ssl_server_dn_match=true
-Doracle.net.ssl_version=1.2  (Not required for 12.2)
-Doracle.net.wallet_location= "(SOURCE=(METHOD=FILE)
(METHOD_DATA=(DIRECTORY=/users/test/wallet_dbname)))"
UCPSample
```

> **Note:**
>
> These are Windows system examples. Add a `\` continuation character if you are setting `-D` properties on multiple lines on UNIX (Linux or a Mac).

Using Java KeyStore

To use Java and the 12.2 or older JDBC Thin Drivers to connect to Autonomous Database with Java KeyStore (JKS), do the following:

1. **Make sure that the prerequisites are met:** See JDBC Thin Driver Connection Prerequisites Connections with Wallets (mTLS) for more information.
2. **Verify the connection**: You can either use a Java program, a servlet, or IDEs to verify the connection to the database. A simple test is to download DataSourceSample.java or UCPSample.java from JDBC code samples and update the connection URL to have the required TNS alias and pass `TNS_ADMIN`, providing the path for `tnsnames.ora` and update the connection URL to have the required TNS alias. Also, in the sample source code update the database username and password. For example:

```
DB_URL="jdbc:oracle:thin:@dbname_high"
```

> **Note:**
>
> If you are using Microsoft Active Directory with Autonomous Database, then update the sample source code to use the Active Directory username and Active Directory user password. See Use Microsoft Active Directory with Autonomous Database for more information.

3. **Compile and Run**: Compile and run the sample to get a successful connection. You need to pass the connection properties as shown. Update the properties with the location where `tnsnames.ora` and JKS files are placed. If you want to pass these connection properties programmatically then refer to DataSourceForJKS.java. For example:

```
java
-Doracle.net.tns_admin=/users/test/wallet_dbname
-Djavax.net.ssl.trustStore=truststore.jks
-Djavax.net.ssl.trustStorePassword=**********
-Djavax.net.ssl.keyStore=keystore.jks
-Djavax.net.ssl.keyStorePassword=************
-Doracle.net.ssl_server_dn_match=true
-Doracle.net.ssl_version=1.2 // Not required for 12.2
```

JDBC Thin Connections with an HTTP Proxy

If the client is behind a firewall and your network configuration requires an HTTP proxy to connect to the internet, you need to use the JDBC Thin Client 18.1 or higher which enables connections through HTTP proxies.

To connect to Autonomous Database through an HTTPS proxy, open and update your `tnsnames.ora` file. Add the HTTP proxy hostname(https_proxy) and port (https_proxy_port) to the connection string. Replace the values with your HTTPS proxy information. For example:

1. Add the HTTP proxy hostname and port to the connection definitions in `tnsnames.ora`. You need to add the https_proxy and https_proxy_port parameters in the address section of connection definitions. For example, the following sets the HTTP proxy to `proxyhostname` and the HTTP proxy port to `80`; replace these values with your HTTP proxy information:

```
db2022adb_high =
       (description=
             (address=
                   (https_proxy=proxyhostname)(https_proxy_port=80)
(protocol=tcps)(port=1522)(host=adb.example.oraclecloud.com)
             )
             (connect_data=(service_name=db2022adb_high.adb.oraclecloud.com)
             )
             (security=security=(ssl_server_dn_match=yes)
             )
       )
```

Notes:

- JDBC Thin client versions earlier than 18.1 do not support connections through HTTP proxy.
- Successful connection depends on specific proxy configurations and the performance of data transfers would depend on proxy capacity. Oracle does not recommend using this feature in Production environments where performance is critical.
- Configuring `tnsnames.ora` for the HTTP proxy may not be enough depending on your organization's network configuration and security policies. For example, some networks require a username and password for the HTTP proxy.
- In all cases, contact your network administrator to open outbound connections to hosts in the `oraclecloud.com` domain using the relevant port without going through an HTTP proxy.

JDBC Thin Connections Without a Wallet (TLS)

Autonomous Database mandates a secure connection that uses Transport Layer Security (TLSv1.2). Depending on the configuration options, Autonomous Database supports mTLS and TLS authentication. This section covers using JDBC Thin Connections with TLS authentication without a wallet.

- JDBC Thin Driver Connection Prerequisites for TLS Connections Without a Wallet
  Applications that use JDBC Thin driver support TLS and mutual TLS (mTLS) authentication. Connecting to an Autonomous Database instance with TLS authentication requires database credentials (username and password) and provides a secure connection, but does not require that you download Oracle wallets or Java KeyStore (JKS) files.
- Using a JDBC URL TLS Connection String for JDBC Thin Driver Without a Wallet
  To connect the database using JDBC Thin Driver with TLS without a wallet, you need to provide a connection string. Each database service has its own TNS Name and connection string.

JDBC Thin Driver Connection Prerequisites for TLS Connections Without a Wallet

Applications that use JDBC Thin driver support TLS and mutual TLS (mTLS) authentication. Connecting to an Autonomous Database instance with TLS authentication requires database credentials (username and password) and provides a secure connection, but does not require that you download Oracle wallets or Java KeyStore (JKS) files.

Note:

See Update your Autonomous Database Instance to Allow both TLS and mTLS Authentication for information on allowing TLS connections.

Perform the following steps before connecting to the database:

1. **Provision Autonomous Database**: Create a database and obtain your database credentials (username and password).

See Provision an Autonomous Database Instance for more information.

2. **Verify your JDK version for security**: If you are using JDK11, JDK10, or JDK9 then you don't need to do anything for this step. If your JDK version is less than JDK8u162 then you need to download the JCE Unlimited Strength Jurisdiction Policy Files. Refer to the `README` file for installation notes. Download the JCE files from Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files 8 Download.
3. **Check JDBC Driver Version**: The JDBC Thin client versions from 12.2.0.1 or higher are supported. However, it is recommended to use the latest release update from 19c or 23ai to take advantage of recent enhancements that simplify connections and provide easy steps for configuration. Obtain one of these from Oracle Database JDBC driver and Companion Jars Downloads.

Using a JDBC URL TLS Connection String for JDBC Thin Driver Without a Wallet

To connect the database using JDBC Thin Driver with TLS without a wallet, you need to provide a connection string. Each database service has its own TNS Name and connection string.

To run an application using the JDBC Thin Driver with TLS authentication without a wallet:

1. Copy a connection string for the Autonomous Database.

   To connect with TLS authentication copy a TLS connection string.

   a. From the Oracle Cloud Infrastructure Console, on the Autonomous Database details page click **Database connection**.
   b. Under **TLS Authentication**, select **TLS** to view the connection strings for connecting with TLS authentication.
   c. Copy a connection string.

   See View TNS Names and Connection Strings for an Autonomous Database Instance for information on viewing and copying connection stings.

   See Database Service Names for Autonomous Database for information on the different databases services for each connection string.
2. Set the `DB_URL` parameter.

   Use the following format for the `DB_URL` parameter:

   ```
   DB_URL=jdbc:oracle:thin:@my_connect_string
   ```

   For example:

   ```
   DB_URL=jdbc:oracle:thin:@(description= (retry_count=20)(retry_delay=3)
   (address=(protocol=tcps)
   (port=1521)(host=adb.region.oraclecloud.com))
   (connect_data=(service_name=u9adutfb2ba8x4d_database_medium.adb.oraclecloud
   .com))
   (security=(ssl_server_dn_match=yes)))
   ```

Preparing for Oracle Call Interface Connections

Using TLS with Oracle Call Interface connections with newer Oracle Instant Client/Oracle Database Client versions, a wallet is not required. For TLS connections with Oracle Call

Interface running on older Oracle Instant Client/Oracle Database Client versions, a wallet is required for both mTLS authenticated connections and TLS authenticated connections.

- Oracle Call Interface (OCI) Connections and Wallets (mTLS)
  Oracle Net Services can find the location of the Autonomous Database wallet using the `WALLET_LOCATION` parameter in the `sqlnet.ora` file.
- Oracle Call Interface (OCI) Connections with TLS Authentication
  Using TLS with Oracle Call Interface, connections with newer Oracle Instant Client/Oracle Database Client versions, a wallet is not required.

Oracle Call Interface (OCI) Connections and Wallets (mTLS)

Oracle Net Services can find the location of the Autonomous Database wallet using the `WALLET_LOCATION` parameter in the `sqlnet.ora` file.

When `WALLET_LOCATION` is used, Oracle Net Services automatically uses the wallet. The wallet is used transparently to the application. See Prepare for Oracle Call Interface, ODBC, and JDBC OCI Connections with Wallets (mTLS) for information on setting `WALLET_LOCATION`.

See Download Client Credentials (Wallets) for information on downloading client credentials for Autonomous Database.

Oracle Call Interface (OCI) Connections with TLS Authentication

Using TLS with Oracle Call Interface, connections with newer Oracle Instant Client/Oracle Database Client versions, a wallet is not required.

Oracle Call Interface (OCI) clients support TLS authentication without a wallet if you are using the following client versions:

- Oracle Instant Client/Oracle Database Client 19.13 - only on Linux x64
- Oracle Instant Client/Oracle Database Client 19.14 (or later) and 21.5 (or later) - only on Linux x64 and Windows

> **Note:**
>
> See Update your Autonomous Database Instance to Allow both TLS and mTLS Authentication for information on allowing TLS connections.

See Prepare for Oracle Call Interface, ODBC, and JDBC OCI Connections Using TLS Authentication to prepare for Oracle Call Interface connections.

When a wallet is required and you set `WALLET_LOCATION` parameter in the `sqlnet.ora` file, Oracle Net Services finds the location of the wallet and uses the wallet (use of the wallet is transparent to the application).

Predefined Database Service Names for Autonomous Database

The `tnsnames.ora` file provided with the credentials zip file contains the database service names that allow you to connect to your database.

The available predefined service names differ depending on your workload. The different services provide different levels of performance and concurrency for Autonomous Database connections.

Workload Type	Database Service Names
Data Warehouse	*dbname*_high *dbname*_medium *dbname*_low See Database Service Names for Autonomous Database for more information.
Transaction Processing	*dbname*_tpurgent *dbname*_tp *dbname*_high *dbname*_medium *dbname*_low See Database Service Names for Autonomous Database for more information.
JSON Database	*dbname*_tpurgent *dbname*_tp *dbname*_high *dbname*_medium *dbname*_low See Database Service Names for Autonomous Database for more information.

Connect with Oracle Analytics Desktop

Oracle Analytics Desktop makes it easy to visualize your data so you can focus on exploring interesting data patterns. Just connect to Autonomous Database, select the elements that you're interested in, and let Oracle Analytics Desktop find the best way to visualize it. Choose from a variety of visualizations to look at data in a specific way.

For details on connecting Autonomous Database with Oracle Analytics Desktop, see *User's Guide for Oracle Analytics Desktop*.

Connect with Oracle Analytics Cloud

Oracle Analytics Cloud is a scalable and secure public cloud service that provides a full set of capabilities to explore and perform collaborative analytics for you, your workgroup, and your enterprise.

For details for connecting Autonomous Database with Oracle Analytics Cloud, see *Visualizing Data and Building Reports in Oracle Analytics Cloud*.

Connection and Networking Options and Features

Autonomous Database provides a number of different connection and networking options and features for connecting to a database.

- Using ACLs, VCNs, and Private Endpoints with Autonomous Database
  Describes the options for restricting network access to your Autonomous Database by specifying access control rules or using a virtual cloud network (private access with a private endpoint).

- Connect with Oracle Cloud Infrastructure FastConnect
  Oracle Cloud Infrastructure FastConnect provides an easy way for you to connect your on-premises network to Autonomous Database using FastConnect Public Peering. FastConnect provides higher-bandwidth options, and a more reliable and consistent networking experience compared to internet-based connections.
- Access Autonomous Database with VCN Transit Routing
  Oracle Cloud Infrastructure supports private access from your on-premises network to a database with virtual cloud network (VCN) Transit Routing.
- Access Autonomous Database with Service Gateway
  Autonomous Database supports private access from Oracle Cloud Infrastructure resources in a VCN through a service gateway.
- Use Database Resident Connection Pooling with Autonomous Database
  Database Resident Connection Pool (DRCP) in Autonomous Database supports easier and more efficient management of open connections. Using DRCP provides you with access to a connection pool in your database that enables a significant reduction in key database resources required to support many client connections and when the database needs to scale for many simultaneous connections.
- Use TCP Fast Open (TFO) to Connect Clients to Autonomous Database
  If your application is sensitive to network latency and you want to decrease the network latency between your application and the database, you can enable TCP Fast Open (TFO).
- Best Practices for Low Latency Connections with Autonomous Database
  Taking steps to reduce the latency for connections between an application and Autonomous Database is critical if your application performs many round-trips between the application and the database.
- View Network Information on the OCI Console
  On the Oracle Cloud Infrastructure Console you can view the network information for your Autonomous Database.

Using ACLs, VCNs, and Private Endpoints with Autonomous Database

Describes the options for restricting network access to your Autonomous Database by specifying access control rules or using a virtual cloud network (private access with a private endpoint).

- About Network Access Options
  Provides an overview of the network access options available when you provision or clone Autonomous Database.
- Overview of Restricting Access with ACLs
  When you select the network access **Secure access from allowed IPs and VCNs only** option when you provision or clone an instance, you can restrict network access by defining an Access Control List (ACL). You can also add, update, or remove an ACL for an active instance.
- Overview of Private Endpoints
  You can specify that Autonomous Database uses a private endpoint inside your Virtual Cloud Network (VCN) in your tenancy. You can configure a private endpoint during provisioning or cloning your Autonomous Database, or you can switch to using private endpoints in existing databases that use public endpoints. This allows you to keep all traffic to and from your Autonomous Database off of the public internet.

About Network Access Options

Provides an overview of the network access options available when you provision or clone Autonomous Database.

> **Note:**
>
> For all of these options, the connections to your Autonomous Database use certificate-based authentication and Secure Sockets Layer (SSL).

When you provision or clone your Autonomous Database you specify one of the following network access options:

- **Secure access from everywhere**: This option assigns a public endpoint, public IP and hostname, to your database. With this selection you have two options:
  - The database is accessible from all IP addresses: this is the default option when you provision or clone Autonomous Database.

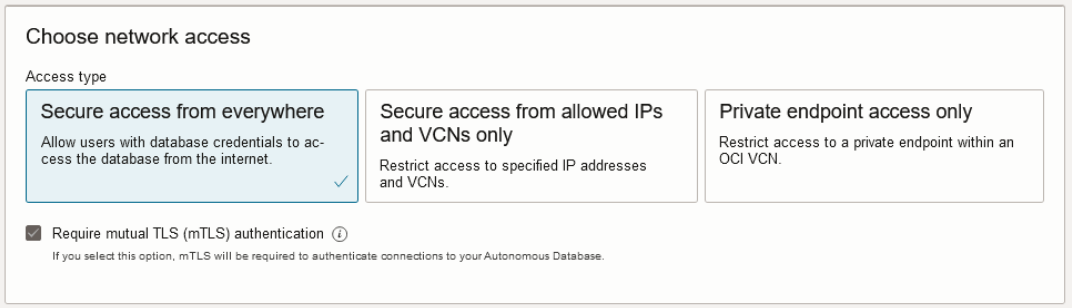
  If you configure your database with the **Secure access from everywhere** option, you can add, modify, or remove ACLs after you provision or clone the database.

  See Overview of Restricting Access with ACLs for more information.
- **Secure access from allowed IPs and VCNs only**: This option assigns a public endpoint, public IP and hostname, to your database. With this selection you must configure access control rules to restrict access by defining access control rules in an Access Control List (ACL). By specifying an ACL, the database will be accessible from allowed set of IP addresses or VCNs.

  If you configure your database with the **Secure access from everywhere** option, you can add, modify, or remove ACLs after you provision or clone the database.

  See Overview of Restricting Access with ACLs for more information.
- **Private endpoint access only**: This option assigns a private endpoint, private IP, and hostname to your database. Specifying this option allows traffic only from the VCN you specify; access to the database from all public IPs or VCNs is blocked. This allows you to define security rules, ingress/egress, at the Network Security Group (NSG) level and to control traffic to your Autonomous Database.

  See Configure Private Endpoints When You Provision or Clone an Instance for more information.



> **Note:**
>
> After you provision or clone your database, you can change the selection you make.

See Configure Network Access with Access Control Rules (ACLs) and Private Endpoints for more information.

Overview of Restricting Access with ACLs

When you select the network access **Secure access from allowed IPs and VCNs only** option when you provision or clone an instance, you can restrict network access by defining an Access Control List (ACL). You can also add, update, or remove an ACL for an active instance.

Specifying an access control list blocks all IP addresses that are not in the ACL list from accessing the database. After you specify an access control list, the database only accepts connections from addresses on the access control list and the database rejects all other client connections.

Depending on where the client machines that connect to your database are located you have the following options with ACLs:

- If your client machines connect to your database through the public internet, then you can use ACLs to specify the client machine's public IP addresses or their public CIDR blocks. In this case only the specified public IP addresses can access your database.
- If the client machines reside in an Oracle Cloud Infrastructure Virtual Cloud Network (VCN), you can configure a Service Gateway to connect to your database. In this case, you can specify the VCN in your ACL, this allows all client machines in that VCN to access your database and blocks all other connections. Furthermore, you can specify the VCN and a list of private IP addresses or CIDR blocks in that VCN. This allows only those client machines with the specified IP addresses or CIDR blocks to access your database and blocks all other connections.

  See VCNs and Subnets for details on Virtual Cloud Networks (VCN).

  See Access to Oracle Services: Service Gateway for details on setting up a Service Gateway.
- If you have on-premises clients that connect to your database through Transit Routing, you can specify the VCN and also the private IP addresses or CIDR blocks of these on-premises clients to access to your database.

  See Transit Routing: Private Access to Oracle Services for details on Transit Routing.
- You can use these options together to set multiple rules to allow access from different types of clients. Multiple rules do not exclude each other.

See Configuring Network Access with Access Control Rules (ACLs) for the steps for configuring network access with ACLs, either when you provision or clone your database, or whenever you want to add, modify or remove ACLs.

Overview of Private Endpoints

You can specify that Autonomous Database uses a private endpoint inside your Virtual Cloud Network (VCN) in your tenancy. You can configure a private endpoint during provisioning or cloning your Autonomous Database, or you can switch to using private endpoints in existing databases that use public endpoints. This allows you to keep all traffic to and from your Autonomous Database off of the public internet.

Specifying the **Private endpoint access only** configuration option only allows traffic from the VCN you specify and blocks access to the database from all public IPs or VCNs. This allows you to define security rules, ingress/egress, at the Network Security Group (NSG) level and to control traffic to your database.

Select the **Allow public access** option when you want to configure an Autonomous Database to use a private endpoint and you also want to allow connections from specific public IP addresses or from specific VCNs (if the VCNs are configured to privately connect to Autonomous Database using a Service Gateway).

See Configure Network Access with Private Endpoints for the steps for configuring network access with a private endpoint, either when you provision or clone your database or when you modify a private endpoint.

Connect with Oracle Cloud Infrastructure FastConnect

Oracle Cloud Infrastructure FastConnect provides an easy way for you to connect your on-premises network to Autonomous Database using FastConnect Public Peering. FastConnect provides higher-bandwidth options, and a more reliable and consistent networking experience compared to internet-based connections.

Use FastConnect to access public services in Oracle Cloud Infrastructure without using the internet, for example, access to Object Storage, or the Oracle Cloud Infrastructure Console and APIs. Without FastConnect, the traffic destined for public IP addresses would be routed over the internet. With FastConnect, that traffic goes over your private physical connection.

For details for connecting Autonomous Database with Oracle Cloud Infrastructure FastConnect see FastConnect Overview.

Access Autonomous Database with VCN Transit Routing

Oracle Cloud Infrastructure supports private access from your on-premises network to a database with virtual cloud network (VCN) Transit Routing.

Transit routing refers to a network setup in which your on-premises network uses a connected virtual cloud network (VCN), through a service gateway on the VCN for connectivity to a database. You connect the on-premises network to the VCN with FastConnect or VPN Connect, and then configure the VCN routing so that traffic transits through the VCN to a database beyond the VCN.

After transit routing is configured, there are no special steps required to access Oracle APEX or Database Actions, or to consume RESTful services published in Oracle REST Data Services.

See Transit Routing: Private Access to Oracle Services for more details and options available with Transit Routing.

See VPN Connect for more information on VPN Connect.

Access Autonomous Database with Service Gateway

Autonomous Database supports private access from Oracle Cloud Infrastructure resources in a VCN through a service gateway.

A service gateway allows connectivity to Autonomous Database from private IP addresses in private subnets without requiring a NAT Gateway in your VCN. After a service gateway is configured, there are no special steps required to connect to Autonomous Database. Use the

same connection steps as described in this chapter, depending on your application type or the client tool you are using to connect to the database.

After a service gateway is configured, there are no special steps required to access Oracle APEX or Database Actions, or to consume RESTful services published in Oracle REST Data Services.

See Access to Oracle Services: Service Gateway for details.

Use Database Resident Connection Pooling with Autonomous Database

Database Resident Connection Pool (DRCP) in Autonomous Database supports easier and more efficient management of open connections. Using DRCP provides you with access to a connection pool in your database that enables a significant reduction in key database resources required to support many client connections and when the database needs to scale for many simultaneous connections.

When you connect to Autonomous Database you choose one of the following depending on values specified in the `tnsnames.ora` configuration file:

- A dedicated server process, which services only one user process.
- A pooled server process, obtained from DRCP, which can service multiple user processes.

To connect with a pooled DRCP server process, do the following:

1. Locate or obtain the `tnsnames.ora` file you are using to connect to your Autonomous Database.

   See Download Client Credentials (Wallets) for more information.

2. Modify the `tnsnames.ora` file to add the server type `SERVER=POOLED`.

   For example:

   ```
   example_high= (description=
        (address=(protocol=tcps)(port=1522)(host=adb.example.oraclecloud.com))
        (connect_data=(service_name=example_high.oraclecloud.com)
   (SERVER=POOLED))
        (security=(ssl_server_dn_match=yes)))
   ```

   When you connect with `(SERVER=POOLED)` specified in the `tnsnames.ora` file you obtain a connection from DRCP.

For Autonomous Database, note the following for working with Database Resident Connection Pools (DRCP):

- DRCP is enabled by default; however using DRCP is optional. To choose a pooled connection specify `SERVER=POOLED` in `tnsnames.ora`. If you do not specify `SERVER=POOLED`, you connect with a dedicated connection.
- You cannot start or stop DRCP.

See Using Database Resident Connection Pool for more information.

Use TCP Fast Open (TFO) to Connect Clients to Autonomous Database

If your application is sensitive to network latency and you want to decrease the network latency between your application and the database, you can enable TCP Fast Open (TFO).

Depending on the characteristics of your application, TFO connections can improve application performance by decreasing network latency for client connections.

- Prerequisites for TCP Fast Open (TFO) Connections
  Lists the prerequisites to use TCP Fast Open (TFO) connections with Autonomous Database.
- JDBC Thin Connections with TCP Fast Open
  Describes the requirements to use JDBC Thin connections with TCP Fast Open (TFO) on Autonomous Database.
- TCP Fast Open Connections with Python
  Describes the requirements to use Python with TCP Fast Open (TFO) connections on Autonomous Database.
- TCP Fast Open Connections with Oracle Call Interface Client for C or SQL*Plus Connections
  Describes the requirements to use Oracle Call Interface Client with TCP Fast Open (TFO) to connect with C or SQL*Plus to Autonomous Database.

Prerequisites for TCP Fast Open (TFO) Connections

Lists the prerequisites to use TCP Fast Open (TFO) connections with Autonomous Database.

To use TCP Fast Open (TFO) to connect to an Autonomous Database instance, the client application requires the following:

- The client must run on a VM in Oracle Cloud Infrastructure.
- One of the following must be true:
  - The Autonomous Database is on a private endpoint.
  - When the Autonomous Database is on a public endpoint, the routing rules between the client VM and the database must be configured so that the traffic goes through a Service Gateway.

    See Access Autonomous Database with Service Gateway and Configure Network Access with Private Endpoints for more information.
- The VM must be running on Linux. Supported platforms are: Oracle Linux x86_64 and Oracle Linux Arm (aarch64).
- The following OS level kernel parameters must be set:
  - As root user, set the following parameters:
    * `sysctl -w net.ipv4.tcp_fastopen=5`
    * `sysctl -w net.ipv4.tcp_fastopen_blackhole_timeout_sec=0`

    Verify the OS level kernel parameter values:

```
cat /proc/sys/net/ipv4/tcp_fastopen
5
cat /proc/sys/net/ipv4/tcp_fastopen_blackhole_timeout_sec
0
```

JDBC Thin Connections with TCP Fast Open

Describes the requirements to use JDBC Thin connections with TCP Fast Open (TFO) on Autonomous Database.

Assure that the client environment and client prerequisites are met. See Prerequisites for TCP Fast Open (TFO) Connections for more information.

To use JDBC with TCP Fast Open (TFO), perform the following steps:

1. Download the latest JDBC-thin driver.

   From the page Oracle Database JDBC driver and Companion Jars Downloads:

   - Select a driver version 23.4 or higher from the available drivers under **Oracle Database 23ai JDBC Driver & UCP Downloads - Free Release**.
   - Select a driver version 19.23 or higher from the available drivers under **Oracle Database 19c JDBC Driver & UCP Downloads - Long Term Release**

2. Download the Oracle JDBC TCP Fast Open (TFO) Library File.

   a. On the page Oracle Database JDBC driver and Companion Jars Downloads, expand **Oracle JDBC TCP Fast Open (TFO) Library v1**

   b. Download the native library .zip for your client platform:

      - JDBC TFO Library for Linux x86_64, library file for “x86_64” for Linux on Intel.
      - JDBC TFO Library for Linux AARCH64 library file for “arm64” for Linux on Arm.

3. Load the native library in the environment where you run your application.

   Set the following environment variables in the library path:

   a. Export `LD_PRELOAD`.

   ```
   export LD_PRELOAD=path_to/libtfojdbc1.so;
   ```

   b. If you are using JDK 8, export `LD_LIBRARY_PATH`.

   For example:

   ```
   export LD_LIBRARY_PATH=path_of_libtfojdbc1.so:path_to/jdk_8_home/jre/lib/
   amd64:path_to/jdk_8_home/jre/lib/amd64/server;
   ```

   c. If you are using JDK 17, export `LD_LIBRARY_PATH`.

   For example:

   ```
   export LD_LIBRARY_PATH=extracted_binaries_path:path_to/jdk_17_home/lib;
   ```

4. Enable the TCP Fast Open (TFO) option for your application.

   The following shows some of the options for enabling TFO in the client connection to the database:

   - In the client application, set the following connection property:

     ```
     properties.put(OracleConnection.CONNECTION_PROPERTY_NET_USE_TCP_FAST_OPEN,
     "true");
     ```

     By default the property `OracleConnection.CONNECTION_PROPERTY_NET_USE_TCP_FAST_OPEN` is set to false (disabled).

   - Add the TFO property in the JDBC URL.

For example:

```
jdbc:oracle:thin:@tcps://xxx.adb.yyy.oraclecloud.com:1521/
zzz.adb.oraclecloud.com?oracle.net.useTcpFastOpen=true
```

Or

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)
(HOST=xxx.adb.yyy.oraclecloud.com) (PORT=1521))
(CONNECT_DATA=(SERVICE_NAME=zzz.adb.oraclecloud.com)))?
oracle.net.useTcpFastOpen=true
```

- Add the following TNS parameter in the `CONNECT_DATA` section of the connection URL.

  For example:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)
(HOST=xxx.adb.yyy.oraclecloud.com)(PORT=1521))
        (CONNECT_DATA=(USE_TCP_FAST_OPEN =ON)
(SERVICE_NAME=zzz.adb.oraclecloud.com)))
```

- Add the following TNS parameter in the `CONNECT_DATA` section of the connection URL.

  For example:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)
(HOST=xxx.adb.yyy.oraclecloud.com)(PORT=1521))(CONNECT_DATA=(TFO=ON)
(SERVICE_NAME=zzz.adb.oraclecloud.com)))
```

5. Run the application.

TCP Fast Open Connections with Python

Describes the requirements to use Python with TCP Fast Open (TFO) connections on Autonomous Database.

Assure that the client environment and client prerequisites are met. See Prerequisites for TCP Fast Open (TFO) Connections for more information.

To use Python with TCP Fast Open (TFO), perform the following steps:

1. Install Python and the `python-oracledb` Driver.

   Use the latest version of `python-orabledb` (Version 2.1.0 or newer is required).

   See Install Python and the python-oracledb Driver for more information.

2. Add the `use_tcp_fast_open` connection parameter when you connect to an Autonomous Database instance.

   For example, use this option with one of:

   - oracledb.connect()
   - oracledb.connect_async()
   - oracledb.create_pool()
   - oracledb.create_pool_async()

   See Connect Python Applications to Autonomous Database for more information.

TCP Fast Open Connections with Oracle Call Interface Client for C or SQL*Plus Connections

Describes the requirements to use Oracle Call Interface Client with TCP Fast Open (TFO) to connect with C or SQL*Plus to Autonomous Database.

Assure that the client environment and client prerequisites are met. See Prerequisites for TCP Fast Open (TFO) Connections for more information.

> **Note:**
>
> You can enable TCP Fast Open (TFO) with either mTLS or TLS connections. See Secure Connections to Autonomous Database with mTLS or with TLS and Update Network Options to Allow TLS or Require Only Mutual TLS (mTLS) Authentication on Autonomous Database for more information.

To use Oracle Call Interface with TCP Fast Open (TFO), perform the following steps:

1. Download the latest Oracle instant client for Linux, either version 23.4 or higher or 19.23 or higher.

   See Oracle Instant Client Downloads for Linux x86-64 (64-bit) for Linux x86-64.

   See Oracle Instant Client Downloads for Linux ARM (aarch64) for Linux ARM.

2. Connect with the `USE_TCP_FAST_OPEN` parameter set to the value `YES`.

   For example, add the parameter `USE_TCP_FAST_OPEN` to `connect_data` in the connection string you use to connect to Autonomous Database:

   ```
   (description= (retry_count=20)(retry_delay=3)(address=(protocol=tcps)
      (port=1522)(host=example.us-phoenix-1.oraclecloud.com))
      (connect_data= (USE_TCP_FAST_OPEN=yes)
   (service_name=example_medium.adb.oraclecloud.com))
          (security=(ssl_server_dn_match=yes)))
   ```

   Or set the `sqlnet.ora` parameter:

   ```
   USE_TCP_FAST_OPEN=yes
   ```

Best Practices for Low Latency Connections with Autonomous Database

Taking steps to reduce the latency for connections between an application and Autonomous Database is critical if your application performs many round-trips between the application and the database.

For example, consider an OLTP application connecting to Autonomous Database and submitting thousands of SQL statements to the database individually to execute a sales order. In this case, the application requires thousands of round-trips, and reducing the latency for each round-trip can considerably speed up the sales order process. For such applications there are best practices that you can follow to reduce the latency for database connections.

- Steps to Reduce Latency for Database Connections
  You can follow these recommendations to reduce the latency for connections between your applications and the database.

- Steps to Reduce Latency for Database Connections for Databases with Autonomous Data Guard
  Provides steps to take to configure an Autonomous Data Guard standby environment, clients and mid-tiers, to reduce latency for database connections when you connect after a failover or after a switchover (when the standby becomes the Primary).
- Conceptual Network Diagram for Low Latency Database Connections
  Shows the conceptual network diagram for low latency connections using public endpoints and private endpoints for your database.

Steps to Reduce Latency for Database Connections

You can follow these recommendations to reduce the latency for connections between your applications and the database.

First determine the database's availability domain. To find an Autonomous Database instance's availability domain, connect as ADMIN and run the following query:

```
SELECT json_value(cloud_identity, '$.AVAILABILITY_DOMAIN')
AVAILABILITY_DOMAIN FROM v$pdbs;
```

For example:

```
SELECT json_value(cloud_identity, '$.AVAILABILITY_DOMAIN') AVAILABILITY_DOMAIN
             FROM v$pdbs;

AVAILABILITY_DOMAIN
-------------------
SoSC:US-ASHBURN-AD-1
```

You can also view the availability domain information on the Oracle Cloud Infrastructure Console. See View Network Information on the OCI Console for more information.

To reduce latency, do the following:

1. Place clients or the mid-tier servers in the same availability domain as the Autonomous Database instance.

   For example, if your application runs on an Oracle Cloud Infrastructure Compute VM, select the same availability domain as the Autonomous Database instance when you create the compute instance.

   If the application runs in another cloud or in an on-premises data center, use OCI FastConnect to reduce the latency for the connection to your OCI region. See FastConnect Overview for more information.
2. Configure the network routing.
   - If you are using an Autonomous Database instance on a public endpoint, configure your network routing so that the connection from the client to the database goes through a Service Gateway.

     See the following for more information.
     - Setting Up a Service Gateway in the Console
     - Access to Oracle Services: Service Gateway

- If you are using an Autonomous Database instance on a private endpoint, you connect to the database using the private endpoint visible in your network, without the need to configure a Service Gateway.

3. Use one-way TLS connections without a wallet.

   As a best practice for lower latency, configure the Autonomous Database instance to allow both mTLS and TLS connections and use TLS connections to connect your application to the database.

   See the following for more information:

   - Secure Connections to Autonomous Database with mTLS or with TLS
   - Update Network Options to Allow TLS or Require Only Mutual TLS (mTLS) Authentication on Autonomous Database

4. Use TCP Fast Open (TFO) to connect to the database.

   See Use TCP Fast Open (TFO) to Connect Clients to Autonomous Database for more information.

Steps to Reduce Latency for Database Connections for Databases with Autonomous Data Guard

Provides steps to take to configure an Autonomous Data Guard standby environment, clients and mid-tiers, to reduce latency for database connections when you connect after a failover or after a switchover (when the standby becomes the Primary).

- Reduce Latency for Database Connections with Local Autonomous Data Guard
  Follow these steps to reduce the latency for the database connections you make when you use Autonomous Data Guard and you either failover or switchover to a local standby database.
- Reduce Latency for Database Connections with Cross-Region Autonomous Data Guard
  Follow these steps to reduce the latency for the database connections you make when you use Autonomous Data Guard and you either fail over or switch over to a cross-region standby database.

Reduce Latency for Database Connections with Local Autonomous Data Guard

Follow these steps to reduce the latency for the database connections you make when you use Autonomous Data Guard and you either failover or switchover to a local standby database.

If you have an Autonomous Data Guard local standby and you are in a region with multiple availability domains, Autonomous Data Guard creates the local standby database in a different availability domain. When you failover or switchover to the standby database, the local standby becomes the Primary database. To prepare for a failover or a switchover it is recommended to have standby clients and mid-tiers available to enable, so that after a failure or after a switchover, your applications can continue working in the event of an availability domain failure.

First, verify that the disaster recovery type for your local peer is Autonomous Data Guard. See Enable Autonomous Data Guard for more information.

Perform the following tasks to configure standby clients and mid-tiers for low latency when you are using Autonomous Data Guard with a local standby in a region with multiple availability domains.

1. Place the standby clients or mid-tiers in the same availability domain as the local standby database (all components should be configured to use the same availability domain).

For example, if your application runs on an Oracle Cloud Infrastructure Compute VM, when you create the compute instance select the same availability domain for the Compute VM as the standby database. This prepares the disaster recovery configuration so that the standby database and the standby Compute VM use the same availability domain after a failover or switchover. This will reduce latency for connections to the database compared to a configuration where the components use different availability domains.

To determine the availability domain of the standby database, connect to the Primary database as the ADMIN user and run the following query:

```
SELECT availability_domain FROM v$pdbs,
     JSON_TABLE(
       cloud_identity,
       '$.AUTONOMOUS_DATA_GUARD[*]'
       COLUMNS (
         standby_type PATH '$.STANDBY_TYPE',
         availability_domain PATH '$.AVAILABILITY_DOMAIN'
       )
     ) jt
WHERE jt.standby_type = 'local';
```

For example, this command shows the availability domain for a local standby database:

```
AVAILABILITY_DOMAIN
-------------------
SoSC:US-ASHBURN-AD-3
```

2. You do not need to do additional network configuration or allow one-way TLS connections for the local standby database. A local standby database has the same setup network configuration as the Primary database.
3. Configure your clients and mid-tiers to use TCP Fast Open.

   See Use TCP Fast Open (TFO) to Connect Clients to Autonomous Database for more information.

Reduce Latency for Database Connections with Cross-Region Autonomous Data Guard

Follow these steps to reduce the latency for the database connections you make when you use Autonomous Data Guard and you either fail over or switch over to a cross-region standby database.

If you add one or more cross-region Autonomous Data Guard standby databases, the cross-region standby databases are added in the regions that you select when you add a cross-region peer. When you failover or switchover to a cross-region Autonomous Data Guard standby database, the cross-region standby becomes the Primary database. To prepare for a regional failover or switchover, it is recommended to have standby clients and mid-tiers available in the remote region. This prepares the clients and the mid-tier in the remote region so that in the case of a failure or after a switchover, your applications can continue working.

First, verify that your disaster recovery includes at least one cross-region Autonomous Data Guard standby. See Add a Cross-Region Standby Database for more information.

Follow these steps to configure clients and mid-tiers for low latency when using Autonomous Data Guard with one or more cross-region standby databases.

1. Place the standby clients or mid-tiers in the same availability domain as the cross-region standby databases.

To determine the availability domains for cross-region Autonomous Data Guard standby databases, connect to the Primary database as the ADMIN user and run the following query:

```
SELECT availability_domain FROM v$pdbs,
     JSON_TABLE(
       cloud_identity,
       '$.AUTONOMOUS_DATA_GUARD[*]'
       COLUMNS (
         standby_type PATH '$.STANDBY_TYPE',
         availability_domain PATH '$.AVAILABILITY_DOMAIN'
       )
     ) jt
WHERE jt.standby_type = 'cross-region';
```

For example, when you have two cross-region standby databases, running this command shows the availability domains for each cross-region standby database:

```
AVAILABILITY_DOMAIN
---------------------
SoSC:PHX-AD-3
SoSC:US-SANJOSE-1-AD-1
```

a. If you have one cross-region standby, the query shows a single availability domain. Place the standby clients and mid-tiers in the same region and use the same availability domain as the cross-region standby database.

   For example, if your application runs on an Oracle Cloud Infrastructure Compute VM, when you create the compute instance select the same availability domain for the Compute VM as the Autonomous Data Guard standby database. This assures that the cross-region standby database and the standby Compute VM are in the same region and use the same availability domain after a failover or switchover.

b. If you have more than one cross-region standby, in each region use the appropriate availability domain that matches the region and the availability domain for each corresponding standby database. You will need to perform this setup multiple times (all components in an individual region should use the same availability domain as the Autonomous Data Guard standby).

If the application runs in another cloud or in an on-premises data center, use OCI FastConnect to reduce the latency for the connection to your OCI region. See FastConnect Overview for more information.

2. Configure the network routing in the region where the standby database resides. Perform this step multiple times if you have more than one cross-region standby database.

   a. If the standby database is on a public endpoint, configure your network routing so that the connection from the clients in the region where the cross-region standby database is, go through a Service Gateway.

      See the following for more information:

      - Setting Up a Service Gateway in the Console
      - Access to Oracle Services: Service Gateway

   b. If the standby database is on a private endpoint, you connect to the database using the private endpoint visible in your network, without the need to configure a Service Gateway.

3. Use one-way TLS connections without a wallet.

   If you configured one-way TLS for your primary database, standby databases will already have on-way TLS configured. You do not need to do any additional configuration on cross-region standby databases.

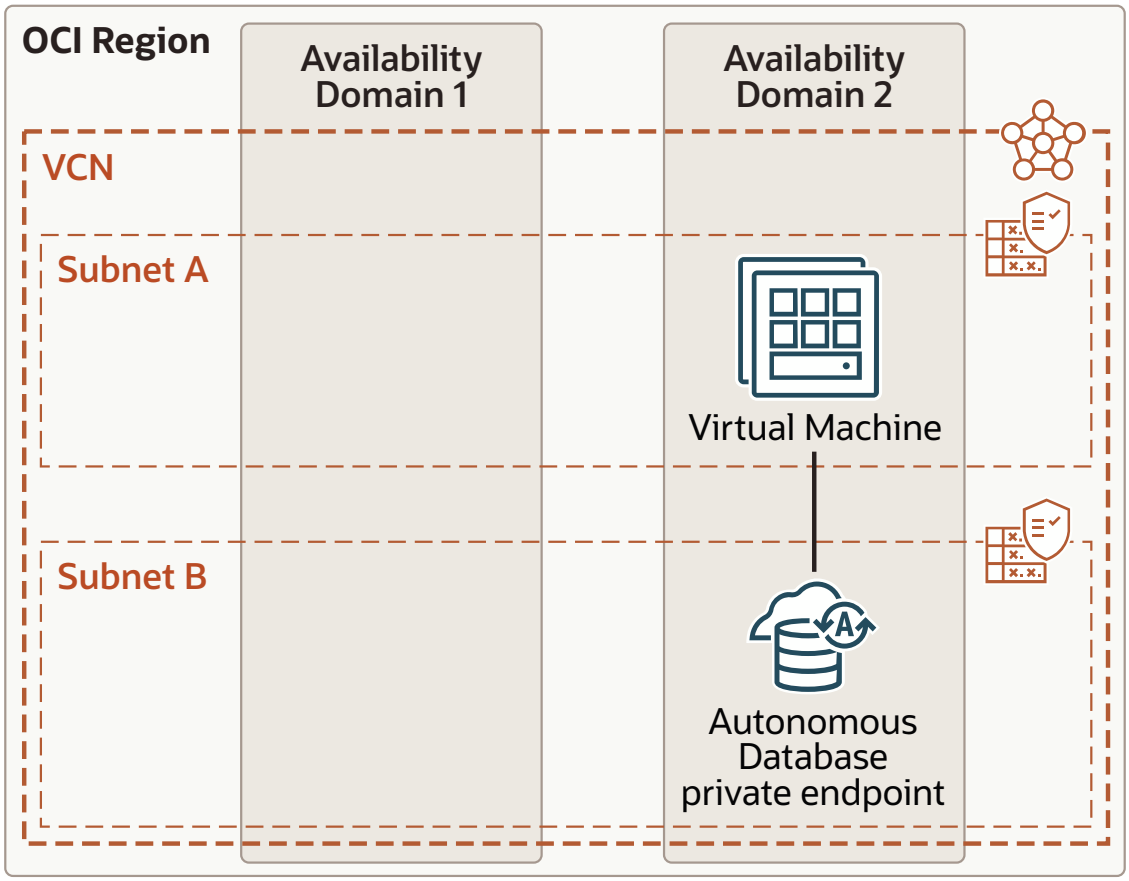
4. Configure your clients and mid-tiers to use TCP Fast Open.

   See Use TCP Fast Open (TFO) to Connect Clients to Autonomous Database for more information.

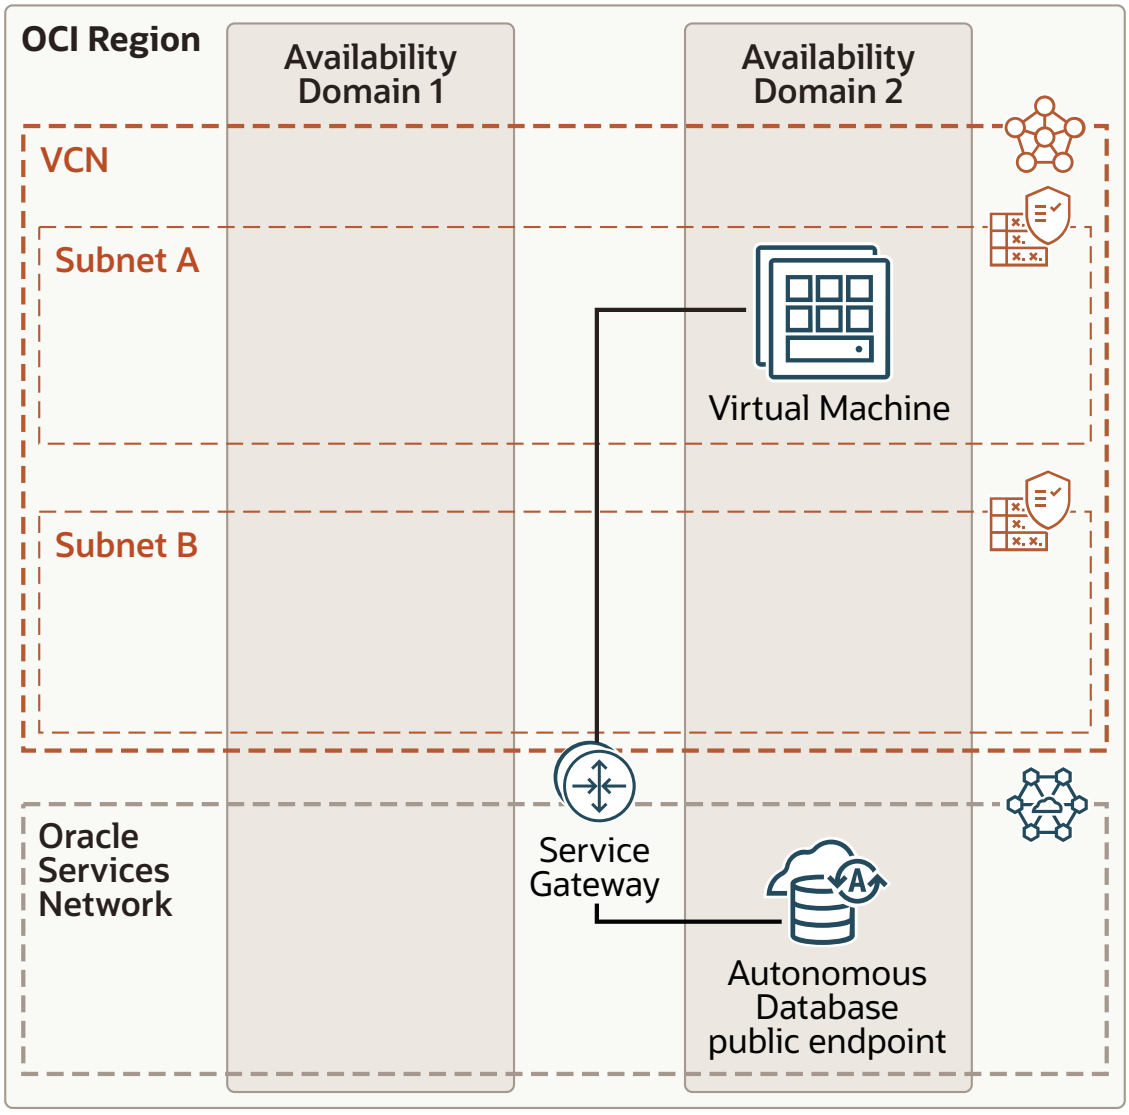
Conceptual Network Diagram for Low Latency Database Connections

Shows the conceptual network diagram for low latency connections using public endpoints and private endpoints for your database.

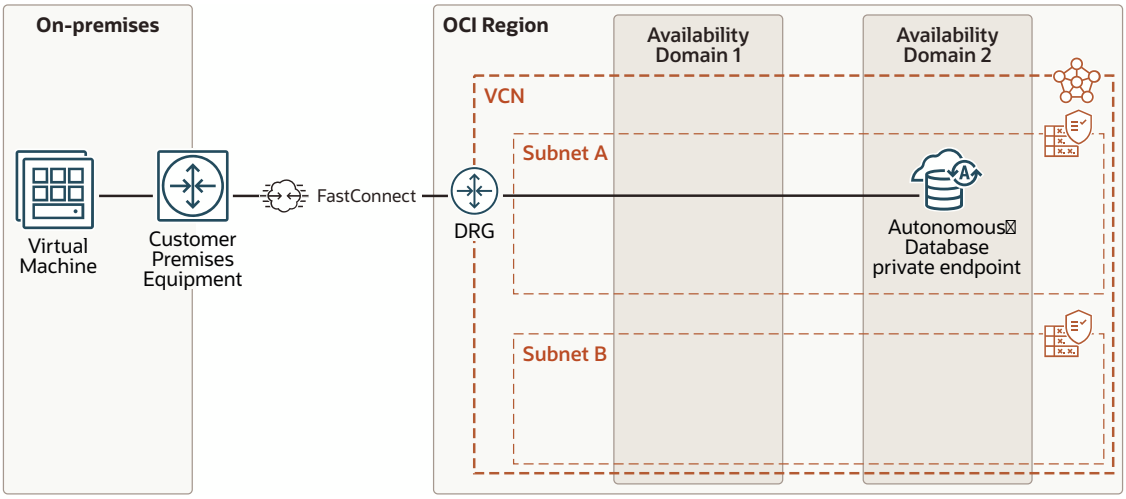
Low Latency Connections Using Private Endpoint with Application Running Inside the OCI Region



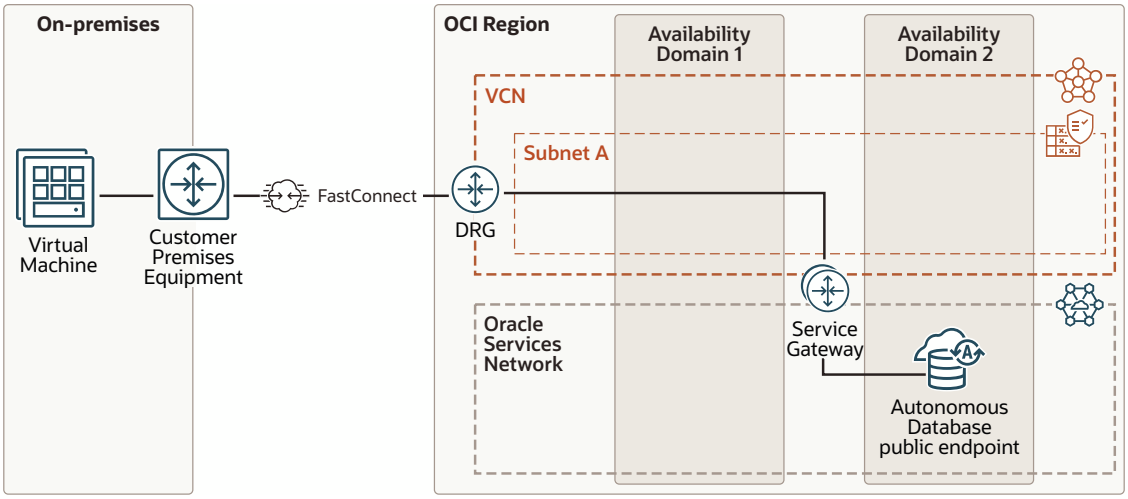
Low Latency Connections Using Public Endpoint with Application Running Inside the OCI Region



Low Latency Connections Using a Private Endpoint with Application Running In On-Premises Data Center Connected to OCI Using FastConnect



Low Latency Connections Using a Public Endpoint with Application Running In Your in On-Premises Data Center Connected to OCI Using FastConnect

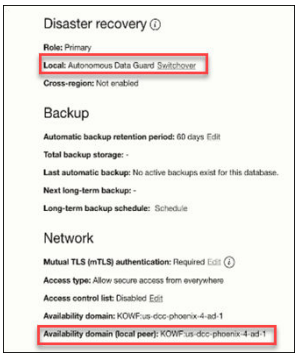


View Network Information on the OCI Console

On the Oracle Cloud Infrastructure Console you can view the network information for your Autonomous Database.

The network information on the OCI Console includes the access type, the availability domain, and other network configuration information for your database.

Field	Description
Mutual TLS	Shows the configuration options for Mutual TLS (mTLS) and TLS for the Autonomous Database instance. See Secure Connections to Autonomous Database with mTLS or with TLS and Update Network Options to Allow TLS or Require Only Mutual TLS (mTLS) Authentication on Autonomous Database for information on Mutual TLS and TLS connections.

Field	Description
Access type	Shows the specified access type for the Autonomous Database instance. See the following for more information: • About Network Access Options • Configure Access Control Lists When You Provision or Clone an Instance • Configure Private Endpoints When You Provision or Clone an Instance
Access control list	This field shows if Access Control Lists (ACLs) of an Autonomous Database Instance are defined. It can display the following two values: • **Enabled**: The field displays Enabled when ACLs of the Autonomous Database Instance are defined. • **Disabled**: The field displays Disabled when ACLs of the Autonomous Database Instance are not defined. See Configuring Network Access with Access Control Rules (ACLs) for more information.
Availability domain	The Availability Domain field displays the availability domain of your Autonomous Database instance. See Regions and Availability Domains for more information on Availability Domains.
Availability domain (local peer)	If there is a local peer database, Autonomous Data Guard standby, this field shows the local standby's availability domain. For example, the following shows the network information with the Autonomous Database instance availability domain and the availability domain for the local Autonomous Data Guard standby:  Disaster recovery Role: Primary Local: Autonomous Data Guard Switchover Cross-region: Not enabled Backup Automatic backup retention period: 60 days Edit Total backup storage: - Last automatic backup: No active backups exist for this database. Next long-term backup: - Long-term backup schedule: Schedule Network Mutual TLS (mTLS) authentication: Required Edit Access type: Allow secure access from everywhere Access control list: Disabled Edit Availability domain: KOWF:us-dcc-phoenix-4-ad-1 Availability domain (local peer): KOWF:us-dcc-phoenix-4-ad-1 You can also view information on Autonomous Data Guard with Local Standby.

The following fields are shown when the Autonomous Database is configured with a private endpoint.

Field	Description
Public access	This field shows if public access for an Autonomous Database Instance on a Private Endpoint is enabled. This field can display the following two values: • **Enabled**: A data recipient on the public internet can access data when the data resides on an Autonomous Database instance on a private endpoint. • **Disabled**: A data recipient on the public internet cannot access data when the data resides on an Autonomous Database instance on a private endpoint.
Virtual cloud network	The Virtual Cloud Network field specifies the VCN for use with an Oracle Cloud Infrastructure Service Gateway. This field includes a link for the VCN associated with the private endpoint. See Access to Oracle Services: Service Gateway.
Subnet	The Subnet field includes a link for the subnet associated with the private endpoint. See VCN and Subnet Management for more information.
Public Endpoint URL	This field shows the public endpoint URL for the resource when public access is enabled.
Private endpoint IP	This field displays the private endpoint IP address for the resource.
Private endpoint URL	This field displays the private endpoint URL for the resource.
Network Security Group	Shows a Network Security Group (NSG) within your VCN. The NSG specifies rules for connections to your Autonomous Database. You can define security rules in an NSG to allow connections to the Autonomous Database instance which creates a virtual firewall for your Autonomous Database. See Network Security Groups for more information.

Manage Users

Describes administration tasks for managing database users on Autonomous Database.

- Create Users on Autonomous Database
  There are several options to create users on Autonomous Database. You can use Oracle Database Actions Database Users card or use client-side tools that connect to the database to create database users.
- Remove Users on Autonomous Database
  To remove users from your database, connect to the database as the ADMIN user using any SQL client tool.
- Manage User Profiles with Autonomous Database
  You can create and alter user profiles in Autonomous Database. After you create or alter a profile, you can specify the profile clause with `CREATE USER` or `ALTER USER`. You can also import existing user profiles from another environment with Oracle Data Pump Import.
- Manage User Roles and Privileges on Autonomous Database
  There are several ways to manage user privileges and roles on Autonomous Database. You can use Oracle Database Actions Database Users card or client-side tools to connect to the database to mange privileges and roles.

- Create and Update User Accounts for Oracle Machine Learning Components on Autonomous Database
  An administrator can add an existing database user account to use with Oracle Machine Learning components or create a new user account and user credentials with the Oracle Machine Learning User Management interface.
- Use Identity and Access Management (IAM) Authentication with Autonomous Database
- Use Microsoft Entra ID with Autonomous Database
  You can configure an Autonomous Database instance for Microsoft Entra ID users to connect using Entra ID (Azure AD) `OAuth2` tokens.
- Use Microsoft Active Directory with Autonomous Database
- Configure Kerberos Authentication with Autonomous Database
  Describes how to configure Kerberos to authenticate Oracle Autonomous Database users.

Create Users on Autonomous Database

There are several options to create users on Autonomous Database. You can use Oracle Database Actions Database Users card or use client-side tools that connect to the database to create database users.

- Create Users on Autonomous Database with Database Actions
  You can quickly create Autonomous Database users with Database Actions.
- Create Users on Autonomous Database - Connecting with a Client Tool
  You can create users by connecting to the database as the ADMIN user using any SQL client tool.
- Unlock User Accounts on Autonomous Database
  If a user account is locked, as the ADMIN user you can unlock the account.
- About User Passwords on Autonomous Database
  Autonomous Database requires strong passwords; the password you specify for a user must meet the minimum default password complexity rules.
- Manage the Administrator Account on Autonomous Database
  The administrator user for Autonomous Database is ADMIN, and this administrator account has several system privileges granted to be able to manage users and other areas of the database.

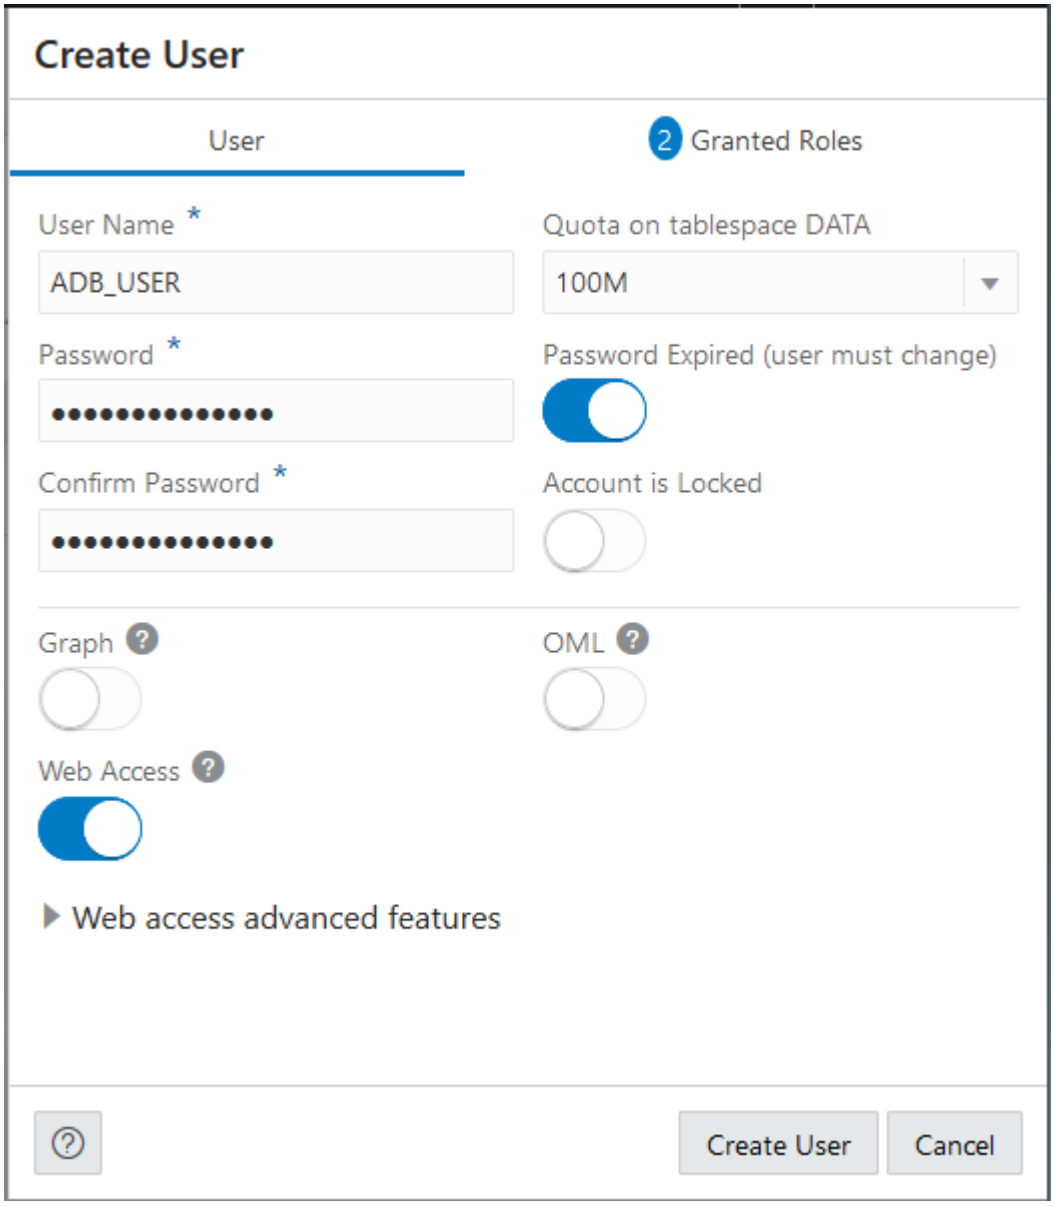
Create Users on Autonomous Database with Database Actions

You can quickly create Autonomous Database users with Database Actions.

First, access Database Actions as the ADMIN user. See Access Database Actions as ADMIN for more information.

1. Click the top left ☰ next to Oracle Database Actions.

   This shows the Database Actions menu, including **Development**, **Data Studio**, **Administration**, **Downloads**, **Monitoring**, and **Related Services**.
2. Under **Administration** click **Database Users**.
3. On the Database Users page, in the All Users area click **+ Create User**.
4. To create a new user, enter a user name, a password, and enter the password again to confirm the password. Also select any options you want to enable for the user: **Graph**, **OML**, or **Web Access**.



5. Set a value for the **Quota on tablespace DATA** for the user.
6. If you want to grant roles for the new user, click the **Granted Roles** tab and select the roles for the user. For example, select DWROLE and CONNECT.
7. Click **Create User**.

   Database Actions shows the User Created confirmation message.

See Manage User Roles and Privileges on Autonomous Database for more information on granting roles and adding or updating privileges for a user.

See The Database Users Page for detailed information on Database Actions Database Users.

If you provide Web Access for the new user, then you need to send a URL to the new user. See Provide Database Actions Access to Database Users for more information.

The administrator needs to provide the credentials wallet to the new user for client-side access. See Connect to Autonomous Database for more information on client-side access credentials.

> **Note:**
>
> Autonomous Database requires strong passwords; the password you specify must meet the default password complexity rules. See About User Passwords on Autonomous Database for more information.

See Create Oracle APEX Workspaces in Autonomous Database for information on creating APEX workspaces.

See Create and Update User Accounts for Oracle Machine Learning Components on Autonomous Database to add user accounts for Oracle Machine Learning Notebooks.

Create Users on Autonomous Database - Connecting with a Client Tool

You can create users by connecting to the database as the ADMIN user using any SQL client tool.

For example, connect using Oracle SQL Developer (see Connect Oracle SQL Developer with a Wallet (mTLS)).

1. Connect as the ADMIN user.
2. Run the following SQL statements:

```
CREATE USER new_user IDENTIFIED BY password;
GRANT CREATE SESSION TO new_user;
```

> **Note:**
>
> `IDENTIFIED` with the `EXTERNALLY` clause is not supported with Autonomous Database.
> In addition, `IDENTIFIED` with the `BY VALUES` clause is not allowed.

This creates *new_user* with connect privileges. This user can now connect to the database and run queries. To grant additional privileges to users, see Manage User Roles and Privileges on Autonomous Database.

The administrator needs to provide the credentials wallet to the user *new_user*. See Connect to Autonomous Database for more information on client credentials.

> **Note:**
>
> Autonomous Database requires strong passwords; the password you specify must meet the default password complexity rules. See About User Passwords on Autonomous Database for more information.

See Provide Database Actions Access to Database Users to add users for Database Actions.

See Create Oracle APEX Workspaces in Autonomous Database for information on creating APEX workspaces.

See Create and Update User Accounts for Oracle Machine Learning Components on Autonomous Database to add user accounts for Oracle Machine Learning components.

Unlock User Accounts on Autonomous Database

If a user account is locked, as the ADMIN user you can unlock the account.

To unlock an account, connect to your database as the ADMIN user and run the following command:

```
ALTER USER username IDENTIFIED BY password ACCOUNT UNLOCK;
```

See *SQL Language Reference* for information on the `ALTER USER` command.

About User Passwords on Autonomous Database

Autonomous Database requires strong passwords; the password you specify for a user must meet the minimum default password complexity rules.

Autonomous Database sets minimum standards for passwords, and the default profile sets parameters to limit the number of failed login attempts.

- The password must be between 12 and 30 characters long and must include at least one uppercase letter, one lowercase letter, and one numeric character.

  Note, the password limit is shown as 60 characters in some help tooltip popups. Limit passwords to a maximum of 30 characters.
- The password cannot contain the *username*.
- The password cannot be one of the last four passwords used for the same *username*.
- The password cannot contain the double quote (") character.
- The password must not be the same password that is set less than 24 hours ago.

To change the password complexity rules and password parameter values you can alter the default profile or create a new profile and assign it to users. See Manage User Profiles with Autonomous Database for more information.

The following are the Autonomous Database default profile password parameter values:

Password Parameter	Description	Value
`FAILED_LOGIN_ATTEMPTS`	The maximum times a user can try to log in and fail before locking the account. This limit applies for regular database user accounts.	10
`PASSWORD_GRACE_TIME`	The number of days after the grace period begins during which a warning is issued and login is allowed.	30
`PASSWORD_LIFE_TIME`	The number of days the same password can be used for authentication.	360
`PASSWORD_LOCK_TIME`	The number of days an account will be locked after the specified number of consecutive failed login attempts.	1
`PASSWORD_REUSE_MAX`	The number of password changes required before the current password can be reused.	4
`PASSWORD_REUSE_TIME`	The number of days before which a password cannot be reused.	1

See Manage User Profiles with Autonomous Database for information on using `CREATE USER` or `ALTER USER` with a profile clause.

See *SQL Language Reference* for information on the `ALTER USER` command.

Manage the Administrator Account on Autonomous Database

The administrator user for Autonomous Database is ADMIN, and this administrator account has several system privileges granted to be able to manage users and other areas of the database.

You can change the administrator user (ADMIN) password and when locked, unlock the administrator user account on Autonomous Database. When you use the APIs to create an Autonomous Database or to reset the ADMIN password, you can optionally use an Oracle Cloud Infrastructure Vault secret to store the password.

See CreateAutonomousDatabase for more information.

See ADMIN User Roles and Privileges for more information on the ADMIN user.

- Set the ADMIN Password in Autonomous Database
  Provides the steps to set the ADMIN password.
- Unlock the ADMIN Account in Autonomous Database
  Shows the steps to unlock the ADMIN user account.
- Use Oracle Cloud Infrastructure Vault Secret for ADMIN Password

Set the ADMIN Password in Autonomous Database

Provides the steps to set the ADMIN password.

From the Oracle Cloud Infrastructure Console, change the password for the ADMIN user by following these steps:

1. On the **Details** page, from the **More actions** drop-down list, select **Administrator password**.
2. On the **Administrator password** panel enter the new password and confirm.
3. Click **Change**.

> **Note:**
>
> You can also use Database Actions to change the password for the ADMIN user. See Manage Users and User Roles on Autonomous Database - Connecting with Database Actions for more information.

The password for the default administrator account, ADMIN, has the same password complexity rules mentioned in the section About User Passwords on Autonomous Database.

Unlock the ADMIN Account in Autonomous Database

Shows the steps to unlock the ADMIN user account.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.

- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

Use the following steps to unlock the ADMIN account by updating the ADMIN password:

1. On the **Details** page, from the **More actions** drop-down list, select **Administrator password**.
2. On the **Administrator password** enter the new password and confirm.
3. Click **Change**.

This operation unlocks the ADMIN account if it was locked.

The password for the default administrator account, ADMIN, has the same password complexity rules mentioned in the section About User Passwords on Autonomous Database.

Use Oracle Cloud Infrastructure Vault Secret for ADMIN Password

When you create or clone an Autonomous Database instance or when you reset the ADMIN password, you can use an Oracle Cloud Infrastructure vault secret to specify the ADMIN password.

Autonomous Database allows you to use the APIs to provide a protected vault secret as the ADMIN password, with secure access to the vault secret granted through Oracle Cloud Infrastructure IAM policies.

> **Note:**
>
> Using an Oracle Cloud Infrastructure vault secret for the ADMIN password is only supported with the APIs.

Oracle Cloud Infrastructure Vault secrets are credentials that you use with Oracle Cloud Infrastructure services. Storing secrets in a vault provides greater security than you might achieve storing them elsewhere, such as in code or in configuration files. By calling database APIs you can use secrets from the Vault Service to set the ADMIN password. The vault secret password option is available when you create or clone an Autonomous Database instance, or when you set or reset the ADMIN password.

You create secrets using the Oracle Cloud Infrastructure Console, CLI, or API.

Notes for using a vault secret to set or reset the ADMIN password:

- In order for Autonomous Database to reach the secret in a vault, the following conditions must apply:
  - The secret must be in `current` or `previous` state.
  - If you specify a secret version in the API call, the specified secret version is used. If you do not specify a secret version, the call uses the latest secret version.
  - You must have the proper user group policy that allows `READ` access to the specific secret in a given compartment. For example:

```
Allow userGroup1 to read secret-bundles in compartment training
```

- The password stored in the secret must conform to Autonomous Database password requirements.

See the following for more information:

- Overview of Vault
- CreateAutonomousDatabase
- UpdateAutonomousDatabaseDetails
- About User Passwords on Autonomous Database
- Managing Secrets

Remove Users on Autonomous Database

To remove users from your database, connect to the database as the ADMIN user using any SQL client tool.

For example, connect using Oracle SQL Developer.

1. Connect as the ADMIN user.
2. Run the following SQL statement:

```
DROP USER user_name CASCADE;
```

This removes *user_name* and the objects owned by that user.

> **Note:**
>
> This removes all *user_name* objects and the data owned by *user_name* is deleted.

You can also remove a database user with Oracle Database Actions. To remove a user, from the Database Users page, click the ⋮ to open the context menu for a user, and then select **Delete**. See Manage Users and User Roles on Autonomous Database - Connecting with Database Actions for more information on using Oracle Database Actions Database Users.

Manage User Profiles with Autonomous Database

You can create and alter user profiles in Autonomous Database. After you create or alter a profile, you can specify the profile clause with `CREATE USER` or `ALTER USER`. You can also import existing user profiles from another environment with Oracle Data Pump Import.

> **Note:**
>
> Autonomous Database has restrictions on the profile clause. See SQL Commands for information on `CREATE PROFILE` and `ALTER PROFILE` restrictions.

To add, modify, or remove a password parameter in a profile, including the `DEFAULT` profile you must have the `ALTER PROFILE` system privilege.

1. To add or alter a profile, as the ADMIN user run either `CREATE PROFILE` or `ALTER PROFILE`. For example:

   ```
   CREATE PROFILE new_profile
     LIMIT PASSWORD_REUSE_MAX 10
     PASSWORD_LOCK_TIME 5;
   ```

   If you are not the ADMIN user, then you must have `CREATE PROFILE` privilege to run `CREATE PROFILE`. If you run `ALTER PROFILE`, then you must have `ALTER PROFILE` privilege.

2. Use the new or altered profile with a `CREATE USER` or `ALTER USER` command. For example:

   ```
   CREATE USER new_user IDENTIFIED BY password PROFILE new_profile;
   GRANT CREATE SESSION TO new_user;
   ```

This creates *new_user* with profile *new_profile* and with connect privileges. The *new_user* can now connect to the database and run queries. To grant additional privileges to users, see Manage User Privileges on Autonomous Database - Connecting with a Client Tool.

See CREATE PROFILE for information on using `CREATE PROFILE` or `ALTER PROFILE`.

You can import existing profiles created in other environments using Oracle Data Pump Import (`impdp`). Any existing profile association with database users is preserved after importing into Autonomous Database. When a newly created user, created from an Oracle Data Pump import, attempts to login for the first time, the login is handled as follows:

- The password complexity restrictions are the same as the restrictions for any user on Autonomous Database.
- If the user's password violates the password complexity requirements, the account is expired with a 30-day grace period. In this case the user is required to change their password before the grace period ends.

> **Note:**
>
> Profile assignments for users with profile `ORA_PROTECTED_PROFILE` and `ORA_ADMIN_PROFILE` cannot be modified.

The following users share the `ORA_PROTECTED_PROFILE` profile and the profile assignment of these users cannot be changed:

- `ADBSNMP`
- `ADB_APP_STORE`
- `DCAT_ADMIN`
- `GGADMIN`
- `RMAN$CATALOG`

The `ADMIN` user is assigned to `ORA_ADMIN_PROFILE`.

When you create or alter a profile, you can specify a Password Verification Function (PVF) to manage password complexity. See Manage Password Complexity on Autonomous Database for more information.

- Manage Password Complexity on Autonomous Database
  You can create a Password Verify Function (PVF) and associate the PVF with a profile to manage the complexity of user passwords.
- Gradual Database Password Rollover for Applications

Manage Password Complexity on Autonomous Database

You can create a Password Verify Function (PVF) and associate the PVF with a profile to manage the complexity of user passwords.

> **Note:**
>
> The minimum password length for a user specified PVF is 8 characters and must include at least one upper case letter, one lower case letter and one numeric character. The minimum password length for the DEFAULT profile is 12 characters (the DEFAULT profile uses the `CLOUD_VERIFY_FUNCTION` PVF). The password cannot contain the username.
> Oracle recommends using a minimum password length of 12 characters. If you define a profile's PVF, and set the minimum password length to less than 12 characters, then tools such as Oracle Database Security Assessment Tool (DBSAT) and Qualys report this as a database security risk.

For example, to specify a PVF for a profile, use the following command:

```
CREATE PROFILE example_profile LIMIT PASSWORD_VERIFY_FUNCTION
ADMIN.EXAMPLE_PVF
```

If the profile is created or altered by any user other than the ADMIN user, then you must grant the `EXECUTE` privilege on the PVF. If you create a PVF and the password check fails, the database reports the `ORA-28219` error.

You can specify an Oracle supplied PVF, from one of the following:

- `CLOUD_VERIFY_FUNCTION` (this is the default password verification function for Autonomous Database):
  This function checks for the following requirements when users create or modify passwords:
  - The password must be between 12 and 30 characters long and must include at least one uppercase letter, one lowercase letter, and one numeric character.
  - The password cannot contain the username.
  - The password cannot be one of the last four passwords used for the same username.
  - The password cannot contain the double quote (") character.
  - The password must not be the same password that is set less than 24 hours ago.
- `ORA12C_STIG_VERIFY_FUNCTION`
  This function checks for the following requirements when users create or modify passwords:
  - The password has at least 15 characters.
  - The password has at least 1 lower case character and at least 1 upper case character.

  - The password has at least 1 digit.
  - The password has at least 1 special character.
  - The password differs from the previous password by at least 8 characters.

  See ora12c_stig_verify_function Password Requirements for more information.

Note the following restrictions for a Password Verification Function (PVF) that you create and assign to a profile:

- If you specify a user profile, the minimum password length depends on how you define the associated PVF, as follows:
  - If a PVF is defined, then the minimum password length enforced is 8 characters with at least one upper case letter, one lower case letter and one numeric character. The password cannot contain the username.
  - If the PVF is defined as `NULL`, then the minimum password length enforced is 8 characters with at least one upper case letter, one lower case letter and one numeric character. The password cannot contain the username.
  - If the profile does not have a PVF defined, then the DEFAULT profile's PVF (`CLOUD_VERIFY_FUNCTION`) is assigned and the minimum password length enforced is 12 characters.
- If you specify a Password Verify Function (PVF) that is more strict than the default `CLOUD_VERIFY_FUNCTION`, then the new verify function is used.
- A PVF that you create must be created as a `DEFINER RIGHTS` PL/SQL function. If a `INVOKER` rights PVF is provided as input to `CREATE` or `ALTER PROFILE`, then `ORA-28220` error is thrown.
- Any PVF that you create must be created in the ADMIN user schema. If a non-ADMIN user owned PVF is provided as input to `CREATE` or `ALTER PROFILE`, then `ORA-28220` error is thrown.
- A PVF cannot be altered or dropped by a non-ADMIN user. That is, any user with the `CREATE` or `DROP ANY PROCEDURE` privilege is not allowed to alter or drop a PVF.
- If the PVF associated with a profile is dropped, then any attempt to change the password for a user who uses the PVF in their profile throws the error `ORA-7443`. Users can still login when the PVF associated with their profile is dropped. However, if a user's password is expired and the PVF is dropped, then the user cannot login.

  To recover from the `ORA-7443` error, the ADMIN user must recreate the dropped PVF and assign it to the profile, or assign an existing PVF to the profile. This allows a user change their password and login.
- The `CREATE ANY PROCEDURE` system privilege and `DROP ANY PROCEDURE` system privilege are audited for PVF security. See the `PROCEDURES` list in Listings of System and Object Privileges for more information.

See Managing the Complexity of Passwords for more information.

Gradual Database Password Rollover for Applications

An application can change its database passwords without an administrator having to schedule downtime.

To accomplish this, you can associate a profile having a non-zero limit for the `PASSWORD_ROLLOVER_TIME` password profile parameter, with an application schema. This allows the database password of the application user to be altered while allowing the older password

to remain valid for the time specified by the `PASSWORD_ROLLOVER_TIME` limit. During the rollover period of time, the application instance can use either the old password or the new password to connect to the database server. When the rollover time expires, only the new password is allowed.

See Managing Gradual Database Password Rollover for Applications for more information.

Manage User Roles and Privileges on Autonomous Database

There are several ways to manage user privileges and roles on Autonomous Database. You can use Oracle Database Actions Database Users card or client-side tools to connect to the database to mange privileges and roles.

- Manage Users and User Roles on Autonomous Database - Connecting with Database Actions
  You can manage user roles for Autonomous Database users with Oracle Database Actions. The same steps also let you modify account settings for a user.
- Manage User Privileges on Autonomous Database - Connecting with a Client Tool
  Autonomous Databases come with a predefined database role named `DWROLE`. This role provides the common privileges for Autonomous Database users. Depending on the usage requirements you may also need to grant individual privileges to users.

Manage Users and User Roles on Autonomous Database - Connecting with Database Actions

You can manage user roles for Autonomous Database users with Oracle Database Actions. The same steps also let you modify account settings for a user.

First, access Database Actions as the ADMIN user. See Access Database Actions as ADMIN for more information.

1. Click the top left ☰ next to Oracle Database Actions.

   This shows the Database Actions menu, including **Development** and **Administration**.
2. Under **Administration** click **Database Users**.
3. On the Database Users page, in the card for the user you want to modify click the ⋮ to open the context menu for the user, then select **Edit**.

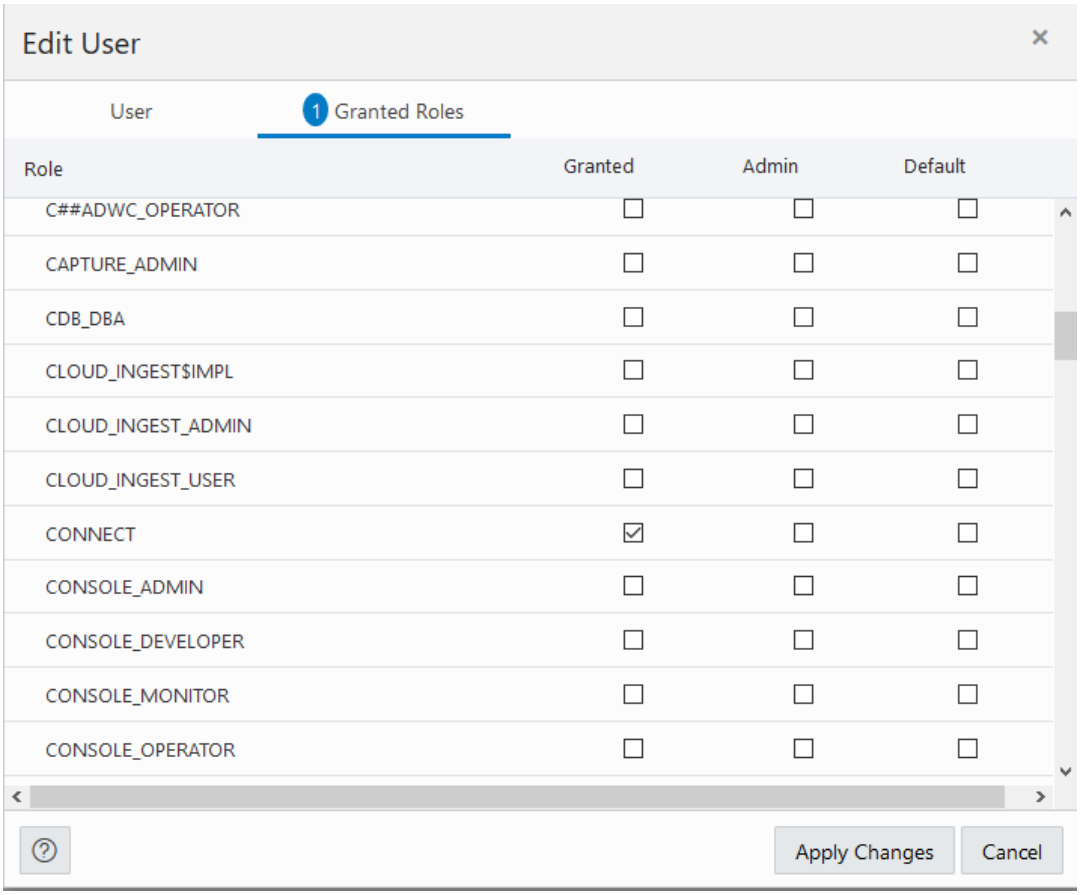
   This shows the Edit User area with the User tab selected.

   > **Note:**
   >
   > If you want to manage the user's account settings, for example if you want to provide Web Access to provide access to Database Actions, or if you want to lock the user's account, you can do this from the User tab.
4. In the Edit User area, click **Granted Roles**.

   This displays the Granted Roles tab with a list of available roles and selection boxes. For each role, you can check **Granted** to grant the role, **Admin** to permit the user to grant the role to other users, and **Default** to use the default settings for **Granted** and **Admin**.
5. Select the roles you want to grant to the user.

   For example, select CONNECT and DWROLE.

For each role, you can select **Granted** to grant the role, **Admin** to permit the user to grant the role to other users, and **Default** to use the default settings for Granted and Admin. A new user is granted CONNECT and RESOURCE roles when **Web Access** is selected.



6. Click **Apply Changes**.

See The Database Users Page for more information on Database Actions Database Users.

See Create Users on Autonomous Database with Database Actions for information on using Database Actions.

Manage User Privileges on Autonomous Database - Connecting with a Client Tool

Autonomous Databases come with a predefined database role named `DWROLE`. This role provides the common privileges for Autonomous Database users. Depending on the usage requirements you may also need to grant individual privileges to users.

1. To grant `DWROLE` role, connect to the database as ADMIN user using any SQL client tool. For example, connect using Oracle SQL Developer (see Connect Oracle SQL Developer with a Wallet (mTLS)).
2. As the ADMIN user grant `DWROLE`. For example, the following command grants `DWROLE` to the user `adb_user`:

```
GRANT DWROLE TO adb_user;
```

3. Grant individual privileges to users with the `GRANT` command instead of or in addition to granting `DWROLE` privileges. See *Oracle Database SQL Language Reference*.

4. If a user needs to load data, do one of the following to add the privileges required to load data:

   - Add quota to a new user with `CREATE USER` or alter the quota for an existing user with `ALTER USER`. For example:

     ```
     CREATE USER sales
          QUOTA 5M on DATA;
     ALTER USER sales
          QUOTA 1G on DATA;
     ```

   - Grant `UNLIMITED TABLESPACE` privileges to a user. For example, the following command grants unlimited tablespace privileges to the user `adb_user`:

     ```
     GRANT UNLIMITED TABLESPACE TO adb_user;
     ```

> **Note:**
>
> Granting `UNLIMITED TABLESPACE` privilege allows a user to use all the allocated storage space. You cannot selectively revoke tablespace access from a user with the `UNLIMITED TABLESPACE` privilege. You can grant selective or restricted access only after revoking the privilege.

The privileges in `DWROLE` are the following:

```
CREATE ANALYTIC VIEW
CREATE ATTRIBUTE DIMENSION
ALTER SESSION
CREATE HIERARCHY
CREATE JOB
CREATE MATERIALIZED VIEW
CREATE MINING MODEL
CREATE PROCEDURE
CREATE SEQUENCE
CREATE SESSION
CREATE SYNONYM
CREATE TABLE
CREATE TRIGGER
CREATE TYPE
CREATE VIEW
READ,WRITE ON directory DATA_PUMP_DIR
EXECUTE privilege on the PL/SQL package DBMS_CLOUD
EXECUTE privilege on OCI PL/SQL SDK
```

Create and Update User Accounts for Oracle Machine Learning Components on Autonomous Database

An administrator can add an existing database user account to use with Oracle Machine Learning components or create a new user account and user credentials with the Oracle Machine Learning User Management interface.

- Create User
  An administrator creates new user accounts and user credentials for Oracle Machine Learning in Database Actions.

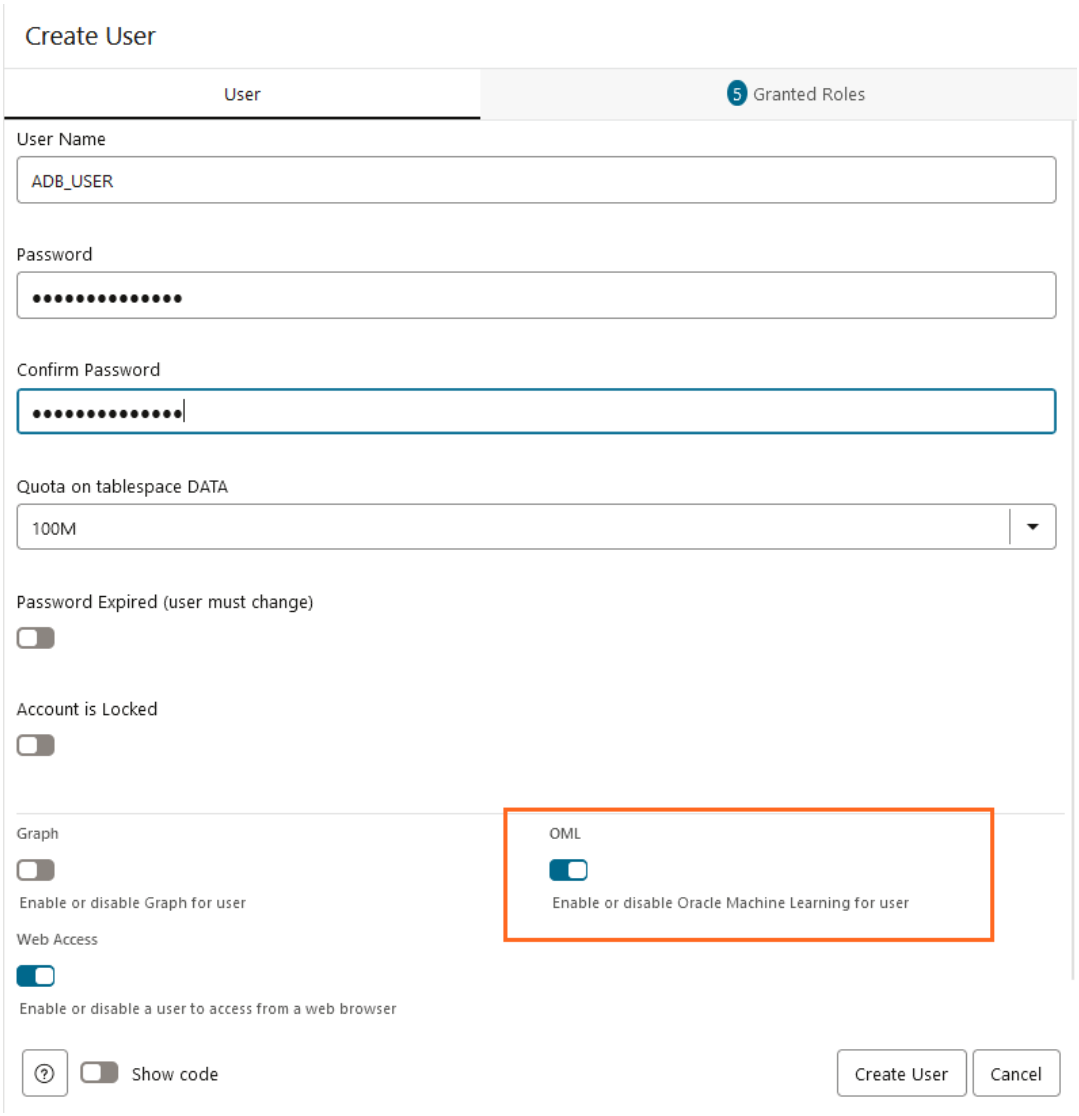
- Add Existing Database User Account to Oracle Machine Learning Components
  As the ADMIN user you can add an existing database user account to provide access to Oracle Machine Learning components.

Create User

An administrator creates new user accounts and user credentials for Oracle Machine Learning in Database Actions.

To create a user account:

1. On the Autonomous Databases page, under the **Display Name** column, select an Autonomous Database.
2. On the Autonomous Database Details page, select **Database Actions** and click **Database Users**.
3. On the Database Users page, in the All Users area click **+ Create User**.
4. To create a new user enter a user name, a password, and enter the password again to confirm the password.
5. Select the options you want for the user and select **OML** to enable Oracle Machine Learning for the user.



6. Click **Create User**.

This creates a new database user and grants the required privileges to use Oracle Machine Learning.

> **Note:**
>
> With a new database user, an administrator needs to issue grant commands on the database to grant table access to the new user for the tables associated with the user's Oracle Machine Learning notebooks.

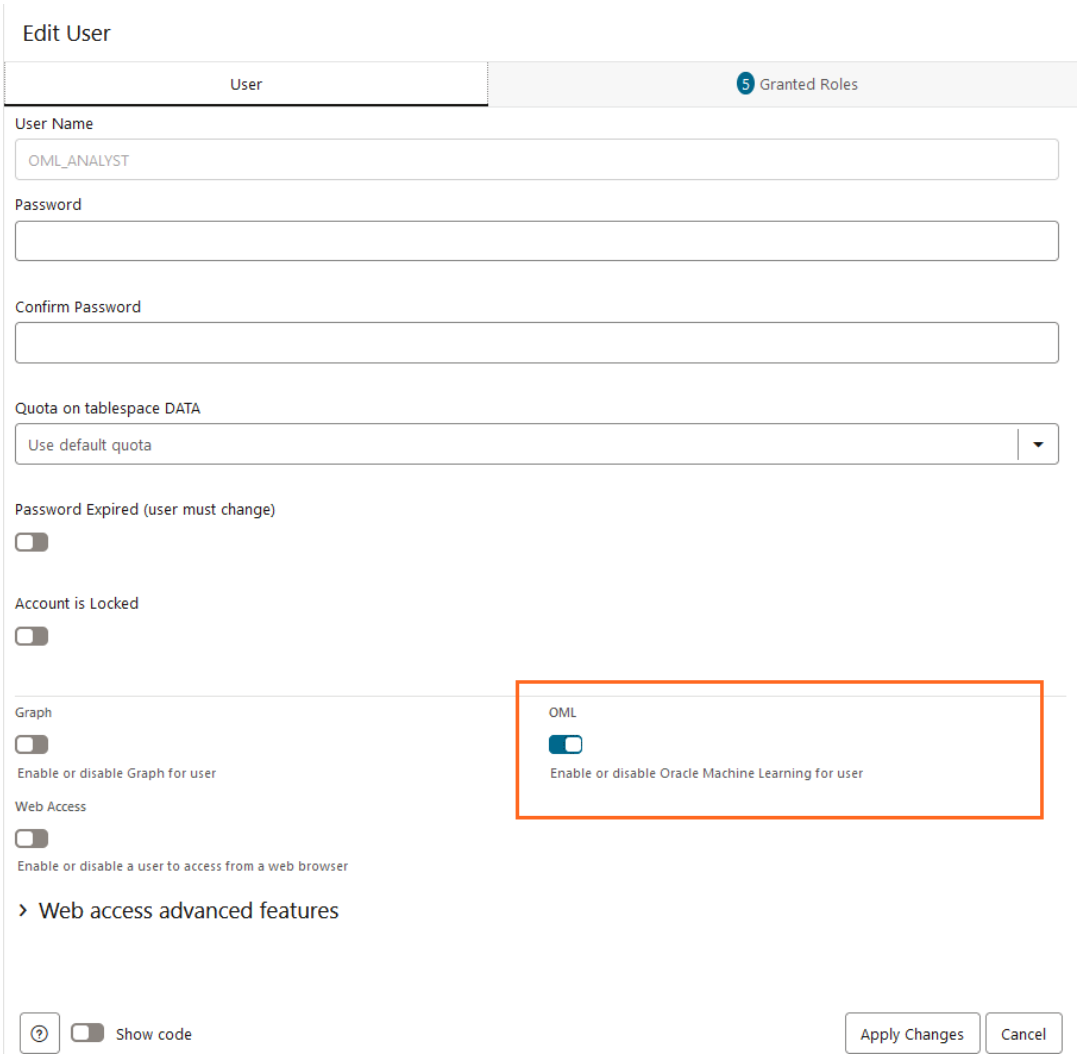
Add Existing Database User Account to Oracle Machine Learning Components

As the ADMIN user you can add an existing database user account to provide access to Oracle Machine Learning components.

To add an existing database user account:

1. On the Autonomous Databases page, under the **Display Name** column, select an Autonomous Database.
2. On the Autonomous Database Details page, select **Database Actions** and click **Database Users**.
3. In the All Users, search for the user of interest or select the user. For example, search the user **OML_ANALYST**.
4. In the user's card, click  and select **Edit**
5. In the Edit User panel, select **OML**.

   For example:



6. Click **Apply Changes**.

This grants the required privileges to use the Oracle Machine Learning application. In Oracle Machine Learning this user can then access any tables the user has privileges to access in the database.

Use Identity and Access Management (IAM) Authentication with Autonomous Database

You can configure Autonomous Database to use Oracle Cloud Infrastructure Identity and Access Management (IAM) authentication and authorization to allow IAM users to access an Autonomous Database with IAM credentials.

> **Note:**
>
> Autonomous Database integration with Oracle Cloud Infrastructure IAM is supported in commercial regions with identity domains as well as in the legacy IAM, which does not include identity domains. IAM with identity domains was introduced with new Oracle Cloud Infrastructure tenancies that were created after November 8, 2021. Autonomous Database supports users and groups in default and non-default identity domains.

- About Identity and Access Management (IAM) Authentication with Autonomous Database
  You can enable an Autonomous Database instance to use Oracle Cloud Infrastructure (IAM) authentication and authorization for users.
- Prerequisites for Identity and Access Management (IAM) Authentication on Autonomous Database
  Describes the prerequisites for enabling IAM user access on Autonomous Database.
- Enable Identity and Access Management (IAM) Authentication on Autonomous Database
  Describes the steps to enable IAM user access on Autonomous Database.
- Create Identity and Access Management (IAM) Groups and Policies for IAM Users
  Describes the steps to write policy statements for an IAM group to enable IAM user access to Oracle Cloud Infrastructure resources, specifically Autonomous Database instances.
- Add IAM Users on Autonomous Database
  To add IAM users to allow access to Autonomous Database, map database global users to IAM groups or users with `CREATE USER` or `ALTER USER` statements (with `IDENTIFIED GLOBALLY AS` clause).
- Add IAM Roles on Autonomous Database
  Optionally, create global roles to provide additional database roles and privileges to IAM users when multiple IAM users are mapped to the same shared global user.
- Create IAM Database Password for IAM Users
  To add an IAM user and allow the IAM user to login to Autonomous Database by supplying a username and password, you must create an IAM database password.
- Connect to Autonomous Database with Identity and Access Management (IAM) Authentication
  After the ADMIN user enables Oracle Cloud Infrastructure IAM on Autonomous Database, users log in to the Autonomous Database instance using their Oracle Cloud Infrastructure IAM credentials or access the database through an Oracle Cloud Infrastructure IAM database token.
- Configure IAM Proxy Authentication
  Proxy authentication allows an IAM user to proxy to a database schema for tasks such as application maintenance.

- Disable Identity and Access Management (IAM) Authentication on Autonomous Database
  Describes the steps to disable IAM external authentication user access for Autonomous Database.
- Accessing a Database Cross-Tenancy Using an IAM Integration
  Users and groups in one tenancy can access DBaaS database instances in another tenancy if policies in both tenancies allow this.
- Notes for Using Autonomous Database Tools with Identity and Access Management (IAM) Authentication
  Provides notes for using Autonomous Database tools with IAM authentication enabled.

About Identity and Access Management (IAM) Authentication with Autonomous Database

You can enable an Autonomous Database instance to use Oracle Cloud Infrastructure (IAM) authentication and authorization for users.

> **Note:**
>
> Autonomous Database integration with Oracle Cloud Infrastructure IAM is supported in commercial regions with identity domains as well as in the legacy IAM, which does not include identity domains. IAM with identity domains was introduced with new Oracle Cloud Infrastructure tenancies that were created after November 8, 2021. Autonomous Database supports users and groups in default and non-default identity domains.

Oracle Cloud Infrastructure IAM integration with Autonomous Database supports the following:

- IAM Database Password Authentication
- Identity and Access Management (IAM) SSO Token Based Authentication

See Authenticating and Authorizing IAM Users for Oracle Autonomous Databases for complete details about the architecture for using IAM users on Autonomous Database.

- IAM Database Password Authentication
  You can enable an Autonomous Database instance to allow user access with an Oracle Cloud Infrastructure IAM database password (using a password verifier).
- Identity and Access Management (IAM) SSO Token Based Authentication
  You can enable an Autonomous Database instance to use Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) SSO tokens.

IAM Database Password Authentication

You can enable an Autonomous Database instance to allow user access with an Oracle Cloud Infrastructure IAM database password (using a password verifier).

> **Note:**
>
> Any supported 12c and above database client can be used for IAM database password access to Autonomous Database.

An Oracle Cloud Infrastructure IAM database password allows an IAM user to log in to an Autonomous Database instance as Oracle Database users typically log in with a user name and password. The user enters their IAM user name and IAM database password. An IAM database password is a different password than the Oracle Cloud Infrastructure Console password. Using an IAM user with the password verifier you can login to Autonomous Database with any supported database client.

For password verifier database access, you create the mappings for IAM users and OCI applications to the Autonomous Database instance. The IAM user accounts themselves are managed in IAM. The user accounts and user groups can be in either the default domain or in a custom, non-default domain.

Identity and Access Management (IAM) SSO Token Based Authentication

You can enable an Autonomous Database instance to use Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) SSO tokens.

For token verifier database access, you create the mappings for IAM users and OCI applications to the Autonomous Database instance. The IAM user accounts themselves are managed in IAM. The user accounts and user groups can be in either the default domain or in a custom, non-default domain.

There are several ways a database client can obtain an IAM database token:

- A client application or tool can request the database token from IAM for the user and can pass the database token through the client API. Using the API to send the token overrides other settings in the database client. Using IAM tokens requires the latest Oracle Database client 19c (at least 19.16). Some earlier clients provide a limited set of capabilities for token access.
- If the application or tool does not support requesting an IAM database token through the client API, the IAM user can first use Oracle Cloud Infrastructure command line interface (CLI) to retrieve the IAM database token and save it in a file location. For example, to use SQL*Plus and other applications and tools using this connection method, you first obtain the database token using the Oracle Cloud Infrastructure (OCI) Command Line Interface (CLI). If the database client is configured for IAM database tokens, when a user logs in with the slash login form, the database driver uses the IAM database token that has been saved in a default or specified file location.
- A client application or tool can use an Oracle Cloud Infrastructure IAM instance principal or resource principal to get an IAM database token, and use the IAM database token to authenticate itself to an Autonomous Database instance.
- IAM users and OCI applications can request a database token from IAM with several methods, including using an API-key. See Configuring a Client Connection for SQL*Plus That Uses an IAM Token for an example. See About Authenticating and Authorizing IAM Users for an Oracle Autonomous Database for a description of other methods such as using a delegation token within an OCI cloud shell.

Prerequisites for Identity and Access Management (IAM) Authentication on Autonomous Database

Describes the prerequisites for enabling IAM user access on Autonomous Database.

If the database is enabled for another external authentication scheme, verify that you want to use IAM on the Autonomous Database instance. There can only be one external authentication scheme enabled at any given time.

If you want to use IAM and another external authentication scheme is enabled, you can either first disable the other external authentication scheme or use the `force` parameter with the value set to true when you enable IAM authentication. See Enable Identity and Access Management (IAM) Authentication on Autonomous Database for information on using the `force` parameter.

Enable Identity and Access Management (IAM) Authentication on Autonomous Database

Describes the steps to enable IAM user access on Autonomous Database.

> **Note:**
>
> Autonomous Database integration with Oracle Cloud Infrastructure IAM is supported in commercial regions with identity domains as well as in the legacy IAM, which does not include identity domains. IAM with identity domains was introduced with new Oracle Cloud Infrastructure tenancies that were created after November 8, 2021. Autonomous Database supports users and groups in default and non-default identity domains.

To enable Autonomous Database to allow IAM users to connect to the database:

1. Perform the prerequisites for IAM authorization and authentication on Autonomous Database. See Prerequisites for Identity and Access Management (IAM) Authentication on Autonomous Database for more information.
2. Use the procedure `DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION` to enable Oracle Cloud Infrastructure IAM authentication.

   When you perform these steps, connect to the Autonomous Database instance as the ADMIN user or as a user with ADMIN privileges.

   For example:

   ```
   BEGIN
      DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
         type => 'OCI_IAM' );
   END;
   /
   ```

   By default the `force` parameter is false. When another external authentication method is enabled and `force` is false, `DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION` reports the following error:

   ```
   ORA-20004: Another external authentication is already enabled.
   ```

   If you want to disable the external authentication that is currently enabled and use IAM authentication instead, include the `force` parameter.

   For example:

   ```
   BEGIN
      DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
         type => 'OCI_IAM',
   ```

```
      force => TRUE );
END;
/
```

This sets the `IDENTITY_PROVIDER_TYPE` system parameter.

For example, you can use the following to verify `IDENTITY_PROVIDER_TYPE`:

```
SELECT NAME, VALUE FROM V$PARAMETER WHERE NAME='identity_provider_type';

NAME                    VALUE
----------------------- -------
identity_provider_type OCI_IAM
```

Create Identity and Access Management (IAM) Groups and Policies for IAM Users

Describes the steps to write policy statements for an IAM group to enable IAM user access to Oracle Cloud Infrastructure resources, specifically Autonomous Database instances.

A policy is a group of statements that specifies who can access particular resources, and how. Access can be granted for the entire tenancy, databases in a compartment, or individual databases. This means you write a policy statement that gives a specific group a specific type of access to a specific type of resource within a specific compartment.

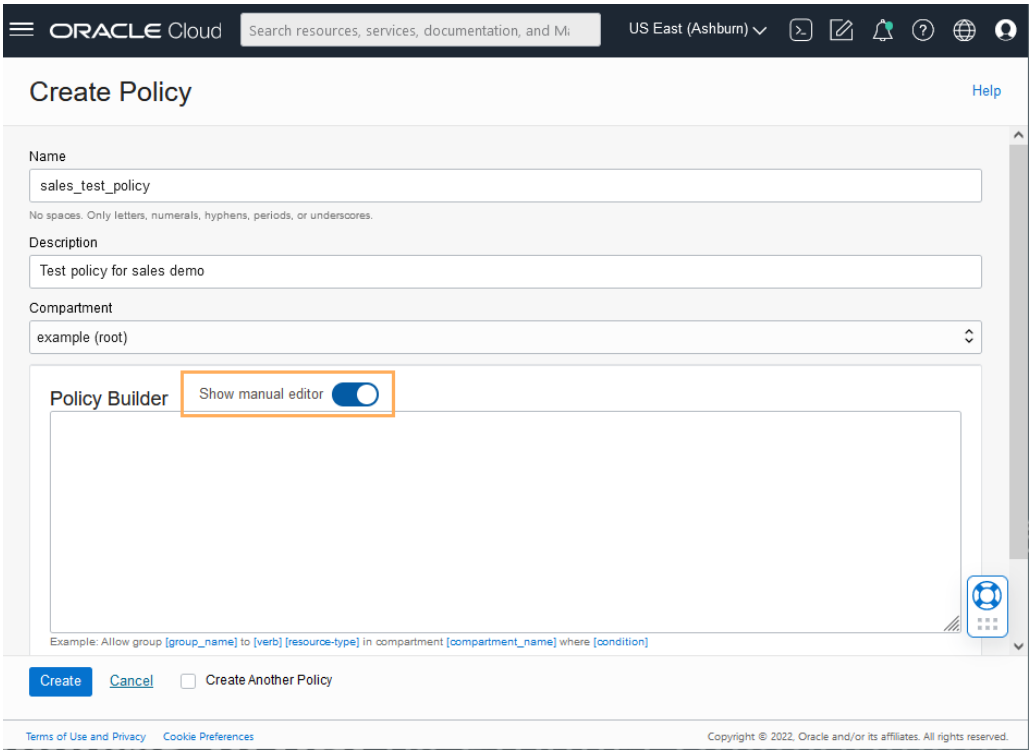
> **Note:**
>
> Defining a policy is required to use IAM tokens to access Autonomous Database. A policy is not required when using IAM database passwords to access Autonomous Database.

To enable Autonomous Database to allow IAM users to connect to the database using IAM tokens:

1. Perform Oracle Cloud Infrastructure Identity and Access Management prerequisites by creating a group and adding users to the group.

   For example, create the group `sales_dbusers`.

   See Managing Groups for more information.
2. Write policy statements to enable access to Oracle Cloud Infrastructure resources.
   a. In the Oracle Cloud Infrastructure console click **Identity & Security**.
   b. Under **Identity & Security** click **Policies**.
   c. To a write policy, click **Create Policy**.
   d. On the Create Policy page, enter a Name and a Description.
   e. On the Create Policy page, select **Show manual editor**.



f. Use the **Policy Builder** to create a policy.

For example to create a policy to allow users in IAM group `DBUsers` to access any Autonomous Database in their tenancy:

```
Allow group DBUsers to use autonomous-database-family in tenancy
```

For example to create a policy that limits members of `DBUsers` group to access Autonomous Databases in compartment `testing_compartment` only:

```
allow group DBUsers to use autonomous-database-family in compartment
testing_compartment
```

For example to create a policy that limits group access to a single database in a compartment:

```
allow group DBUsers to use autonomous-database-family in compartment
testing_compartment
               where target.id =
'ocid1.autonomousdatabase.oc1.iad.aaaabbbbcccc...b5678ca'
```

Refer to Creating an IAM Policy to Authorize Users Authenticating with Tokens in *Database Security Guide* for more information on IAM policies to access the database.

g. Click **Create**.

See Managing Policies for more information on policies.

Notes for creating policies for use with IAM users on Autonomous Database:

- Policies can allow IAM users to access Autonomous Database instances across the entire tenancy, in a compartment, or can limit access to a single Autonomous Database instance.

- You can use either instance principal or resource principal to retrieve database tokens to establish a connection from your application to an Autonomous Database instance. If you are using an instance principal or resource principal, you must map a dynamic group. Thus, you cannot exclusively map instance and resource principals; you only can map them through a shared mapping and putting the instance or resource instance in an IAM dynamic group.

  You can create Dynamic Groups and reference dynamic groups in the policies you create to access Oracle Cloud Infrastructure. See Configure Policies and Roles to Access Resources and Managing Dynamic Groups for details.

Add IAM Users on Autonomous Database

To add IAM users to allow access to Autonomous Database, map database global users to IAM groups or users with `CREATE USER` or `ALTER USER` statements (with `IDENTIFIED GLOBALLY AS` clause).

The authorization of IAM users to an Autonomous Database instance works by mapping IAM global users (schemas) to IAM users (exclusive mapping) or IAM groups (shared schema mapping).

To authorize IAM users on an Autonomous Database instance:

1. Log in as the ADMIN user to the database that is enabled to use IAM (the ADMIN user has the required `CREATE USER` and `ALTER USER` system privileges that you need for these steps).
2. Create a mapping between the Autonomous Database user (schema) with `CREATE USER` or `ALTER USER` statements and include the `IDENTIFIED GLOBALLY AS` clause, specifying the IAM group name.

   Use the following syntax to map a global user to an IAM group:

   ```
   CREATE USER global_user IDENTIFIED GLOBALLY AS
   'IAM_GROUP_NAME=IAM_GROUP_NAME';
   ```

   For example, to map an IAM group named `db_sales_group` to a shared database global user named `sales_group`:

   ```
   CREATE USER sales_group IDENTIFIED GLOBALLY AS
        'IAM_GROUP_NAME=db_sales_group';
   ```

   This creates a shared global user mapping. The mapping, with global user `sales_group`, is effective for all users in the IAM group. Thus, anyone in the `db_sales_group` can log in to the database using their IAM credentials (through the shared mapping of the `sales_group` global user).

   The following example shows how to accomplish this for a non-default domain:

   ```
   CREATE USER shared_sales_schema IDENTIFIED GLOBALLY AS
        'IAM_GROUP_NAME=sales_domain/db_sales_group';
   ```

3. If you want to create additional global user mappings for other IAM groups or users, follow these steps for each IAM group or user.

> **Note:**
>
> Database users that are not `IDENTIFIED GLOBALLY` can continue to login as before, even when the Autonomous Database is enabled for IAM authentication.

To exclusively map a local IAM user to an Oracle Database Global User:

1. Log in as the ADMIN user to the database that is enabled to use IAM (the ADMIN user has the required `CREATE USER` and `ALTER USER` system privileges that you need for these steps).
2. Create a mapping between the Autonomous Database user (schema) with `CREATE USER` or `ALTER USER` statements and include the `IDENTIFIED GLOBALLY AS` clause, specifying the IAM local IAM user name.

   For example, to create a new database global user named `peter_fitch` and map this user to an existing local IAM user named `peterfitch`:

   ```
   CREATE USER peter_fitch IDENTIFIED GLOBALLY AS
   'IAM_PRINCIPAL_NAME=peterfitch'
   ```

   The following example shows how to create the user by specifying a non-default domain, `sales_domain`:

   ```
   CREATE USER peter_fitch2 IDENTIFIED GLOBALLY AS
   'IAM_PRINCIPAL_NAME=sales_domain/peterfitch';
   ```

Add IAM Roles on Autonomous Database

Optionally, create global roles to provide additional database roles and privileges to IAM users when multiple IAM users are mapped to the same shared global user.

The use of global roles is optional when using either an exclusive IAM mapping to user (schema) or shared user mapping on Autonomous Database. For example, all privileges and roles can be granted to the shared schema and all IAM users who map to the shared schema would be granted the privileges and roles assigned to the shared schema.

You can use a global role to optionally differentiate users who use the same shared schema. For example, a set of users can all have the same shared schema and the shared schema could have the `CREATE SESSION` privilege. Then global roles can be used to provide differentiated privileges and roles assigned to different groups of users who all use the same shared schema.

Granting additional roles to IAM users in Autonomous Database works by mapping Autonomous Database global roles to IAM groups.

To map Autonomous Database global roles to IAM groups:

1. Log in as the ADMIN user to the database that is enabled to use IAM (the ADMIN user has the required `CREATE USER` and `ALTER USER` system privileges that you need for these steps).
2. Set database authorization for Autonomous Database roles with `CREATE ROLE` or `ALTER ROLE` statements and include the `IDENTIFIED GLOBALLY AS` clause, specifying the IAM group name.

Use the following syntax to map a global role to an IAM group:

```
CREATE ROLE global_role IDENTIFIED GLOBALLY AS
    'IAM_GROUP_NAME=IAM_GROUP_of_WHICH_the_IAM_USER_IS_a_MEMBER';
```

For example, to map an IAM group named `ExporterGroup` to a shared database global role named `export_role`:

```
CREATE ROLE export_role IDENTIFIED GLOBALLY AS
     'IAM_GROUP_NAME=ExporterGroup';
```

The following example shows how to create the role by specifying a non-default domain, `sales_domain`:

```
CREATE ROLE export_role IDENTIFIED GLOBALLY AS
     'IAM_GROUP_NAME=sales_domain/ExporterGroup';
```

All members of the `ExporterGroup` in the `sales_domain` domain will be authorized with the database global role `export_role` when they log in to the database.

3. Use `GRANT` statements to grant the required privileges or other roles to the global role.

```
GRANT CREATE SESSION TO export_role;
GRANT DWROLE TO export_role;
```

4. If you want an existing database role to be associated with an IAM group, then use `ALTER ROLE` statement to alter the existing database role to map the role to an IAM group. Use the following syntax to alter an existing database role to map it to an IAM group:

```
ALTER ROLE existing_database_role
   IDENTIFIED GLOBALLY AS 'IAM_GROUP_NAME=IAM_Group_Name';
```

If you want to add additional global role mappings for other IAM groups, follow these steps for each IAM group.

Create IAM Database Password for IAM Users

To add an IAM user and allow the IAM user to login to Autonomous Database by supplying a username and password, you must create an IAM database password.

See Working with IAM Database Passwords for more information.

Connect to Autonomous Database with Identity and Access Management (IAM) Authentication

After the ADMIN user enables Oracle Cloud Infrastructure IAM on Autonomous Database, users log in to the Autonomous Database instance using their Oracle Cloud Infrastructure IAM credentials or access the database through an Oracle Cloud Infrastructure IAM database token.

After you enable Oracle Cloud Infrastructure IAM user access, you can also log in to the Autonomous Database using your local database account username and password (non-global database user account).

You can use a database client to access an Autonomous Database instance as an Oracle Cloud Infrastructure IAM user. To use a client with Oracle Cloud Infrastructure IAM username and password credentials and a password verifier, the database client must be 12c or newer.

Alternatively, you can use an Oracle Cloud Infrastructure IAM database token to access an Autonomous Database instance. Using IAM tokens requires the latest Oracle Database client 19c (at least 19.16). Some earlier clients provide a limited set of capabilities for token access.

The following examples show password verifier with SQL*Plus to access the database with an Oracle Cloud Infrastructure IAM username and password and the steps required to use SQL*Plus with an Oracle Cloud Infrastructure IAM database token.

> **Note:**
>
> If your Autonomous Database instance is in Restricted Mode, only the users with the `RESTRICTED SESSION` privilege such as ADMIN can connect to the database.

You can use an Oracle Cloud Infrastructure IAM database token to access an Autonomous Database instance with supported clients, including the following:

- JDBC-Thin with support for IAM Token Authentication is supported with the following:
  - JDBC version 19.13.0.0.1 (or later): See JDBC and UCP Downloads for JDBC drivers.
  - JDBC version 21.4.0.0.1 (or later): See JDBC and UCP Downloads for JDBC drivers.

  See Support for IAM Token-Based Authentication for more information:
- SQL*Plus and Oracle Instant Client: Supported with SQL*Plus and Instant Client on Linux versions 19.13 or later, and Instant Client on Linux versions 21.4 or later.
  See Identity and Access Management (IAM) Token-Based Authentication for more information.
- The database client can also be configured to retrieve a database token using the IAM username and IAM database password.

  See Client Connections That Use a Token Requested by an IAM User Name and Database Password for more information.
- .NET clients (latest version of Linux or Windows). .NET software components are available as a free download from the following sites:
  - Oracle Data Access Components - .NET Downloads
  - NuGet Gallery
  - Visual Studio Code Marketplace
- About Connecting to an Autonomous Database Instance Using IAM
  IAM users can connect to the Autonomous Database instance by using either an IAM database password verifier or an IAM token.
- Configuring a Client Connection for SQL*Plus That Uses an IAM Database Password
  You can configure SQL*Plus to use an IAM database password.
- Configuring a Client Connection for SQL*Plus That Uses an IAM Token
  You can configure a client connection for SQL*Plus that uses an IAM token.

- Use Instance Principal to Access Autonomous Database with Identity and Access Management (IAM) Authentication
  After the ADMIN user enables Oracle Cloud Infrastructure IAM on Autonomous Database, an application can access the database through an Oracle Cloud Infrastructure IAM database token using an instance principal.

About Connecting to an Autonomous Database Instance Using IAM

IAM users can connect to the Autonomous Database instance by using either an IAM database password verifier or an IAM token.

Using the IAM database password verifier is similar to the Oracle Database password authentication process. However, instead of the password verifier (encrypted hash of the password) being stored in the Oracle database, the verifier is instead stored as part of the Oracle Cloud Infrastructure (OCI) IAM user profile.

The second connection method, the use of an IAM token for the database, is more modern. The use of token-based access is a better fit for Cloud resources such as Autonomous Database. The token is based on the strength that the IAM endpoint can enforce. This can be multi-factor authentication, which is stronger than the use of passwords alone. Another benefit of using tokens is that the password verifier (which is considered sensitive) is never stored or available in memory. A TCPS (TLS) connection is required when using tokens for database access.

> **Note:**
>
> You cannot configure native network encryption when passing an IAM token. Only Transport Layer Security (TLS) by itself is supported, not native network encryption or native network encryption with TLS.

- Client Connections That Use an IAM Database Password Verifier
  After you have configured the authorization needed for the IAM user, this user can log in using existing client application, such as SQL*Plus or SQLcl without additional configuration.
- Client Connections That Use a Token Requested by a Client Application or Tool
  For IAM token access to the Autonomous Database, the client application or tool requests a database token from IAM for the IAM user.

Client Connections That Use an IAM Database Password Verifier

After you have configured the authorization needed for the IAM user, this user can log in using existing client application, such as SQL*Plus or SQLcl without additional configuration.

The IAM user enters the IAM user name and IAM database password (not the Oracle Cloud Infrastructure (OCI) console password) using any currently supported database client. The only constraint is that the database client version be either Oracle Database release 12.1.0.2 or later to use Oracle Database 12c passwords. The database client must be able to use the `12C` password verifier. Using the `11G` verifier encryption is not supported with IAM. No special client or tool configuration is needed for the IAM user to connect to the OCI DBaaS instance.

Client Connections That Use a Token Requested by a Client Application or Tool

For IAM token access to the Autonomous Database, the client application or tool requests a database token from IAM for the IAM user.

The client application will pass the database token directly to the database client through the database client API.

If the application or tool has not been updated to request an IAM token, then the IAM user can use Oracle Cloud Infrastructure (OCI) command line interface (CLI) to request and store the database token. You can request a database access token (`db-token`) using the following credentials:

- Security tokens (with IAM authentication), delegation tokens (in the OCI cloud shell) and `API-keys`, which are credentials that represent the IAM user to enable the authentication
- Instance principal tokens, which enable instances to be authorized actors (or principals) to perform actions on service resources after authenticating
- Resource principal token, which is a credential that enables the application to authenticate itself to other Oracle Cloud Infrastructure services
- Using an IAM user name and IAM database password (can only be requested by database client).

When the IAM users logs into the client with a slash `/` login and the `OCI_IAM` parameter is configured (`sqlnet.ora`, `tnsnames.ora`, or as part of a connect string), then the database client retrieves the database token from a file. If the IAM user submits a user name and password, the connection will use the IAM database verifier access described for client connections that use IAM database password verifiers. The instructions in this guide show how to use the OCI CLI as a helper for the database token. If the application or tool has been updated to work with IAM, then follow the instructions for the application or tool. Some common use cases include the following: SQLPlus on-premises, SQLcl on-premises, SQL*Plus in Cloud Shell, or applications that use SEP wallets.

Configuring a Client Connection for SQL*Plus That Uses an IAM Database Password

You can configure SQL*Plus to use an IAM database password.

- As the IAM user, log in to the Autonomous Database instance by using the following syntax:

  ```
  CONNECT user_name@db_connect_string
  Enter password: password
  ```

  In this specification, *user_name* is the IAM user name. There is a limit of 128 bytes for the combined *domain_name*/*user_name*.

  The following example shows how IAM user `peter_fitch` can log in to an Autonomous Database instance.

  ```
  sqlplus /nolog
  connect peter_fitch@db_connect_string
  Enter password: password
  ```

  Some special characters will require double quotation marks around *user_name* and *password*. For example:

  ```
  "peter_fitch@example.com"@db_connect_string

  "IAM database password"
  ```

Configuring a Client Connection for SQL*Plus That Uses an IAM Token

You can configure a client connection for SQL*Plus that uses an IAM token.

1. Ensure you have an IAM user account.
2. Check with an IAM administrator and an Oracle Database administrator to ensure you have a policy allowing you to access the database in the compartment or your tenancy and that you are mapped to a global schema in the database.
3. If your application or tool does not support direct IAM integration, then download, install, and configure the OCI CLI. (See *OCI Command Line Interface Quickstart*.) Set up an API key as part of the OCI CLI configuration and select default values.
   a. Set up the API key access for the IAM user.
   b. Retrieve the `db-token`. For example:
      - Retrieving a `db-token` with an `API-key` using the Oracle Cloud Infrastructure (OCI) command-line interface:

        ```
        oci iam db-token get
        ```

      - Retrieving a `db-token` with a security (or session) token:

        ```
        oci iam db-token get --auth security_token
        ```

        If the security token has expired, a window will appear so the user can log in to OCI again. This generates the security token for the user. OCI CLI will use this refreshed token to get the `db-token`.
      - Retrieving a `db-token` with a delegation token: When you log in to the cloud shell, the delegation token is automatically generated and placed in the `/etc` directory. To get this token, run the following command in the cloud shell:

        ```
        oci iam db-token get
        ```

      - Retrieving an instance token by using the OCI command-line interface:

        ```
        oci iam db-token get --auth instance_principal
        ```

   c. The database client can also be configured to retrieve a database token using the IAM username and IAM database password.

      See Client Connections That Use a Token Requested by an IAM User Name and Database Password for more information.

   See Required Keys and OCIDs for more information.
4. Ensure that you are using the latest release updates for the Oracle Database client releases 19c, 21c, or 23ai.

   This configuration only works with the Oracle Database client release 19c, 21c, or 23ai.
5. Follow the existing process to download the wallet from the Autonomous Database and then follow the directions for configuring it for use with SQL*Plus.
   a. Confirm that DN matching is enabled by looking for `SSL_SERVER_DN_MATCH=ON` in `sqlnet.ora`.

> **Note:**
>
> Partial or full DN matching is required when sending a token from the database client to Autonomous Database. If Autonomous Database is using a private endpoint, you need to specify a host value for the connect string parameter. Using an IP address for the host parameter in the connect string will not work with DN matching and the IAM token will not be sent to the database.
> See Private Endpoints Configuration Examples on Autonomous Database for configuration information on how to set the host parameter when using a private endpoint.

b. Configure the database client to use the IAM token by adding `TOKEN_AUTH=OCI_TOKEN` to the `sqlnet.ora` file. Because you will be using the default locations for the database token file, you do not need to include the token location.

The `TOKEN_AUTH` and `TOKEN_LOCATION` values in the `tnsnames.ora` connect strings take precedence over the `sqlnet.ora` settings for that connection. For example, for the connect string, assuming that the token is in the default location (`~/.oci/db-token` for Linux):

```
(description=
  (retry_count=20)(retry_delay=3)
  (address=(protocol=tcps)(port=1522)
  (host=example.us-phoenix-1.oraclecloud.com))

(connect_data=(service_name=aaabbbccc_exampledb_high.example.oraclecloud.co
m))
  (security=(ssl_server_dn_match=yes)
  (TOKEN_AUTH=OCI_TOKEN)))
```

After the connect string is updated with the `TOKEN_AUTH` parameter, the IAM user can log in to the Autonomous Database instance by running the following command to start SQL*Plus. You can include the connect descriptor itself or use the name of the descriptor from the `tnsnames.ora` file.

```
connect /@exampledb_high
```

Or:

```
connect /@(description=
  (retry_count=20)(retry_delay=3)
  (address=(protocol=tcps)(port=1522)
  (host=example.us-phoenix-1.oraclecloud.com))

(connect_data=(service_name=aaabbbccc_exampledb_high.example.oraclecloud.com))
  (security=(ssl_server_cert_dn="CN=example.uscom-east-1.oraclecloud.com,
     O=Example Corporation,
     L=Redwood City, ST=California, C=US")
  (TOKEN_AUTH=OCI_TOKEN)))
```

The database client is already configured to get a `db-token` because `TOKEN_AUTH` has already been set, either through the `sqlnet.ora` file or in a connect string. The database client gets the `db-token` and signs it using the private key and then sends the token to the Autonomous

Database. If an IAM user name and IAM database password are specified instead of slash `/`, then the database client will connect using the password instead of using the `db-token`.

Use Instance Principal to Access Autonomous Database with Identity and Access Management (IAM) Authentication

After the ADMIN user enables Oracle Cloud Infrastructure IAM on Autonomous Database, an application can access the database through an Oracle Cloud Infrastructure IAM database token using an instance principal.

See Accessing the Oracle Cloud Infrastructure API Using Instance Principals for more information.

Configure IAM Proxy Authentication

Proxy authentication allows an IAM user to proxy to a database schema for tasks such as application maintenance.

- About Configuring IAM Proxy Authentication
  IAM users can connect to Oracle DBaaS by using proxy authentication.
- Configure Proxy Authentication for the IAM User
  To configure proxy authentication for an IAM user, the IAM user must already have a mapping to a global schema (exclusive or shared mapping). A separate database schema for the IAM user to proxy to must also be available.
- Validate the IAM User Proxy Authentication
  You can validate the IAM user proxy configuration for both password and token authentication methods.

About Configuring IAM Proxy Authentication

IAM users can connect to Oracle DBaaS by using proxy authentication.

Proxy authentication is typically used to authenticate the real user and then authorize them to use a database schema with the schema privileges and roles in order to manage an application. Alternatives such as sharing the application schema password are considered insecure and unable to audit which actual user performed an action.

A use case can be in an environment in which a named IAM user who is an application database administrator can authenticate by using their credentials and then proxy to a database schema user (for example, `hrapp`). This authentication enables the IAM administrator to use the `hrapp` privileges and roles as user `hrapp` in order to perform application maintenance, yet still use their IAM credentials for authentication. An application database administrator can sign in to the database and then proxy to an application schema to manage this schema.

You can configure proxy authentication for both the password authentication and token authentication methods.

Configure Proxy Authentication for the IAM User

To configure proxy authentication for an IAM user, the IAM user must already have a mapping to a global schema (exclusive or shared mapping). A separate database schema for the IAM user to proxy to must also be available.

After you ensure that you have this type of user, alter the database user to allow the IAM user to proxy to it.

1. Log in to the Autonomous Database instance as a user who has the `ALTER USER` system privileges.
2. Grant permission for the IAM user to proxy to the local database user account.

   An IAM user cannot be referenced in the command so the proxy must be created between the database global user (mapped to the IAM user) and the target database user.

   In the following example, `hrapp` is the database schema to proxy to, and `peterfitch_schema` is the database global user exclusively mapped to user `peterfitch`.

   ```
   ALTER USER hrapp GRANT CONNECT THROUGH peterfitch_schema;
   ```

At this stage, the IAM user can log in to the database instance using the proxy. For example, to connect using a password verifier:

```
CONNECT peterfitch[hrapp]@connect_string
Enter password: password
```

To connect using a token:

```
CONNECT [hrapp]/@connect_string
```

Validate the IAM User Proxy Authentication

You can validate the IAM user proxy configuration for both password and token authentication methods.

1. Log in to the Autonomous Database instance as a user who has the `CREATE USER` and `ALTER USER` system privileges.
2. Connect as the IAM user and run the `SHOW USER` and `SELECT SYS_CONTEXT` commands.

   For example, suppose you want to check the proxy authentication of the IAM user `peterfitch` when they proxy to database user `hrapp`. Run the following queries after you proxy to the database using an IAM user. Depending on how you authenticate and access the database, you will get different values for these queries.

   - For password authentication, assuming the IAM user is in the default domain:

     ```
     CONNECT peterfitch[hrapp]/password\!@connect_string
     SHOW USER;
     --The output should be USER is "HRAPP"
     SELECT SYS_CONTEXT('USERENV','AUTHENTICATION_METHOD') FROM DUAL;
     --The output should be "PASSWORD_GLOBAL_PROXY"
     SELECT SYS_CONTEXT('USERENV','PROXY_USER') FROM DUAL;
     --The output should be "PETERFITCH_SCHEMA"
     ```

```
SELECT SYS_CONTEXT('USERENV','CURRENT_USER') FROM DUAL;
--The output should be "HRAPP"
```

- For token authentication, for a user who is in a non-default domain, `sales_domain`:

```
CONNECT [hrapp]/@connect_string
SHOW USER;
--The output should be USER is "HRAPP"
SELECT SYS_CONTEXT('USERENV','AUTHENTICATION_METHOD') FROM DUAL;
--The output should be "TOKEN_GLOBAL_PROXY"
SELECT SYS_CONTEXT('USERENV','PROXY_USER') FROM DUAL;
--The output should be "PETERFITCH_SCHEMA"
SELECT SYS_CONTEXT('USERENV','CURRENT_USER') FROM DUAL;
--The output should be "HRAPP"
```

Disable Identity and Access Management (IAM) Authentication on Autonomous Database

Describes the steps to disable IAM external authentication user access for Autonomous Database.

You can disable IAM user access on your Autonomous Database instance as follows:

1. Run the `DBMS_CLOUD_ADMIN.DISABLE_EXTERNAL_AUTHENTICATION` procedure.

   For example:

   ```
   BEGIN
      DBMS_CLOUD_ADMIN.DISABLE_EXTERNAL_AUTHENTICATION;
   END;
   /
   ```

2. If you also want to update access to IAM from the resource, in this case the Autonomous Database instance, you may need to remove or modify the IAM group and the policies you set up to allow access to IAM from the Autonomous Database instance.

Accessing a Database Cross-Tenancy Using an IAM Integration

Users and groups in one tenancy can access DBaaS database instances in another tenancy if policies in both tenancies allow this.

- About Cross-Tenancy Access for IAM Users to DBaaS Instances
  Cross-tenancy access to an Oracle Cloud Infrastructure (OCI) DBaaS instance is similar to a single tenancy scenario except that tenancy information is required for mappings and token requests and a policy is required in both tenancies to allow this cross tenancy database resource access.
- Configuring Policies
  You must create policies in both the user tenancy and the database resource tenancy to allow cross-tenancy database access.
- Mapping Database Schemas and Roles to Users and Groups in Another Tenancy
  When you perform this type of mapping, you must add the tenancy OCID to the mapping information so the database knows it is cross-tenancy access.

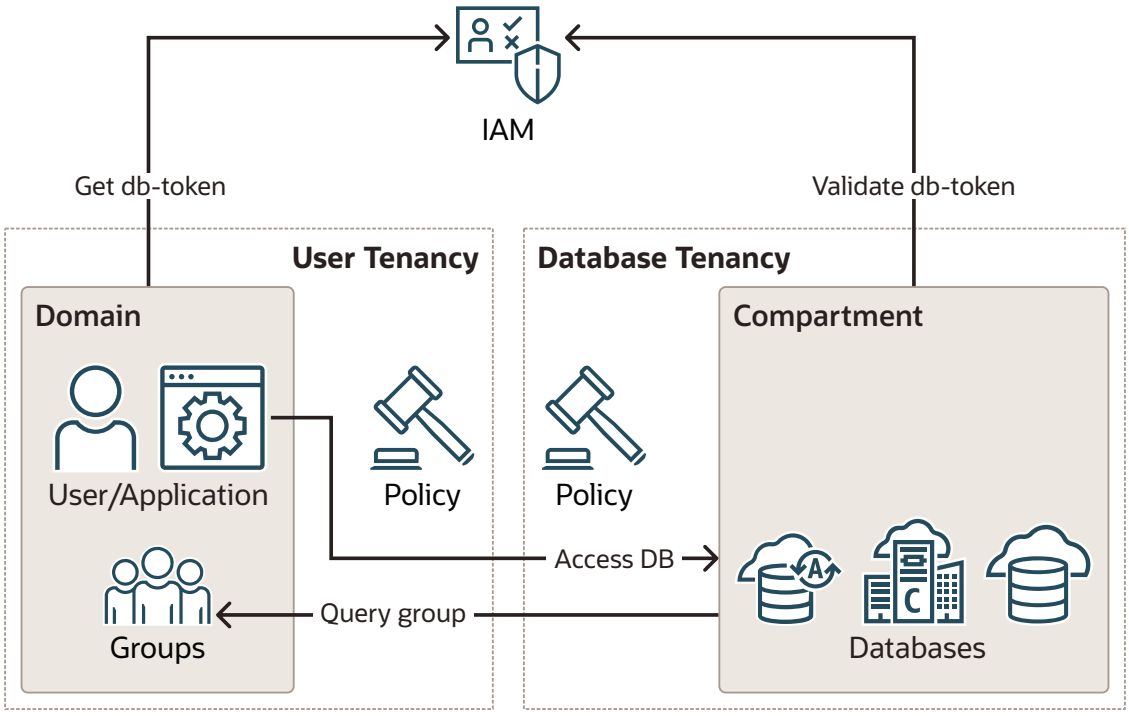
- Configuring Database Clients for Cross-Tenancy Access
  You can configure some database clients directly.
- Requesting Cross-Tenancy Tokens Using the OCI Command-Line Interface
  You must add the `--scope` parameter to the Oracle Cloud Infrastructure (OCI) command-line interface command to get a `db-token` for a cross-tenancy request. If the database you are accessing is in a different region than the user tenancy home region, then the region must also be added to the OCI CLI command using the --region parameter.

About Cross-Tenancy Access for IAM Users to DBaaS Instances

Cross-tenancy access to an Oracle Cloud Infrastructure (OCI) DBaaS instance is similar to a single tenancy scenario except that tenancy information is required for mappings and token requests and a policy is required in both tenancies to allow this cross tenancy database resource access.

The following figure illustrates the process for a cross-tenancy access to an OCI DBaaS instance.

Figure 3-1 Cross-Tenancy Access to an OCI DBaaS Instance



The cross-tenancy process is as follows:

1. The policy is required in both tenancies to endorse and admit access cross tenancy.
2. The IAM principal (user or application) requests a db-token for a cross-tenancy resource.
3. The `db-token` is returned and is used to access the database in a different tenancy
4. The database will make a cross-tenancy group query for the user's groups and map principal to global schema and optional global roles.

You must subscribe the user tenancy to the same regions in which the databases are located. For example, if the databases in the database tenancy are in the `PHX` and `IAD` regions, then

you must subscribe the user tenancy to these regions. This is not the home region, just the additional subscribed regions in the user tenancy.

Configuring Policies

You must create policies in both the user tenancy and the database resource tenancy to allow cross-tenancy database access.

- Configuring the Source User Tenancy
  Two policies are required to allow cross-tenancy access in the user tenancy.
- Configuring the Target Database Resource Tenancy
  The database tenancy will need matching policies to enable access to the users from the user tenancy as well as allow its own databases to query group information in the user tenancy
- Policy Examples for Cross-Tenancy Access
  Examples include using a `WHERE` clause to refine the cross-tenancy configuration, and other methods of performing this type of configuration.

Configuring the Source User Tenancy

Two policies are required to allow cross-tenancy access in the user tenancy.

The first policy is to allow a user tenancy group to access a database in a different tenancy. The second policy allows a database in the database tenancy to query group information in the user tenancy.

1. In the OCI console, select **Identity & Security**.
2. Under **Identity**, select **Policies**.
3. Click **Create Policy** and in the Policy Builder select **Show manual editor**.
4. Use the `DEFINE` statement to make it easier to read the actual policies.

   For example:

   ```
   DEFINE tenancy database_tenancy as ocid1.tenancy.OCID
   ```

5. Endorse the tenancy group `domainA/xt_db_users` to use database_connections in tenancy `database_tenancy`.

   This allows users of the `xt_db_users` group in `domainA` to access any database in tenancy `database_tenancy`.

   ```
   ENDORSE group domainA/xt_db_users to use database-connections in tenancy
   database_tenancy
   ```

6. Use the `ADMIT` statement to create an Admit policy to allow any database in the database tenancy to query group information for specific IAM users in the user tenancy.

   ```
   ADMIT any-user of tenancy database_tenancy to {GROUP_MEMBERSHIP_INSPECT,
   AUTHENTICATION_INSPECT} in tenancy
   ```

Configuring the Target Database Resource Tenancy

The database tenancy will need matching policies to enable access to the users from the user tenancy as well as allow its own databases to query group information in the user tenancy

1. In the OCI console, select **Identity & Security**.
2. Under **Identity**, select **Policies**.
3. Click **Create Policy** and in the Policy Builder select **Show manual editor**.
4. Use `DEFINE` to make it easier to troubleshoot and read the policies.

   ```
   DEFINE tenancy user_tenancy as ocid1.tenancy.OCID
   DEFINE group xt_db_users as ocid1.group.defg
   ```

5. Use `ADMIT` to create an Admit policy in the tenancy to match the Endorse policy from the user tenancy.

   The Admit policy must match the `ENDORSE` policy in the user tenancy so that it can enable users from the `user_tenancy` to access databases in this tenancy.

   ```
   ADMIT group xt_db_users of tenancy user_tenancy to use database-
   connections in tenancy
   ```

6. Create an Endorse policy, which will match the Admit policy created in the User tenancy.

   The Endorse policy will enable databases in the database tenancy to query group information from the user_tenancy.

   ```
   ENDORSE any-user to {GROUP_MEMBERSHIP_INSPECT, AUTHENTICATION_INSPECT} in
   tenancy user_tenancy
   ```

While using `any-user` makes it easy to understand the required policies, Oracle recommends that you use stronger constraints in addition to or instead of using `any-user`. The `any-user` option will allow any principal or resource to query user groups in the `user_tenancy`. Ideally, you should limit this to just allowing the database resources (resource principals) to make the group queries. You can do this by adding a `WHERE` clause to the policies or by adding a dynamic group that limits it to the members of the dynamic group. Defining every possible way to specify dynamic groups and policies is outside the scope of this topic. You can find more information from these sources:

- Managing Dynamic Groups
- Managing Policies

Policy Examples for Cross-Tenancy Access

Examples include using a `WHERE` clause to refine the cross-tenancy configuration, and other methods of performing this type of configuration.

You can add a `WHERE` clause to limit the database resources allowed to make the cross-tenancy group query:

```
ADMIT any-user of tenancy db_tenancy to {GROUP_MEMBERSHIP_INSPECT,
AUTHENTICATION_INSPECT} in tenancy where request.principal.type = 'dbsystem'
```

This Admit policy allows any Base Database Service (resource type: `dbsystem`) in the `db_tenancy` to query a user's group information from the user tenancy. Resource type names are in the table below.

A similar method can be done by putting the same resource type into a dynamic group:

```
dynamic group: db_principals
any {resource.type = 'dbsystem', resource.type = 'vmcluster', resource.type =
'cloudvmcluster'}
```

The dynamic group in the preceding example includes database instances for Oracle Base Database Service (`dbsystem`), Oracle Exadata Cloud@Customer (`vmcluster`), and Oracle Exadata Database Service (`cloudvmcluster`).

This example uses a dynamic group instead of `any-user`:

```
ADMIT dynamic group db_principals of tenancy db_tenancy to
{GROUP_MEMBERSHIP_INSPECT, AUTHENTICATION_INSPECT} in tenancy
```

You can also add all resource principals in a compartment using `resource.compartment.id`. However, this might also allow other non-database resource principals to make the cross-tenancy group query. The following table provides a mapping of the various resource types with the DBaaS platform name:

DBaaS Platform Name	Resource Type Name
ADB-S	`autonomousdatabase`
ADB-D (OPC)	`cloudautonomousvmcluster*`
Base DBS	`dbsystem`
ExaCS	`cloudvmcluster`
ExaCC	`vmcluster`

* Older ADBD instances may still be using the `autonomousexainfrastructure` resource type.

Mapping Database Schemas and Roles to Users and Groups in Another Tenancy

When you perform this type of mapping, you must add the tenancy OCID to the mapping information so the database knows it is cross-tenancy access.

Use a full colon to separate the tenancy OCID when you use the `CREATE USER` and `CREATE ROLE` statements in SQL*Plus.

- To use the `CREATE USER` statement to perform the mapping:

  The following examples show exclusive and shared schema mapping with principals and groups in default and non-default domains. When using default domains, you do not need to include a domain name.

  ```
  CREATE USER schema1 IDENTIFIED GLOBALLY
  AS 'IAM_PRINCIPAL_NAME=ocid1.tenancy.OCID:example_domain/
  peter.fitch@oracle.com';

  CREATE USER schema2 IDENTIFIED GLOBALLY
  AS 'IAM_PRINCIPAL_NAME=ocid1.tenancy.OCID:peter.fitch@oracle.com';

  CREATE USER qa_db_user_group IDENTIFIED GLOBALLY
  AS 'IAM_GROUP_NAME=ocid1.tenancy.OCID:example_domain/xt_db_users';
  ```

```
CREATE USER qa_sales_user_group IDENTIFIED GLOBALLY
AS 'IAM_GROUP_NAME=ocid1.tenancy.OCID:sales_users';

CREATE USER xt_ip_user IDENTIFIED GLOBALLY
AS 'IAM_PRINCIPAL_OCID=ocid1.instance.region1.sea.OCID';
GRANT CREATE SESSION TO xt_ip_user;

CREATE USER xt_iam_dg IDENTIFIED GLOBALLY
AS 'IAM_GROUP_NAME=ocid1.tenancy.region1.OCID:sales_principals';
GRANT CREATE SESSION TO xt_iam_dg;
```

- To use the `CREATE ROLE` statement to perform the mapping:

  The following examples show global role mapping with groups in default and non-default domains. When using default domains, you do not need to include a domain name.

```
CREATE ROLE globalrole1 IDENTIFIED GLOBALLY
AS 'IAM_GROUP_NAME=ocid1.tenancy.abcdef:example_domain/xt_db_users';

CREATE ROLE globalrole2 IDENTIFIED GLOBALLY
AS 'IAM_GROUP_NAME=ocid1.tenancy.abcdef:sales_users';
```

Configuring Database Clients for Cross-Tenancy Access

You can configure some database clients directly.

The database tenancy must be identified in either the connect string or in `sqlnet.ora` if the client is configured to directly get the access token from OCI IAM. Review client-specific documentation for specific parameter values (JDBC-thin, ODP.NET-core, managed).

Requesting Cross-Tenancy Tokens Using the OCI Command-Line Interface

You must add the `--scope` parameter to the Oracle Cloud Infrastructure (OCI) command-line interface command to get a `db-token` for a cross-tenancy request. If the database you are accessing is in a different region than the user tenancy home region, then the region must also be added to the OCI CLI command using the --region parameter.

See Optional Parameters for more details about using the optional parameters of the `oci get` command.

You can scope it for the entire tenancy or scope it to a compartment or database in the tenancy. When scoping for cross tenancy compartment or database, you do not need to also add the tenancy information because the compartment and database OCIDs are unique across OCI.

Certain clients can request the tokens directly from MSEI. Refer to their documentation on setting the parameters to get the MSEI `OAuth2` access tokens.

Notes for Using Autonomous Database Tools with Identity and Access Management (IAM) Authentication

Provides notes for using Autonomous Database tools with IAM authentication enabled.

- Oracle APEX is not supported for IAM users with Autonomous Database. See Create Oracle APEX Workspaces in Autonomous Database for information on using regular database users with Autonomous Database.

- Database Actions is not supported for IAM users with Autonomous Database. See Provide Database Actions Access to Database Users for information on using regular database users with Autonomous Database.
- Oracle Machine Learning Notebooks and other components are not supported for IAM Authorized users with Autonomous Database. See Add Existing Database User Account to Oracle Machine Learning Components for information on using regular database users with Autonomous Database.

Use Microsoft Entra ID with Autonomous Database

You can configure an Autonomous Database instance for Microsoft Entra ID users to connect using Entra ID (Azure AD) `OAuth2` tokens.

> **Note:**
>
> Microsoft recently changed the name of Microsoft Azure AD to Microsoft Entra ID.

- About Integrating Oracle Autonomous Database with Microsoft Entra ID
  Oracle Autonomous Database and Microsoft Entra ID can be configured to allow users and applications to connect to the database using their Entra ID credentials.
- Enable Microsoft Entra ID Authentication on Autonomous Database
  A Microsoft Entra ID administrator and an Autonomous Database administrator perform steps to configure Entra ID authentication on Autonomous Database.
- Role and Schema Mapping for Entra ID Authentication on Autonomous Database
  Microsoft Entra ID (Azure AD) users are mapped to one database schema and optionally to one or more database roles. After mapping Entra ID users, user can connect to the Autonomous Database instance.
- Microsoft Entra ID Client Configuration and Access for Autonomous Database
  After you configure Entra ID on your Autonomous Database and you map Entra ID users, there are numerous ways that a user can configure a client to the Autonomous Database instance using Entra ID tokens.
- Testing the Accessibility of the Entra ID Endpoint
  You must ensure that your Oracle Database can access the Entra ID endpoint.

About Integrating Oracle Autonomous Database with Microsoft Entra ID

Oracle Autonomous Database and Microsoft Entra ID can be configured to allow users and applications to connect to the database using their Entra ID credentials.

Azure users and applications can log in with Entra ID Single Sign On (SSO) credentials to access the database. This is done with an Entra ID `OAuth2` access token that the user or application first requests from Entra ID. This `OAuth2` access token contains the user identity and database access information and is then sent to the database. Refer to Refer to the Microsoft article Passwordless authentication options for Azure Active Directory for information about configuring multi-factor and passwordless authentication.

You can perform this integration in the following Oracle Database environments:

- On-premises Oracle Database release 19.18 and later, excluding 21c
- All Oracle Database server platforms: Linux, Windows, AIX, Solaris, and HPUX
- Oracle Autonomous Database Serverless

- Oracle Autonomous Database on Dedicated Exadata Infrastructure
- Oracle Autonomous Database on Exadata Cloud@Customer
- Oracle Exadata Database Service on Dedicated Infrastructure
- Oracle Exadata Database Service on Cloud@Customer
- Oracle Base Database Service

The instructions for configuring Entra ID use the term "Oracle Database" to encompass these environments.

This type of integration enables the Azure user to access an Oracle Autonomous Database instance. Azure users and applications can log in with Entra ID Single Sign On (SSO) credentials to get an Entra ID `OAuth2` access token to send to the database.

The Entra ID administrator creates and registers Oracle Autonomous Database with Entra ID. Within Entra ID, this is called an app registration, which is short for application registration. This is the digital information that Entra ID must know about the software that is using Entra ID. The Entra ID administrator also creates application (app) roles for the database app registration in Entra ID. App roles connect Azure users, groups, and applications to database schemas and roles. The Entra ID administrator assigns Azure users, groups, or applications to the app roles. These app roles are mapped to a database global schema or a global role or to both a schema and a role. An Azure user, group, or application that is assigned to an app role will be mapped to a database global schema, global role, or to both a schema and a role. An Oracle global schema can also be mapped exclusively to an Azure user. An Azure guest user (non-organization user) or an Entra ID service principal (application) can only be mapped to a database global schema through an Entra ID app role. An Oracle global role can only be mapped from an Azure app role and cannot be mapped from an Azure user.

Oracle Autonomous Database tools including Oracle APEX, Database Actions, Oracle Graph Studio, and Oracle Database API for MongoDB are not compatible with using Entra ID tokens to connect with the database.

Tools and applications that are updated to support Entra ID tokens can authenticate users directly with Entra ID and pass the database access token to the Oracle Autonomous Database instance. You can configure existing database tools such as SQL*Plus to use an Entra ID token from a file location. In these cases, Entra ID tokens can be retrieved using tools like Microsoft PowerShell or Azure CLI and put into a file location. An Entra ID `OAuth2` database access tokens are issued with an expiration time. The Oracle Database client driver will ensure that the token is in a valid format and that it has not expired before passing it to the database. The token is scoped for the database, which means that there is information in the token about the database where the token will be used. The app roles the Entra ID principal was assigned to in the database Entra ID app registration are included as part of the access token. The directory location for the Entra ID token should only have enough permission for the user to write the token file to the location and the database client to retrieve these files (for example, just read and write by the user). Because the token allows access to the database, it should be protected within the file system.

Azure users can request a token from Entra ID using a number of methods to open an Azure login window to enter their Entra ID credentials.

Oracle Autonomous Database accepts tokens representing the following Entra ID principals:

- Azure user, who is registered user in the Entra ID tenancy
- Guest user, who is registered as a guest user in the Entra ID tenancy
- Service, which is the registered application connecting to the database as itself with the client credential flow (connection pool use case)

Oracle Autonomous Database supports the following Entra ID authentication flows:

- Interactive flow (also called authorization code flow) using Proof Key for Code Exchange (PKCE), most commonly used for human users (not applications) to authenticate to Entra ID in a client environment with a browser
- Client credentials, which are for database applications that connect as themselves (and not the end-user)
- On-Behalf-Of (OBO), where an application requests an access token on behalf of a logged-in user to send to the database
- Resource owner password credential (ROPC), which is not recommended for production use, but can be used in test environments where a pop-up browser user authentication would be difficult to incorporate. ROPC needs the Entra ID user name and password credential to be part of the token request call.

> **Note:**
>
> The DBaaS integration with Microsoft Entra ID does not support users with administrative privileges (`SYSDBA`, `SYSOPER`, `SYSBACKUP`, `SYSDG`, `SYSKM`, and `SYSRAC`).

Topics

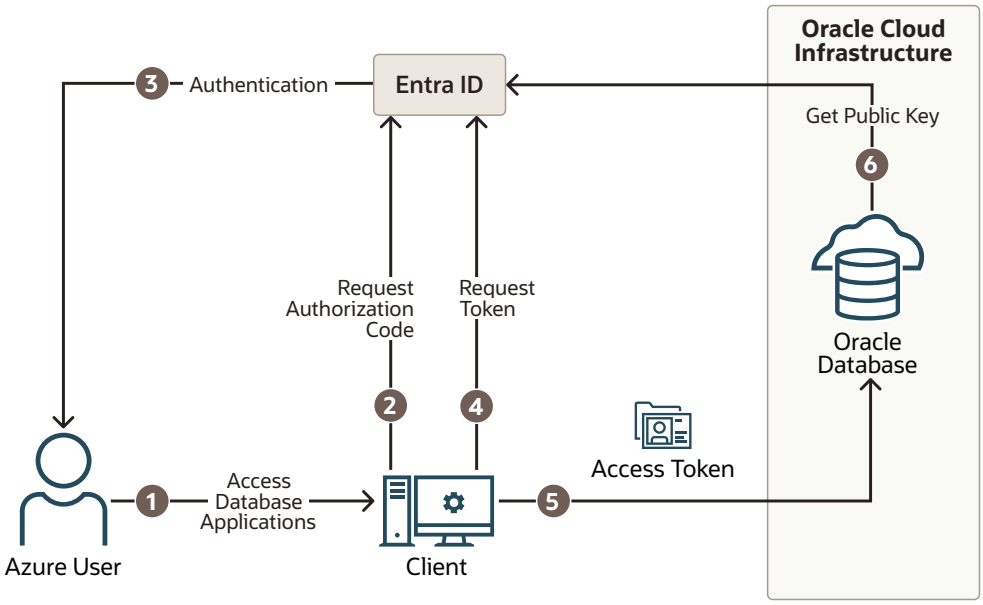
- Architecture of Oracle Database Integration with Microsoft Entra ID
  Microsoft Azure Active Directory access tokens follow the OAuth 2.0 standard with extensions.
- Azure Users Mapping to an Oracle Database Schema and Roles
  Microsoft Azure users must be mapped to an Oracle Database schema and have the necessary privileges (through roles) before being able to authenticate to the Oracle Database instance.
- Use Cases for Connecting to an Oracle Database Using Entra ID
  Oracle Database supports several use cases for connecting to the database.

Architecture of Oracle Database Integration with Microsoft Entra ID

Microsoft Azure Active Directory access tokens follow the OAuth 2.0 standard with extensions.

The Entra ID access token will be needed before you access the database from the database client (for example, with SQLPlus or SQLcl). The Oracle clients (for example, OCI, JDBC, and ODP) can be configured to pick up an Entra ID token from a file location or the token can be passed to the client through the database client API. An Azure user can use a script (examples available from Microsoft) to retrieve a token and put it into a file location for the database client to retrieve. Applications can use the Azure SDK to get an access token and pass the token through the database client API. Command-line tools such as Microsoft PowerShell or the Azure command-line interface can be used to retrieve the Entra ID token if the application cannot directly get the token.

The following diagram is a generalized flow diagram for OAuth 2.0 standard, using the `OAuth2` token. See Authentication flow support in MSAL in the Microsoft Entra ID documentation for more details about each supported flow.



The authorization code flow is an OAuth2 standard and is described in detail as part of the standards. There are two steps in the flow. The first step authenticates the user and retrieves the authorization code. The second step uses the authorization code to get the database access token.

1. The Azure user requests access to the resource, the Oracle Database instance.
2. The database client or application requests an authorization code from Entra ID.
3. Entra ID authenticates the Azure user and returns the authorization code.
4. The helper tool or application uses the authorization code with Entra ID to exchange it for the `OAuth2` token.
5. The database client sends the `OAuth2` access token to the Oracle database. The token includes the database app roles the user was assigned to in the Entra ID app registration for the database.
6. The Oracle Database instance uses the Entra ID public key to verify that the access token was created by Entra ID.

Both the database client and the database server must be registered with the **app registrations** feature in the Azure Active Directory section of the Azure portal. The database client must be registered with Entra ID app registration. Permission must also be granted to allow the database client to get an access token for the database.

Azure Users Mapping to an Oracle Database Schema and Roles

Microsoft Azure users must be mapped to an Oracle Database schema and have the necessary privileges (through roles) before being able to authenticate to the Oracle Database instance.

In Microsoft Azure, an Entra ID administrator can assign users, groups, and applications to the database app roles.

Exclusively mapping an Entra ID user to a database schema requires the database administrator to create and manage a database schema for the lifecycle of the user (joining, moving, leaving). The database administrator must create the schema when the user joins the organization. The database administrator must also modify the privileges and roles that are

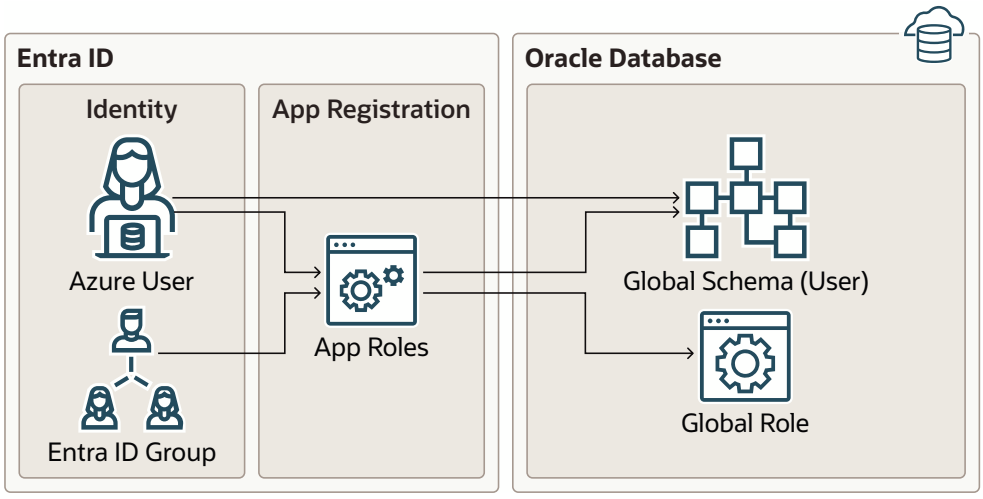
granted to the database schema to align them with the tasks the Azure user is assigned to. When the Azure user leaves the organization, the database administrator must drop the database schema so that an unused account is not left on the database. Using the database app roles enables the Entra ID administrator to control access and roles by assigning users to app roles that are mapped to global schemas and global roles. This way, user access to the database is managed by Entra ID administrators and database administrators do not need to create, manage, and drop schemas for every user.

An Azure user can be mapped to a database schema (user) either exclusively or through an app role.

- **Creating an exclusive mapping between an Azure user and an Oracle Database schema.** In this type of mapping, the database schema must be created for the Azure user. Database privileges and roles that are needed by the Azure user must be granted to the database schema. The database schema not only must be created when the Azure user is authorized to the database, but the granted privileges and roles must be modified as the Entra ID roles and tasks change. Finally, the database schema must be dropped when the Azure user leaves the organization.
- **Creating a shared mapping between an Entra ID app role and an Oracle Database schema.** This type of mapping, which is more common than exclusive mappings, is for Azure users who have been assigned directly to the app role or is a member of an Entra ID group that is assigned to the app role. The app role is mapped to an Oracle Database schema (shared schema mapping). Shared schema mapping allows multiple Azure users to share the same Oracle Database schema so a new database schema is not required to be created every time a new user joins the organization. This operational efficiency allows database administrators to focus on database application maintenance, performance, and tuning tasks instead of configuring new users, updating privileges and roles, and removing accounts.

In addition to database roles and privileges being granted directly to the mapped global schema, additional roles and privileges can be granted through mapped global roles. Different Azure users mapped to the same shared global schema may need different privileges and roles. Azure app roles can be mapped to Oracle Database global roles. Azure users who are assigned to the app role or are a member of an Entra ID group that is assigned to the app role will be granted the Oracle Database global role when they access the database.

The following diagram illustrates the different types of assignments and mappings that are available.



These mappings are as follows:

- An Azure user can be mapped directly to an Oracle Database global schema (user).
- An Azure user, Entra ID group, or application is assigned to an app role, which is then mapped to either an Oracle Database global schema (user) or a global role.

Use Cases for Connecting to an Oracle Database Using Entra ID

Oracle Database supports several use cases for connecting to the database.

- **OAuth2 authorization code flow:** This is the most common flow for human users. The client directs the Azure user to Entra ID to get the authorization code. This code is used to get the database access token. See the Microsoft Azure article Microsoft identity platform and OAuth 2.0 authorization code flow.
- **Resource owner password credentials (ROPC):** This flow is not recommended for production servers. It is useful for test software that cannot work with a pop-up authentication window. It is used in non-graphic user interface environments when a pop-up window cannot be used to authenticate a user.
- **Client credentials:** This flow is used for applications to connect with the database. The application must register with Entra ID app registration and needs a client ID and client password. These client credentials must be used to get the database access token from Entra ID when the application connects to the database. The application can pass the token through the file system or through the database client API.
- **On-behalf-of (OBO) token:** An Azure application requests an OBO token for a logged in user. The OBO token will also be an access token for the database with the Azure user identity and assigned app roles for the database. This enables the Azure user to log in to the database as the user and not the application. Only an application can request an OBO token for its Azure user and pass it to the database client through the API.

Enable Microsoft Entra ID Authentication on Autonomous Database

A Microsoft Entra ID administrator and an Autonomous Database administrator perform steps to configure Entra ID authentication on Autonomous Database.

- Registering the Oracle Database Instance with a Microsoft Entra ID Tenancy
  A user with Entra ID administrator privileges uses Microsoft Entra ID to register the Oracle Database instance with the Microsoft Entra ID tenancy.
- Enabling Microsoft Entra ID v2 Access Tokens
  Oracle Database supports integration with the v1 and v2 Azure AD `OAuth2` access token.
- Managing App Roles in Microsoft Entra ID
  In Entra ID, you can create and manage app roles that will be assigned to Azure users and groups and also be mapped to Oracle Database global schemas and roles.
- Configuring Microsoft Entra ID as an External Identity Provider for Autonomous Database
  An Autonomous Database administrator can enable Entra ID as an external identity provider on an Autonomous Database instance.

Enabling Microsoft Entra ID v2 Access Tokens

Oracle Database supports integration with the v1 and v2 Azure AD `OAuth2` access token.

Oracle Database supports the Entra ID v2 token as well as the default v1 token. However, to use the Entra ID v2 token, you must perform some additional steps to ensure it works with the

Oracle Database. You can use this token with applications that are registered in the Azure portal using the **App registrations** experience.

When you use the Azure AD v2 `OAuth2` access token, the credential flow continues to work as it did before without any changes. However, the `upn:` claim must be added when you use v2 tokens with the interactive flow.

1. Check the version of the Entra ID access token that you are using.
2. Log in to the Microsoft Entra ID portal.
3. Search for and select **Entra ID**.
4. Under **Manage**, select **App registrations**.
5. Choose the application for which you want to configure optional claims based on your scenario and desired outcome.
6. Under **Manage**, select **Token configuration**.
7. Click **Add optional claim** and select **upn**.

When you use v2 tokens, the `aud:` claim only reflects the APP ID value. You do not need to set the `https:domain` prefix to the APP ID URI when v2 tokens are being used. This simplifies the configuration for the database because the default APP ID URI can be used.

- Checking the Entra ID Access Token Version
  You can check the version of the Entra ID access token that your site uses by using the JSON Web Tokens web site.

Related Topics

- Checking the Entra ID Access Token Version
  You can check the version of the Entra ID access token that your site uses by using the JSON Web Tokens web site.

Checking the Entra ID Access Token Version

You can check the version of the Entra ID access token that your site uses by using the JSON Web Tokens web site.

By default, Entra ID v1 access token, but your site may have chosen to use v2. Oracle Database supports v1 tokens and Autonomous Database Serverless supports v2 tokens, as well. If you want to use the v2 access tokens, then you can enable their use for the Oracle database. To find the version of the Entra ID access token that you are using, you can either check with your Entra ID administrator, or confirm the version from the JSON Web Tokens website, as follows.

1. Go to the JSON Web Tokens website.

   ```
   https://jwt.io/
   ```

2. Copy and paste the token string into the **Encoded** field.
3. Check the **Decoded** field, which displays information about the token string.

   Near or at the bottom of the field, you will see a claim entitled `ver`, which indicates either of the following versions:

   - `"ver": "1.0"`
   - `"ver": "2.0"`

Related Topics

- Enabling Microsoft Entra ID v2 Access Tokens
  Oracle Database supports integration with the v1 and v2 Azure AD `OAuth2` access token.

Configuring Microsoft Entra ID as an External Identity Provider for Autonomous Database

An Autonomous Database administrator can enable Entra ID as an external identity provider on an Autonomous Database instance.

To enable Entra ID as an external identity provider:

1. Log in to the Autonomous Database instance as a user who has the `EXECUTE` privilege on the `DBMS_CLOUD_ADMIN` PL/SQL package. The `ADMIN` user has this privilege.
2. Run the `DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION` procedure with the Entra ID required parameters.

```
BEGIN
  DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
      type   =>'AZURE_AD',
      params => JSON_OBJECT('tenant_id' VALUE 'tenant_id',
                            'application_id' VALUE 'application_id',
                            'application_id_uri' VALUE
'application_id_uri'),
      force => TRUE
  );
END;
```

In this procedure the Entra ID parameters are:

- `type`: Specifies the external authentication provider. For Entra ID, as shown, use `'AZURE_AD'`.
- `params`: Values for the required Entra ID parameters are available from the Azure portal on the app registration Overview pane for Azure Active Directory. The required `params` for Entra ID are:
  - `tenant_id`: Tenant ID of the Azure Account. Tenant Id specifies the Autonomous Database instance's Entra ID application registration.
  - `application_id`: Azure Application ID created in Entra ID to assign roles/schema mappings for external authentication in the Autonomous Database instance.
  - `application_id_uri`: Unique URI assigned to the Azure Application.
    This it the identifier for the Autonomous Database instance. The name must be domain qualified (this supports cross tenancy resource access).

    The maximum length for this parameter is 256 characters.
- `force`: Set this parameter to `TRUE` if another `EXTERNAL AUTHENTICATION` method is configured for the Autonomous Database instance and you want to disable it.

For example:

```
BEGIN
  DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
      type   =>'AZURE_AD',
      params => JSON_OBJECT('tenant_id' VALUE '29981886-6fb3-44e3-82',
                            'application_id' VALUE '11aa1a11-aaa',
```

```
                              'application_id_uri' VALUE 'https://
example.com/111aa1aa'),
      force  => TRUE
  );
END;
```

This sets the `IDENTITY_PROVIDER_TYPE` system parameter.

For example, you can use the following to verify `IDENTITY_PROVIDER_TYPE`:

```
SELECT NAME, VALUE FROM V$PARAMETER WHERE NAME='identity_provider_type';

NAME                   VALUE
---------------------- --------
identity_provider_type AZURE_AD
```

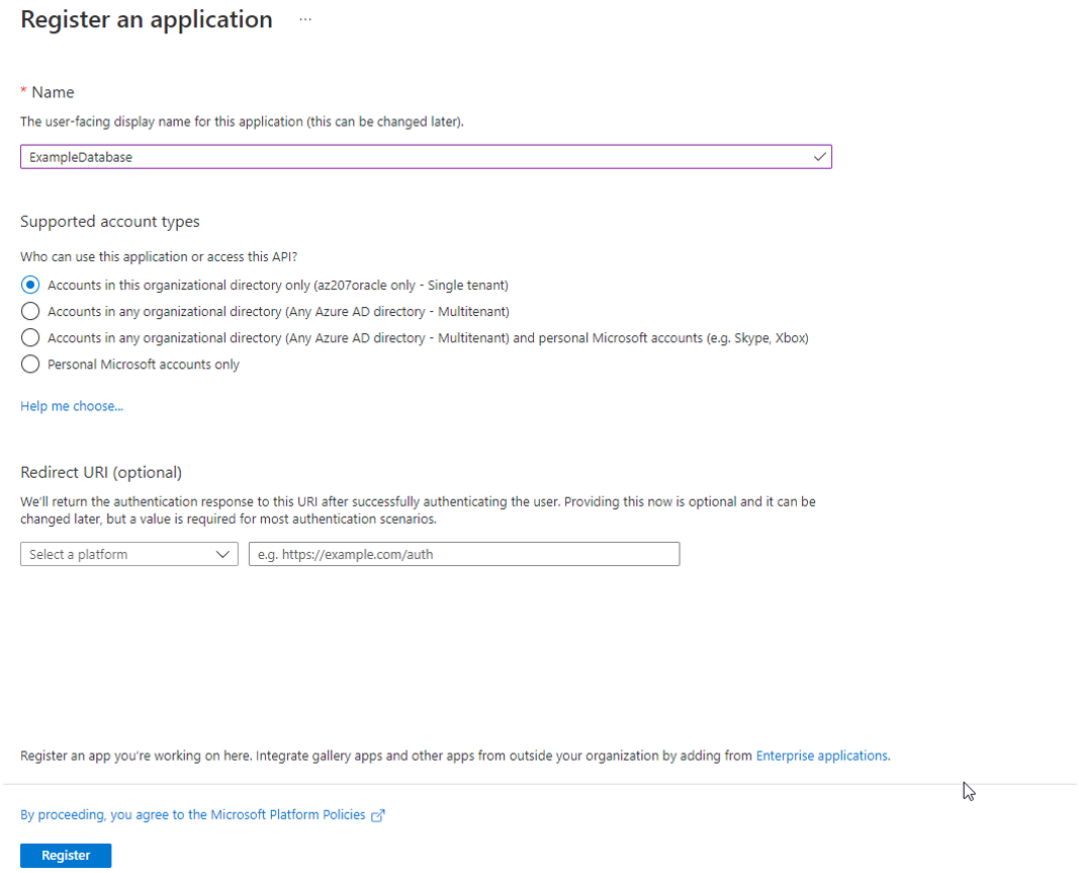
See ENABLE_EXTERNAL_AUTHENTICATION Procedure for more information.

Registering the Oracle Database Instance with a Microsoft Entra ID Tenancy

A user with Entra ID administrator privileges uses Microsoft Entra ID to register the Oracle Database instance with the Microsoft Entra ID tenancy.

1. Log in to the Azure portal as an administrator who has Microsoft Entra ID privileges to register applications.
2. In the Azure Active directory admin center page, from the left navigation bar, select **Azure Active Directory**.
3. In the MS - App registrations page, select **App registrations** from the left navigation bar.
4. Select **New registration**.

The Register an application window appears.



5. In the Register an application page, enter the following Oracle Database instance registration information:
   - In the **Name** field, enter a name for the Oracle Database instance connection (for example, *`Example Database`*).
   - Under Supported account types, select the account type that matches your use case.
     - **Accounts in this organizational directory only (*tenant_name* only - Single tenant)**
     - **Accounts in any organizational directory (Any Entra ID directory - Multitenant)**
     - **Accounts in any organizational directory (Any Entra ID directory - Multitenant) and personal Microsoft accounts (e.g. Skype, Xbox)**
     - **Personal Microsoft accounts only**
6. Bypass the Redirect URI (Optional) settings. You do not need to create a redirect URI because Entra ID does not need one for the database server.
7. Click **Register**.

   After you click **Register**, Entra ID displays the app registration's Overview pane, which will show the Application (client) ID under Essentials. This value is a unique identifier for the application in the Microsoft identity platform. Note the term Application refers to the Oracle Database instance.
8. Register a scope for the database app registration.

A scope is a permission to access the database. Each database will need a scope so that clients can establish a trust with the database by requesting permission to use the database scope. This allows the database client to get access tokens for the database.

a. In the left navigation bar, select **Expose an API**.

b. Under Set the App ID URI, in the **Application ID URI** field, enter the app ID URI for the database connection using the following format, and then click **Save**:

   your_tenancy_url/application_(client)_id

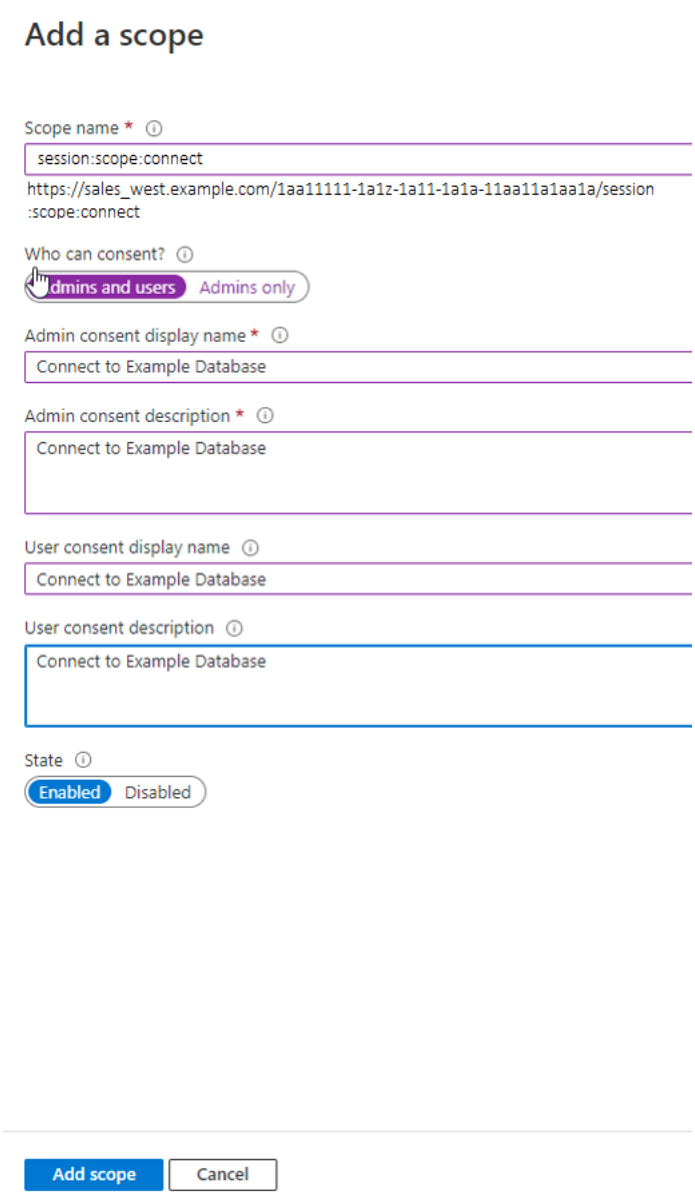
   In this specification:

   - *your_tenancy_url* must include `https` as the prefix and the fully qualified domain name of your Entra ID tenancy.
   - *application_(client)_id* is the ID that was generated when you registered the Oracle Database instance with Entra ID. It is displayed in the Overview pane of the app registration.

   For example:

   ```
   https://sales_west.example.com/1aa11111-1a1z-1a11-1a1a-11aa11a1aa1a
   ```

c. Select **Add a scope** and then enter the following settings:



- **Scope name** specifies a name for the scope. Enter the following name:

  ```
  session:scope:connect
  ```

  This name can be any text. However, a scope name must be provided. You will need to use this scope name later when you give consent to the database client application to access the database.
- **Who can consent** specifies the necessary permissions. Select **Admins and users**, or for higher restrictions, **Admins only**.
- **Admin consent** display name describes the scope's purpose (for example, `Connect to Oracle`), which only administrators can see.
- **Admin consent display name** describes the scope's purpose (for example, `Connect to Example Database`), which only administrators can see.

- **User consent display name** is a short description of the purpose of the scope (for example, `Connect to Example Database`), which users can see if you specify **Admins and users** in **Who can consent**.
- **User consent description** is a more detailed description of the purpose of the scope (for example, `Connect to Example Database`), which users can see if you specify **Admins and users** in **Who can consent**.
- **State** enables or disables the connection. Select **Enabled**.

After you complete these steps, you are ready to add one or more Azure app roles, and then perform the mappings of Oracle schemas and roles.

Related Topics

- Quickstart: Register an application with the Microsoft identity platform

Managing App Roles in Microsoft Entra ID

In Entra ID, you can create and manage app roles that will be assigned to Azure users and groups and also be mapped to Oracle Database global schemas and roles.

- Creating a Microsoft Entra ID App Role
  Azure users, groups, and applications that need to connect to the database will be assigned to the database app roles.
- Assigning Users and Groups to the Microsoft Entra ID App Role
  Before Microsoft Azure users can have access to the Oracle database, they must first be assigned to the app roles that will be mapped to Oracle Database schema users or roles.
- Assigning an Application to an App Role
  An application that must connect to the database using the client credential flow must to be assigned to an app role.

Creating a Microsoft Entra ID App Role

Azure users, groups, and applications that need to connect to the database will be assigned to the database app roles.

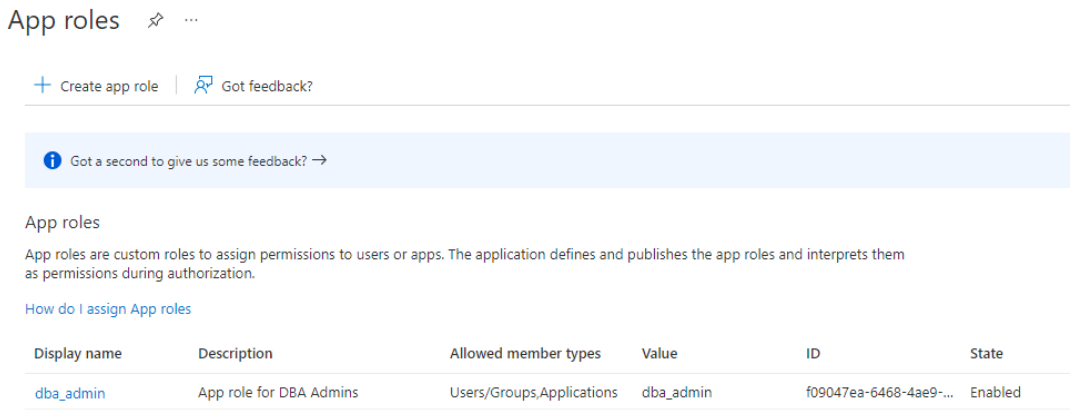
See the Microsoft Azure article Create and assign a custom role in Azure Active Directory for detailed steps on how to create an app role. The following steps describe how to create the app role for use with an Oracle database.

1. Log in to Entra ID as an administrator who has privileges for creating app roles.
2. Access the Oracle Database app registration that you created.
   a. Use the **Directory + subscription** filter to locate the Entra ID tenant that contains the Oracle Database app registration.
   b. Select **Azure Active Directory**.
   c. Under **Manage**, select **App registrations**, and then select the Oracle Database instance that you registered earlier.
3. Under **Manage**, select **App roles**.
4. In the App roles page, select **Create app role**.
5. In the Create app role page, enter the following information:
   - **Display name** is the displayed name of the role (for example, `HR App Schema`). You can include spaces in this name.

- **Value** is the actual name of the role (for example, `HR_APP`). Ensure that this setting matches exactly the string that is referenced in the database mapping to a schema or role. Do not include spaces in this name.
- **Description** provides a description of the purpose of this role.
- **Do you want to enable this app role?** enables you to activate the role.

6. Click **Apply**.

   The app role appears in the App roles pane.



Assigning Users and Groups to the Microsoft Entra ID App Role

Before Microsoft Azure users can have access to the Oracle database, they must first be assigned to the app roles that will be mapped to Oracle Database schema users or roles.

See the Microsoft Azure article Add app roles to your application and receive them in the token for detailed steps assigning users and groups to an app role. The following steps explain how to do this for an Oracle database.

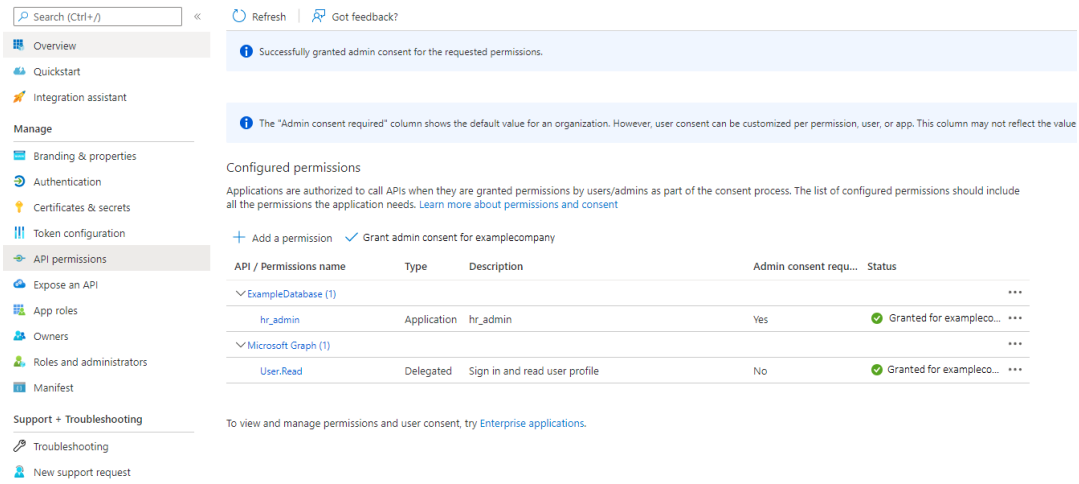
1. Log in to Entra ID as an administrator who has privileges for assigning Azure users and Entra ID groups to app roles.
2. In enterprise applications, find the name of the Oracle Database app registration that you created. This is automatically created when you create an app registration.
   a. Use the **Directory + subscription** filter to locate the Azure Active Directory tenant that contains the Oracle connection.
   b. Select **Azure Active Directory**.
   c. Under **Manage**, select **Enterprise applications**, and then select the Oracle Database app registration name that you registered earlier.
3. Under Getting Started, select **Assign users and groups**.
4. Select **Add user/group**.
5. In the Add assignment window, select **Users and groups** to display a list of users and security groups.
6. From this list, select the users and groups that you want to assign to the app role, and then click **Select**.
7. In the Add assignment window, select **Select a role** to display a list of the app roles that you have created.

8. Select the app role and then select **Select**.
9. Click **Assign**.

Assigning an Application to an App Role

An application that must connect to the database using the client credential flow must to be assigned to an app role.

1. Log in to Entra ID as an administrator who has privileges for assigning Azure users and Entra ID groups to app roles.
2. Access the app registration for the application.
3. Under Manage, select **API permissions**.
4. In the Configured permissions area, select **+ Add a permission**.
5. In the Request API permission pane, select the **My APIs** tab.
6. Select the Oracle Database app that you want to give permission for this application to access. Then select the **Application permissions** option.
7. Select the database app roles to assign to the application and then click the **Add Permission** box at the bottom of the screen to assign the app roles and close the dialog box. Ensure that the app roles that you just assigned appear under Configured permissions.



8. Select **Grant admin consent for *tenancy*** to grant consent for the tenancy users, then select **Yes** in the confirmation dialog box.

Related Topics

- Configure the admin consent workflow

Role and Schema Mapping for Entra ID Authentication on Autonomous Database

Microsoft Entra ID (Azure AD) users are mapped to one database schema and optionally to one or more database roles. After mapping Entra ID users, user can connect to the Autonomous Database instance.

- Exclusively Mapping an Oracle Database Schema to a Microsoft Azure User
  You can exclusively map an Oracle Database schema to a Microsoft Azure user.

- Mapping a Shared Oracle Schema to an App Role
  In this mapping, an Oracle schema is mapped to an app role. Therefore, anyone who has that app role would get the same shared schema.
- Mapping an Oracle Database Global Role to an App Role
  Oracle Database global roles that are mapped to Entra ID app roles give Azure users and applications additional privileges and roles above those that they have been granted through their login schemas.

Exclusively Mapping an Oracle Database Schema to a Microsoft Azure User

You can exclusively map an Oracle Database schema to a Microsoft Azure user.

1. Log in to the Oracle Database instance as a user who has been granted the `CREATE USER` or `ALTER USER` system privilege.
2. Run the `CREATE USER` or `ALTER USER` statement with the `IDENTIFIED GLOBALLY AS` clause specifying the Azure user name.

   For example, to create a new database schema user named `peter_fitch` and map this user to an existing Azure user named `peter.fitch@example.com`:

   ```
   CREATE USER peter_fitch IDENTIFIED GLOBALLY AS
   'AZURE_USER=peter.fitch@example.com';
   ```

3. Grant the `CREATE SESSION` privilege to the user.

   ```
   GRANT CREATE SESSION TO peter_fitch;
   ```

Mapping a Shared Oracle Schema to an App Role

In this mapping, an Oracle schema is mapped to an app role. Therefore, anyone who has that app role would get the same shared schema.

1. Log in to the Oracle Database instance as a user who has the `CREATE USER` or `ALTER USER` system privilege.
2. Run the `CREATE USER` or `ALTER USER` statement with the `IDENTIFIED GLOBALLY AS` clause specifying the Azure application role name.

   For example, to create a new database global user account (schema) named `dba_azure` and map it to an existing Entra ID application role named `AZURE_DBA`:

   ```
   CREATE USER dba_azure IDENTIFIED GLOBALLY AS 'AZURE_ROLE=AZURE_DBA';
   ```

Mapping an Oracle Database Global Role to an App Role

Oracle Database global roles that are mapped to Entra ID app roles give Azure users and applications additional privileges and roles above those that they have been granted through their login schemas.

1. Log in to the Oracle Database instance as a user who has been granted the `CREATE ROLE` or `ALTER ROLE` system privilege
2. Run the `CREATE ROLE` or `ALTER ROLE` statement with the `IDENTIFIED GLOBALLY AS` clause specifying the name of the Entra ID application role.

For example, to create a new database global role named `widget_sales_role` and map it to an existing Entra ID application role named `WidgetManagerGroup`:

```
CREATE ROLE widget_sales_role IDENTIFIED GLOBALLY AS
'AZURE_ROLE=WidgetManagerGroup';
```

Microsoft Entra ID Client Configuration and Access for Autonomous Database

After you configure Entra ID on your Autonomous Database and you map Entra ID users, there are numerous ways that a user can configure a client to the Autonomous Database instance using Entra ID tokens.

If you use a wallet for your client connection (mTLS), download the wallet from the Autonomous Database instance and then follow the directions for configuring it for use with your client.

Confirm that DN matching is enabled by looking for `SSL_SERVER_DN_MATCH=ON` in `sqlnet.ora`. Partial or full DN matching is required when sending a token from the database client to Autonomous Database.

> **Note:**
>
> If Autonomous Database is using a private endpoint, you need to specify a host value for the connect string parameter. Using an IP address for the host parameter in the connect string will not work with DN matching and the Entra ID token will not be sent to the database.
> See Private Endpoints Configuration Examples on Autonomous Database for configuration information on how to set the host parameter when using a private endpoint.

See Configuring Entra ID Client Connections to the Oracle Database for more information.

Testing the Accessibility of the Entra ID Endpoint

You must ensure that your Oracle Database can access the Entra ID endpoint.

If your database client is configured to get Microsoft Entra ID OAuth2 tokens, then the database client must be able to access the Entra ID endpoint. Run the following command to check if you have internet access:

```
curl https://login.windows.net/common/discovery/keys
```

A status code of 200 indicates success.

Check with your IT help desk for the proxy information if you weren't successful running this command.

For an Oracle database to accept Entra ID `OAuth2` tokens, the database must request the public key from the Microsoft Entra ID endpoint.

- Run the following test to determine if the database can connect with the Microsoft Entra ID endpoint:

```
SET SERVEROUTPUT ON SIZE 40000
DECLARE
  req UTL_HTTP.REQ;
  resp UTL_HTTP.RESP;
BEGIN
  UTL_HTTP.SET_WALLET(path => 'system:');
  req := UTL_HTTP.BEGIN_REQUEST('https://login.windows.net/common/
discovery/keys');
  resp := UTL_HTTP.GET_RESPONSE(req);
  DBMS_OUTPUT.PUT_LINE('HTTP response status code: ' || resp.status_code);
  UTL_HTTP.END_RESPONSE(resp);
END;
/
```

If this test is successful, then a `PL/SQL procedure successfully completed` message appears.

If the following messages appear, then it means that a database network access control list (ACL) policy blocked your test and you will need to temporarily set an access control list policy to allow you to test this:

```
ORA-29273: HTTP request failed
ORA-24247: network access denied by access control list (ACL)
```

1. Set the ACL as follows:

```
BEGIN
DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
  host => '*',
  ace  =>  xs$ace_type(privilege_list => xs$name_list('connect'),
                       principal_name => 'username_placeholder',
                       principal_type => xs_acl.ptype_db));
END;
/
```

   Replace *username_placeholder* with the user name of the database user who is running the test. For example:

```
BEGIN
DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
  host => '*',
  ace  =>  xs$ace_type(privilege_list => xs$name_list('connect'),
                       principal_name => 'DBA_DEBRA',
                       principal_type => xs_acl.ptype_db));
END;
/
```

2. Try running the test again.

3. Remove the ACL, because you now no longer need it. For example, assuming your user name is `dba_debra`:

```
BEGIN
DBMS_NETWORK_ACL_ADMIN.REMOVE_HOST_ACE(
  host => '*',
  ace  =>  xs$ace_type(privilege_list => xs$name_list('connect'),
                       principal_name => 'DBA_DEBRA',
                       principal_type => xs_acl.ptype_db));
END;
/
```

If the database cannot connect with the Microsoft Entra ID endpoint, even after you set the ACL policy, you will most likely need to set the `HTTP_PROXY` package for your database. Review the topics listed in Related Topics, depending if you are using a default Oracle Database environment or an Oracle Real Application Clusters RAC environment. Your network administrator should be able to tell you what the correct `HTTP_PROXY` setting should be.

Use Microsoft Active Directory with Autonomous Database

You can configure Autonomous Database to authenticate and authorize Microsoft Active Directory users.

This configuration allows Active Directory users to access Autonomous Database using their Active Directory credentials including passwords and Kerberos.

- Prerequisites to Configure CMU with Microsoft Active Directory on Autonomous Database
- Configure CMU with Microsoft Active Directory on Autonomous Database
- Kerberos Authentication for CMU with Microsoft Active Directory
- Add Microsoft Active Directory Roles on Autonomous Database
  To add Active Directory roles, map the database global roles to Active Directory groups with `CREATE ROLE` or `ALTER ROLE` statements (and include the `IDENTIFIED GLOBALLY AS` clause).
- Add Microsoft Active Directory Users on Autonomous Database
  To add Active Directory users to access a database, map database global users to Active Directory groups or users with `CREATE USER` or `ALTER USER` statements (with `IDENTIFIED GLOBALLY AS` clause).
- Tools Restrictions with Active Directory on Autonomous Database
- Connect to Autonomous Database with Active Directory User Credentials
  After the ADMIN user completes the CMU Active Directory configuration steps and creates global roles and global users, users log in to the database using their Active Directory username and password.
- Verify Active Directory User Connection Information with Autonomous Database
  When users log in to the database using their Active Directory username and password, you can verify and audit the user activity.
- Remove Active Directory Users and Roles on Autonomous Database
  To remove Active Directory users and roles from Autonomous Databases, use standard database commands. This does not remove the related Active Directory users or groups that were mapped from the dropped database users or roles.

- Disable Active Directory Access on Autonomous Database
  Describes the steps to remove the CMU configuration from your Autonomous Database (and disable the LDAP access from your Autonomous Database to Active Directory).

Prerequisites to Configure CMU with Microsoft Active Directory on Autonomous Database

You can configure Autonomous Database to authenticate and authorize Microsoft Active Directory users.

Depending on where the Active Directory servers reside, there are two options for configuring Autonomous Database with Centrally Managed Users (CMU) with Microsoft Active Directory:

- **Active Directory (AD) servers publicly accessible**: the Active Directory servers are accessible from Autonomous Database through the public internet.
- **Active Directory (AD) servers reside on a private endpoint**: the Active Directory servers reside on a private endpoint and are not accessible from Autonomous Database through the public internet. For this case, an extra configuration step is required as shown in last step in Configure CMU with Microsoft Active Directory on Autonomous Database where you set the database property `ROUTE_OUTBOUND_CONNECTIONS`.

> **Note:**
>
> See Use Microsoft Entra ID with Autonomous Database for information on using Azure Active Directory with Autonomous Database. The CMU option supports Microsoft Active Directory servers but does not support the Azure Active Directory service.

The integration of Autonomous Database with Centrally Managed Users (CMU) provides integration with Microsoft Active Directory. CMU with Active Directory works by mapping Oracle database global users and global roles to Microsoft Active Directory users and groups.

The following are required prerequisites to configure the connection from Autonomous Database to Active Directory:

- You must have Microsoft Active Directory installed and configured. See AD DS Getting Started for more information.
- You must create an Oracle service directory user in Active Directory. See Connecting to Microsoft Active Directory for information on the Oracle service directory user account.
- An Active Directory system administrator must have installed Oracle password filter on the Active Directory servers, and set up Active Directory groups with Active Directory users to meet your requirements.

> **Note:**
>
> This is not required if you are using Kerberos authentication for CMU Active Directory. See Kerberos Authentication for CMU with Microsoft Active Directory for more information.

If you use password authentication with CMU Active Directory for Autonomous Database, you must use the included utility `opwdintg.exe` to install the Oracle password filter on

Active Directory, extend the schema, and create three new `ORA_VFR` groups for three types of password verifier generation. See Connecting to Microsoft Active Directory for information on installing the Oracle password filter.

- You need the CMU configuration database wallet, `cwallet.sso` and the CMU configuration file `dsi.ora` to configure CMU for your Autonomous Database:
  - If you have configured CMU for an on-premise database, you can obtain these configuration files from your on-premise database server.
  - If you have not configured CMU for an on-premise database, you need to create these files. Then you upload the configuration files to the cloud to configure CMU on your Autonomous Database instance. You can validate the wallet and the `dsi.ora` by configuring CMU for an on-premise database and verifying that an Active Directory user can successfully log on to the on-premise database with these configuration files.

  For details on the wallet file for CMU, see Create the Wallet for a Secure Connection and Verify the Oracle Wallet.

  For details on the `dsi.ora` file for CMU, see Creating the dsi.ora File.

  For details on configuring Active Directory for CMU and troubleshooting CMU for on-premise databases, see How To Configure Centrally Managed Users For Database Release 18c or Later Releases (Doc ID 2462012.1).
- Port 636 of the Active Directory servers must be open to Autonomous Database in Oracle Cloud Infrastructure. This allows Autonomous Database to access the Active Directory servers.
- When the Active Directory servers are on a public endpoint:
  - The Active Directory servers must be accessible from Autonomous Database through the public internet.
  - You can also extend your on-premise Active Directory to Oracle Cloud Infrastructure where you can set up Read Only Domain Controllers (RODCs) for the on-premise Active Directory. This allows you to use the RODCs in Oracle Cloud Infrastructure to authenticate and authorize on-premise Active Directory users for access to Autonomous Databases.

    See Extend Active Directory integration in Hybrid Cloud for more information.

Configure CMU with Microsoft Active Directory on Autonomous Database

You can configure Autonomous Database to authenticate and authorize Microsoft Active Directory users.

To configure Autonomous Database for CMU to connect to Active Directory:

> **Note:**
>
> When you perform the configuration steps, connect to the database as the ADMIN user.

1. Verify if another external authentication scheme is enabled on your database, and disable it.

   You can continue with CMU-AD configuration on top of Kerberos to provide CMU-AD Kerberos authentication for Microsoft Active Directory users.

   See Kerberos Authentication for CMU with Microsoft Active Directory for more information.

2. Upload the CMU configuration files, including the database wallet file, `cwallet.sso` and the CMU configuration file, `dsi.ora` to your Object Store. This step depends on the Object Store you use.

   The `dsi.ora` configuration file contains the information to find the Active Directory servers.

   If you are using Oracle Cloud Infrastructure Object Store, see Putting Data into Object Storage for details on uploading files.

3. Run `DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION` procedure and pass in a location URI with the `params` JSON argument. You must place the configuration files `cwallet.sso` and `dsi.ora` in the Object Storage location specified in the `location_uri` parameter.

   For example:

   ```
   BEGIN
      DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
          type     => 'CMU',
          params   => JSON_OBJECT('location_uri' value 'https://
   objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o',
                                   'credential_name' value
   'my_credential_name')
      );
   END;
   /
   ```

   Oracle recommends that you store the CMU configuration files in a private bucket in your Object Store.

   In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

   The `credential_name` you use in this step is the credentials to access the Object Store.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

   If the `location_uri` is a pre-authenticated URL or a pre-signed URL, then supplying a `credential_name` is not required.

   The procedure creates a directory object named `CMU_WALLET_DIR` in your database and copies the CMU configuration files from the Object Store location to the directory object. This procedure also sets the database property `CMU_WALLET` to the value `'CMU_WALLET_DIR'` and sets the `LDAP_DIRECTORY_ACCESS` parameter value to the value `PASSWORD` to enable access from the Autonomous Database instance to Active Directory.

4. After you enable CMU authentication, remove the CMU configuration files including the database wallet `cwallet.sso` and the CMU configuration file `dsi.ora` from Object Store. You can use local Object Store methods to remove these files or use `DBMS_CLOUD.DELETE_OBJECT` to delete the files from Object Store.

5. When Active Directory servers are on a private endpoint, perform additional configuration steps to provide access to the private endpoint.

   a. Set the database property `ROUTE_OUTBOUND_CONNECTIONS` to the value `'PRIVATE_ENDPOINT'`.

See Enhanced Security for Outbound Connections with Private Endpoints for more information.

b. Validate that the CMU-AD configuration file `dsi.ora` includes host names. When `ROUTE_OUTBOUND_CONNECTIONS` is set to `'PRIVATE_TARGET'`, IP addresses cannot be specified in `dsi.ora`.

Note for CMU with Active Directory on Autonomous Database:

- Only "password authentication" or Kerberos authentication are supported for CMU with Autonomous Database. When you are using CMU authentication with Autonomous Database, other CMU authentication methods such as Azure AD, OCI IAM, and PKI are not supported.

See Disable Active Directory Access on Autonomous Database for instructions to disable the access from Autonomous Database to Active Directory.

See ENABLE_EXTERNAL_AUTHENTICATION Procedure for information on `DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION`.

See Configuring Centrally Managed Users with Microsoft Active Directory for more information on configuring CMU with Microsoft Active Directory.

Kerberos Authentication for CMU with Microsoft Active Directory

You can configure Autonomous Database to use Kerberos authentication for CMU with Microsoft Active Directory users. This configuration allows CMU Active Directory (CMU-AD) users to access an Autonomous Database instance using Kerberos credentials.

Kerberos can be configured with or without CMU-AD. Just configuring Kerberos requires you to create and maintain a database user for every Kerberos user. Configuring Kerberos with CMU allows you to map an Active Directory group of Kerberos users to a single database user, shared schema, so database access can be controlled by Active Directory group membership. See Configure Kerberos Authentication with Autonomous Database for details on configuring Kerberos without CMU-AD.

> **Note:**
>
> When implementing both the Kerberos authentication and CMU-AD for authorization, Oracle recommends implementing Kerberos authentication first, and then adding CMU-AD authorization.

1. Enable Kerberos in your Autonomous Database instance using Microsoft Active Directory Kerberos server.

   Only Microsoft Active Directory Kerberos servers are supported for Kerberos when you configure Kerberos authentication with CMU-AD.

   a. To enable Kerberos authentication for your Autonomous Database, you must obtain the Kerberos configuration files: `krb.conf` and the service key table file `v5srvtab`.

      To generate these files when you configure Kerberos authentication with CMU-AD you need the server host name. You can obtain the value of the server host from the attribute `PUBLIC_DOMAIN_NAME` in the `CLOUD_IDENTITY` column of `V$PDBS`. This value is different from the Fully Qualified Domain Name (FQDN) for a database on a private endpoint.

Use the following command to obtain the server host name:

```
SELECT guid ||'/'|| json_value(cloud_identity, '$.PUBLIC_DOMAIN_NAME')
    "KSERVICE/KINSTANCE" FROM v$pdbs;
```

You can use a command such as the following to generate the service key table file:

```
ktpass -princ ORACLE/
DATABASE_SERVER_HOST_NAME.DATABASE_SERVER_HOST_DOMAIN@ACTIVE_DIRECTORY_D
EFAULT_DOMAIN
                -pass ACTIVE_DIRECTORY_PASSWORD
                -mapuser DATABASE_SERVER_HOST_NAME
                -crypto ALL
                -ptype KRB5_NT_PRINCIPAL
                -out database.keytab
```

For example:

```
ktpass -princ ORACLE/user.example.com@example.com
                -pass password -mapuser dbexamplekrb
                -crypto ALL -ptype KRB5_NT_PRINCIPAL -out
database.keytab
```

For more information on these files and steps required to obtain them, see Configuring Kerberos Authentication.

b. Copy the Kerberos configuration files `krb.conf` and `v5srvtab` to a bucket in your Object Store.

   This step differs depending on the Object Store you use.

   If you are using Oracle Cloud Infrastructure Object Store, see Putting Data into Object Storage for details on uploading files.

c. Run `DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION` to enable Kerberos external authentication.

   For example:

```
BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
       type     => 'KERBEROS',
       params   => JSON_OBJECT('location_uri' value 'https://
objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/
bucketname/o',
                                'credential_name' value
'my_credential_name')
   );
END;
/
```

> **Note:**
>
> Oracle recommends that you store the Kerberos configuration files in a private bucket in your Object Store.

In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

The `credential_name` you use in this step is the credentials for the Object Store.

See Enable Kerberos Authentication on Autonomous Database for more information.

d. Confirm that Kerberos is configured and enabled.

```
SELECT property_value FROM database_properties
      WHERE property_name='KERBEROS_DIRECTORY';
```

2. Enable and configure CMU-AD on Autonomous Database.

a. Upload the CMU configuration files, including the database wallet file, `cwallet.sso` and the CMU configuration file, `dsi.ora` to your Object Store.

You upload the `cwallet.sso` so that the CMU-AD configuration has the credentials for the Autonomous Database instance to connect to the Active Directory service account.

The `dsi.ora` configuration file contains the information to find the Active Directory servers.

This step differs depending on the Object Store you use.

If you are using Oracle Cloud Infrastructure Object Store, see Putting Data into Object Storage for details on uploading files.

b. Run the `DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION` procedure and pass in a location URI with the `params` JSON argument. You must place the configuration files `cwallet.sso` and `dsi.ora` in the Object Storage location specified in the `location_uri` parameter.

For example:

```
BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
       type     => 'CMU',
       params   => JSON_OBJECT('location_uri' value 'https://
objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/
bucketname/o',
                               'credential_name' value
'my_credential_name')
   );
END;
/
```

Oracle recommends that you store the CMU configuration files in a private bucket in your Object Store.

In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

The `credential_name` you use in this step is the credentials to access the Object Store.

Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

If the `location_uri` is a pre-authenticated URL or a pre-signed URL, then supplying a `credential_name` is not required.

See Prerequisites to Configure CMU with Microsoft Active Directory on Autonomous Database for more information.

c. Confirm that CMU-AD is configured and enabled.

```
SELECT property_value FROM database_properties
      WHERE property_name='CMU_WALLET';
```

3. After completing Step 1 and Step 2, verify that the configuration of Kerberos authentication with CMU-AD is complete.

   a. Login to the Autonomous Database instance as an Active Directory user so that `SYS_CONTEXT` information populated in your `USERENV`.

   b. Query the `SYS_CONTEXT USERENV`.

```
SELECT SYS_CONTEXT('USERENV','AUTHENTICATION_METHOD') FROM DUAL;

SYS_CONTEXT('USERENV','AUTHENTICATION_METHOD')

--------------------------------------------------------------------------------
--------
KERBEROS_GLOBAL
```

When you have Kerberos authentication configured and enabled without CMU-AD, this query returns: `KERBEROS`. See Configure Kerberos Authentication with Autonomous Database for more information.

When you have CMU-AD authentication configured without Kerberos, this query returns: `PASSWORD_GLOBAL`. See Prerequisites to Configure CMU with Microsoft Active Directory on Autonomous Database for more information.

Notes for using Kerberos authentication with CMU-AD:

- You do not need to add the password filter when using Kerberos authentication with CMU-AD. See Prerequisites to Configure CMU with Microsoft Active Directory on Autonomous Database for more information.
- Adding or removing Active Directory users is supported, in the same manner as with CMU with Active Directory when you are using password authentication. See Add Microsoft Active Directory Users on Autonomous Database for more information.
- The existing restrictions about authenticating against the Autonomous Database Built-in Tools with CMU with Active Directory password also apply to CMU with Active Directory with Kerberos authentication. See Tools Restrictions with Active Directory on Autonomous Database for more information.

- Use `DBMS_CLOUD_ADMIN.DISABLE_EXTERNAL_AUTHENTICATION` to disable CMU-AD with Kerberos authentication. See DISABLE_EXTERNAL_AUTHENTICATION Procedure for more information.
- When the CMU-AD servers are on a private endpoint, to use CMU-AD with Kerberos authentication the server host name used for generating the key tab must be set to the value of attribute `PUBLIC_DOMAIN_NAME` in the `CLOUD_IDENTITY` column of `V$PDBS`. This value is different from the FQDN for a private endpoint database.

Add Microsoft Active Directory Roles on Autonomous Database

To add Active Directory roles, map the database global roles to Active Directory groups with `CREATE ROLE` or `ALTER ROLE` statements (and include the `IDENTIFIED GLOBALLY AS` clause).

To add global roles for Active Directory groups on Autonomous Database:

1. Log in as the ADMIN user to the database that is configured to use Active Directory (the ADMIN user has the `CREATE ROLE` and `ALTER ROLE` system privileges that you need for these steps).
2. Set the database authorization for Autonomous Database roles with `CREATE ROLE` or `ALTER ROLE` statement. Include the `IDENTIFIED GLOBALLY AS` clause and specify the DN of an Active Directory group.

   Use the following syntax to map a directory user group to a database global role:

   ```
   CREATE ROLE global_role IDENTIFIED GLOBALLY AS
        'DN_of_an_AD_GROUP_of_WHICH_the_AD_USER_IS_a_MEMBER';
   ```

   For example:

   ```
   CREATE ROLE widget_sales_role IDENTIFIED GLOBALLY AS
        'CN=widget_sales_group,OU=sales,DC=production,DC=example,DC=com';
   ```

   In this example all members of the `widget_sales_group` are authorized with the database role `widget_sales_role` when they log in to the database.
3. Use `GRANT` statements to grant the required privileges or other roles to the global role.

   For example:

   ```
   GRANT CREATE SESSION TO WIDGET_SALES_ROLE;
   GRANT DWROLE TO WIDGET_SALES_ROLE;
   ```

   `DWROLE` is a predefined role that has common privileges defined. See Manage User Privileges on Autonomous Database - Connecting with a Client Tool for information on setting common privileges for Autonomous Database users.
4. If you want to make an existing database role to be associated with an Active Directory group, then use `ALTER ROLE` statement to alter the existing database role to map the role to an Active Directory group.

Use the following syntax to alter an existing database role to map it to an Active Directory group:

```
ALTER ROLE existing_database_role
   IDENTIFIED GLOBALLY AS
'DN_of_an_AD_GROUP_of_WHICH_the_AD_USER_IS_a_MEMBER';
```

5. If you want to create additional global role mappings for other Active Directory groups, follow these steps for each Active Directory group.

For more information on configuring roles with Microsoft Active Directory, see *Configuring Authorization for Centrally Managed Users* in Oracle Database 19c Security Guide or Oracle Database 23ai Security Guide.

Add Microsoft Active Directory Users on Autonomous Database

To add Active Directory users to access a database, map database global users to Active Directory groups or users with `CREATE USER` or `ALTER USER` statements (with `IDENTIFIED GLOBALLY AS` clause).

The integration of Autonomous Database with Active Directory works by mapping Microsoft Active Directory users and groups directly to Oracle database global users and global roles.

To add global users for Active Directory groups or users on Autonomous Database:

1. Log in as the ADMIN user to the database that is configured to use Active Directory (the ADMIN user has the required `CREATE USER` and `ALTER USER` system privileges that you need for these steps).
2. Set database authorization for Autonomous Database users with `CREATE USER` or `ALTER USER` statements and include the `IDENTIFIED GLOBALLY AS` clause, specifying the DN of an Active Directory user or group.

   Use the following syntax to map a directory user to a database global user:

   ```
   CREATE USER global_user IDENTIFIED GLOBALLY AS 'DN_of_an_AD_USER';
   ```

   Use the following syntax to map a directory group to a database global user:

   ```
   CREATE USER global_user IDENTIFIED GLOBALLY AS
       'DN_of_an_AD_GROUP_of_WHICH_the_AD_USER_IS_a_MEMBER';
   ```

   For example, to map a directory group named `widget_sales_group` in the `sales` organization unit of the `production.example.com` domain to a shared database global user named `WIDGET_SALES`:

   ```
   CREATE USER widget_sales IDENTIFIED GLOBALLY AS
        'CN=widget_sales_group,OU=sales,DC=production,DC=example,DC=com';
   ```

   This creates a shared global user mapping. The mapping, with global user `widget_sales`, is effective for all users in the Active Directory group. Thus, anyone in the `widget_sales_group` can log in to the database using their Active Directory credentials (through the shared mapping of the `widget_sales` global user).

3. If you want Active Directory users to use an existing database user, own its schema, and own its existing data, then use `ALTER USER` to alter an existing database user to map the user to an Active Directory group or user.

   - Use the following syntax to alter an existing database user to map it to an Active Directory user:

     ```
     ALTER USER existing_database_user IDENTIFIED GLOBALLY AS
     'DN_of_an_AD_USER';
     ```

   - Use the following syntax to alter an existing database user to map it to an Active Directory group:

     ```
     ALTER USER existing_database_user
          IDENTIFIED GLOBALLY AS
     'DN_of_an_AD_GROUP_of_WHICH_the_AD_USER_IS_a_MEMBER';
     ```

4. If you want to create additional global user mappings for other Active Directory groups or users, follow these steps for each Active Directory group or user.

For more information on configuring roles with Microsoft Active Directory, see *Configuring Authorization for Centrally Managed Users* in Oracle Database 19c Security Guide or Oracle Database 23ai Security Guide.

Tools Restrictions with Active Directory on Autonomous Database

Notes for using Autonomous Database tools with Active Directory:

- Oracle APEX is not supported for Active Directory users with Autonomous Database. See Create Oracle APEX Workspaces in Autonomous Database for information on using regular database users with Autonomous Database.
- Database Actions is not supported for Active Directory users with Autonomous Database. See Provide Database Actions Access to Database Users for information on using regular database users with Autonomous Database.
- Oracle Machine Learning Notebooks are not supported for Active Directory users with Autonomous Database. See Add Existing Database User Account to Oracle Machine Learning Components for information on using regular database users with Autonomous Database.

Connect to Autonomous Database with Active Directory User Credentials

After the ADMIN user completes the CMU Active Directory configuration steps and creates global roles and global users, users log in to the database using their Active Directory username and password.

> **Note:**
>
> Do not log in using a Global User name. Global User names do not have a password and connecting with a Global User name will not be successful. You must have a global user mapping in your Autonomous Database in order to log in to the database. You cannot log in to the database with only global role mappings.

- To log in to the database using an Active Directory username and password, connect as follows:

  ```
  CONNECT "AD_DOMAIN\AD_USERNAME"/
  AD_USER_PASSWORD@TNS_ALIAS_OF_THE_AUTONOMOUS_DATABASE;
  ```

  For example:

  ```
  CONNECT "production\pfitch"/password@adbname_medium;
  ```

  You need to include double quotes when the Active Directory domain is included along with the username, as with this example: `"production\pfitch"`.

  In this example, the Active Directory username is `pfitch` in domain `production`. The Active Directory user is a member of `widget_sales_group` group which is identified by its DN `'CN=widget_sales_group,OU=sales,DC=production,DC=example,DC=com'`.

After configuring CMU with Active Directory on Autonomous Database and setting up Active Directory authorization, with global roles and global users, you can connect to your database using any of the connection methods described in Connect to Autonomous Database. When you connect, if you want to use an Active Directory user then use Active Directory user credentials. For example, provide a username in this form, "*AD_DOMAIN\AD_USERNAME*" (double quotes must be included), and use your *AD_USER_PASSWORD* for the password.

If your Autonomous Database instance is in Restricted Mode, this mode only allows users with the `RESTRICTED SESSION` privilege to connect to the database. The ADMIN user has this privilege. You can use restricted access mode to perform administrative tasks such as indexing, data loads, or other planned activities. See Change Autonomous Database Operation Mode to Read/Write Read-Only or Restricted for more information.

Verify Active Directory User Connection Information with Autonomous Database

When users log in to the database using their Active Directory username and password, you can verify and audit the user activity.

For example, when the user `pfitch` logs in:

```
CONNECT "production\pfitch"/password@exampleadb_medium;
```

The Active Directory user's log on username (samAccountName) is `pfitch` and `widget_sales_group` is the Active Directory Group name, and `widget_sales` is the database global user.

After `pfitch` logs in to the database, the command `SHOW USER` shows the global user name:

```
SHOW USER;

USER is "WIDGET_SALES"
```

The following command shows the DN (Distinguished Name) of the Active Directory user:

```
SELECT SYS_CONTEXT('USERENV', 'ENTERPRISE_IDENTITY') FROM DUAL;
```

For example you can verify this centrally managed user's enterprise identity:

```
SQL> SELECT SYS_CONTEXT('USERENV', 'ENTERPRISE_IDENTITY') FROM DUAL;

SYS_CONTEXT('USERENV','ENTERPRISE_IDENTITY')
--------------------------------------------------------------------------------
cn=Peter Fitch,ou=sales,dc=production,dc=examplecorp,dc=com
```

The following command shows the "`AD_DOMAIN\AD_USERNAME`":

```
SELECT SYS_CONTEXT('USERENV', 'AUTHENTICATED_IDENTITY') FROM DUAL;
```

For example, the Active Directory authenticated user identity is captured and audited when the user logs on to the database:

```
SQL> SELECT SYS_CONTEXT('USERENV', 'AUTHENTICATED_IDENTITY') FROM DUAL;

SYS_CONTEXT('USERENV','AUTHENTICATED_IDENTITY')
--------------------------------------------------------------------------------
production\pfitch
```

See *Verifying the Centrally Managed User Logon Information* in Oracle Database 19c Security Guide or Oracle Database 23ai Security Guide for more information.

Remove Active Directory Users and Roles on Autonomous Database

To remove Active Directory users and roles from Autonomous Databases, use standard database commands. This does not remove the related Active Directory users or groups that were mapped from the dropped database users or roles.

To remove users or roles from Autonomous Database:

1. Log in to the database that is configured to use Active Directory as a user who has been granted the `DROP USER` or `DROP ROLE` system privilege.
2. Drop the global users or the global roles that are mapped to Active Directory groups or users with `DROP USER` or `DROP ROLE` statement.

   See Remove Users on Autonomous Database for more information.

Disable Active Directory Access on Autonomous Database

Describes the steps to remove the CMU configuration from your Autonomous Database (and disable the LDAP access from your Autonomous Database to Active Directory).

After you configure your Autonomous Database instance to access CMU Active Directory, you can disable the access as follows:

1. Connect to the Autonomous Database as the ADMIN user.
2. Use the `DBMS_CLOUD_ADMIN.DISABLE_EXTERNAL_AUTHENTICATION` to disable CMU authentication.

> **Note:**
>
> To run this procedure you must be logged in as ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.

For example:

```
BEGIN
   DBMS_CLOUD_ADMIN.DISABLE_EXTERNAL_AUTHENTICATION;
END;
/
```

This disables CMU authentication on your Autonomous Database instance.

See DISABLE_EXTERNAL_AUTHENTICATION Procedure for more information.

Configure Kerberos Authentication with Autonomous Database

Describes how to configure Kerberos to authenticate Oracle Autonomous Database users.

- About Kerberos Authentication
  You can configure Oracle Autonomous Database to use Kerberos network authentication protocol to authenticate database users. Kerberos is a strong network authentication protocol. It uses secret-key cryptography to enable strong authentication by providing user-to-server authentication.
- Components of the Kerberos Authentication System
  Provides an overview of the Kerberos authentication system.
- Enable Kerberos Authentication on Autonomous Database
  Shows the steps to enable Kerberos authentication on your Autonomous Database instance.
- Disable Kerberos Authentication on Autonomous Database
  Shows the steps to disable Kerberos authentication for your Autonomous Database instance.
- Notes for Kerberos Authentication on Autonomous Database
  Provides notes on using Kerberos Authentication for Autonomous Database.

About Kerberos Authentication

You can configure Oracle Autonomous Database to use Kerberos network authentication protocol to authenticate database users. Kerberos is a strong network authentication protocol. It uses secret-key cryptography to enable strong authentication by providing user-to-server authentication.

- Oracle Autonomous Database support for Kerberos provides the benefits of single sign-on and centralized authentication of Oracle users. Kerberos is a trusted third-party authentication system that relies on shared secrets. It presumes that the third party is secure, and provides single sign-on capabilities, centralized password storage, database link authentication, and enhanced PC security. It does this through a Kerberos authentication server.

- The Kerberos system revolves around the concept of a ticket. A ticket is a set of electronic information that identifies a user or a service. A ticket identifies you and your network access privileges.
- In Kerberos-based authentication, you transparently send a request for a ticket to a Key Distribution Center (KDC). The Key Distribution Center authenticates you and grants you a ticket to access the database.

Components of the Kerberos Authentication System

Provides an overview of the Kerberos authentication system.

- A realm establishes an authentication administrative domain. Each realm has its own Kerberos database which contains the users and services for that particular administrative domain.
- Tickets are issued by the Key Distribution Center (KDC). Clients present tickets to the Database Server to demonstrate the authenticity of their identity. Each ticket has expiration and a renewal time.
- Keytabs stores long-term keys for one or more principals. A keytab file is generated by invoking the tool `kadmin.local` (for MIT Key Distribution Center) or `ktpass` (for Active Directory Key Distribution Center).
- Principals are the entries in the Key Distribution Center database. Each user, host or service is given a principal. A principal is a unique identity to which the Key Distribution Center can assign tickets.
- Kerberos support in Autonomous Database uses these values for various components that make up a service principal's name:

Component of Service Principal	Value in Autonomous Database
`kinstance`	You can obtain this value from the attribute `PUBLIC_DOMAIN_NAME` in the `CLOUD_IDENTITY` column of `V$PDBS`. This value is different from the Fully Qualified Domain Name (FQDN) for a database on a private endpoint. Use the following query to obtain the `kinstance`: `SELECT json_value(cloud_identity, '$.PUBLIC_DOMAIN_NAME') "KINSTANCE" FROM v$pdbs;`

Component of Service Principal	Value in Autonomous Database

kservice	On Autonomous Database you have two options for the kservice value:
	• Use the database GUID: If you do not provide a Kerberos service name with DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION, the default service name used is the GUID of the Autonomous Database instance. In this case, when you create the keytab file, use the GUID as the value for the service name.
	Because in the default case the Keytab file uses a service name that is a GUID, which is instance specific, when you use the default service name you must generate different Keytab files for each Autonomous Database instance.
	Use the following command to obtain the GUID (case is significant):
	`SELECT GUID FROM v$pdbs;`
	• Use a custom name: Set the service name when you want to use the same Keytab files on multiple Autonomous Database instances. When you use a custom name you do not need to create and upload different Keytab files for each Autonomous Database instance. When you use a custom name you must specify the params kerberos_service_name parameter with DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION. Specifying this parameter is not required when you use the GUID for the service name.
	After Kerberos is enabled on your Autonomous Database instance, use the following query to view the Kerberos service name:
	`SELECT SYS_CONTEXT('USERENV', 'KERBEROS_SERVICE_NAME') FROM DUAL;`

REALM	Any realm supported by your KDC. REALM must always be in uppercase.

To enable Kerberos authentication for your Autonomous Database, you must keep your Kerberos configuration files (krb.conf) and service key table file (v5srvtab) ready. For more information on these files and steps to obtain them, please see Configuring Kerberos Authentication.

Enable Kerberos Authentication on Autonomous Database

Shows the steps to enable Kerberos authentication on your Autonomous Database instance.

To run DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION you must be logged in as ADMIN user or have the EXECUTE privilege on DBMS_CLOUD_ADMIN.

To use DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION to enable Kerberos authentication:

To enable Kerberos authentication for your Autonomous Database, you must obtain the Kerberos configuration files: krb.conf and the service key table file v5srvtab. For more information on these files and steps required to obtain them, see Configuring Kerberos Authentication .

1. Copy the Kerberos configuration files krb.conf and v5srvtab to a bucket in your Object Store.

   If you are using Oracle Cloud Infrastructure Object Store, see Putting Data into Object Storage for details on uploading files.

2. Run DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION procedure and pass in a location URI with the params JSON argument. You must place the configuration files

`krb.conf` and `v5srvtab` in the Object Storage location specified in the `location_uri` parameter.

For example:

```
BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
       type     => 'KERBEROS',
       params   => JSON_OBJECT(
               'location_uri' value 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o',
               'credential_name' value 'my_credential_name'));
END;
/
```

> **Note:**
>
> Oracle recommends that you store the Kerberos configuration files in a private bucket in your Object Store.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

The `credential_name` you use in this step is the credentials for the Object Store.

Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

If the `location_uri` is a pre-authenticated URL then supplying a `credential_name` is not required.

This creates a directory object named `KERBEROS_DIR` in your database and uses the credential to download the Kerberos configuration files from the Object Store location to the directory object.

You can specify the `params kerberos_service_name` parameter to specify a Kerberos service name. For example:

```
BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
       type     => 'KERBEROS',
       params   => JSON_OBJECT(
               'location_uri' value 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o',
               'credential_name' value 'my_credential_name'
               'kerberos_service_name' value 'oracle' ));
END;
/
```

3. After you enable Kerberos authentication, remove the configuration `krb.conf` and `v5srvtab` from Object Store. You can use local Object Store methods to remove these files or use `DBMS_CLOUD.DELETE_OBJECT` to delete the files from Object Store.

See Navigate to Oracle Cloud Infrastructure Object Storage and Create Bucket for more information on Object Storage.

See ENABLE_EXTERNAL_AUTHENTICATION Procedure for more information.

Disable Kerberos Authentication on Autonomous Database

Shows the steps to disable Kerberos authentication for your Autonomous Database instance.

- Run `DBMS_CLOUD_ADMIN.DISABLE_EXTERNAL_AUTHENTICATION` procedure to disable Kerberos authentication. To run the procedure, you must be logged in as ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.

  ```
  BEGIN
     DBMS_CLOUD_ADMIN.DISABLE_EXTERNAL_AUTHENTICATION;
  END;
  /
  ```

  This disables the Kerberos authentication (or any external authentication scheme specified) for Oracle Autonomous Database.

  See DISABLE_EXTERNAL_AUTHENTICATION Procedure for more information.

Notes for Kerberos Authentication on Autonomous Database

Provides notes on using Kerberos Authentication for Autonomous Database.

- If you enable Kerberos authentication for your Autonomous Database, you can still use password-based database authentication for your database.
- Kerberos authentication is not supported for these tools:
  - Oracle Database API for MongoDB
  - Oracle REST Data Services
  - Oracle Machine Learning
  - APEX
  - Oracle Graph Studio
  - Oracle Database Actions
- You can enable Kerberos authentication to authenticate the ADMIN user. You can use the Reset Password functionality on the Oracle Cloud Infrastructure Console to reset the ADMIN user's password and regain access if a corrupted keytab file causes ADMIN user's authentication to fail.
- The default value for the maximum **clock skew** in Autonomous Database is 300 seconds (5 minutes). You cannot change the default **clock skew** value.

Manage Credentials or Configure Policies and Roles to Access Resources

Describes the methods in which Autonomous Database connects to other resources, either with credentials, or by configuring policies and roles for access.

- Manage Credentials
  You can create credentials, list credentials, or delete credentials in your Autonomous Database.
- Use Vault Secret Credentials
- Configure Policies and Roles to Access Resources
  You may use the following methods to access cloud resources securely without storing user credentials: Oracle Cloud Infrastructure Resource Principals, AWS Amazon Resource Names (ARN)s, Azure service principal, or Google service accounts.

Manage Credentials

You can create credentials, list credentials, or delete credentials in your Autonomous Database.

- Create Credentials to Access Cloud Services
  To access services in the Cloud, such as Cloud Object Store, you first need to create credentials in your Autonomous Database.
- Create Credentials to Access Cloud Services in Cloud Shell
  Use Cloud Shell to run a script that creates OCI Native Credentials and Auth Token in your Autonomous Database.
- List Credentials
  `DBMS_CLOUD` provides the ability to store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`. You can list credentials from the view `ALL_CREDENTIALS`.
- Delete Credentials
  `DBMS_CLOUD` provides the ability to store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`. You can remove credentials with `DBMS_CLOUD.DROP_CREDENTIAL`.

Create Credentials to Access Cloud Services

To access services in the Cloud, such as Cloud Object Store, you first need to create credentials in your Autonomous Database.

1. Create and store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`. For example:

```
SET DEFINE OFF
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'adb_user@example.com',
    password => 'password',
    comments => 'credential to access object store'
  );
END;
/
```

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. After you store the credentials you can then use the same credential name for all data loads.

   For detailed information about the parameters, see CREATE_CREDENTIAL Procedure.

Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

> **Note:**
>
> Some tools like SQL*Plus and SQL Developer use the ampersand character (`&`) as a special character. If you have the ampersand character in your password use the `SET DEFINE OFF` command in those tools as shown in the example to disable the special character and get the credential created properly.

2. With the credential you created in Step 1, you can access Object Store or other cloud resources from Autonomous Database using a procedure such as `DBMS_CLOUD.COPY_DATA`, `DBMS_CLOUD.EXPORT_DATA`, `DBMS_CLOUD_PIPELINE` if you are using a Data Pipeline, or other procedures that require `DBMS_CLOUD` credentials.

Create Credentials to Access Cloud Services in Cloud Shell

Use Cloud Shell to run a script that creates OCI Native Credentials and Auth Token in your Autonomous Database.

To create credentials using the `DBMS_CLOUD` package, see Create Credentials to Access Cloud Services and CREATE_CREDENTIAL Procedure.

- About the Create Credential Cloud Shell Script
  Run the Create Credential Cloud Shell script in the Cloud Shell developer tool to generate credential scripts to run in your Autonomous Database.
- Example: Create OCI Native Credentials
  This example uses Cloud Shell to run the Create Credential script to create OCI Native Credential scripts.
- Example: Create OCI Native Credentials and Run in Autonomous Database
  This example uses the Create Credential script to create an OCI Native Credential script that is run in your Autonomous Database.
- Create Auth Token Used for Swift Credential
  If you want to create an Auth Token/Swift credential, though Oracle recommends the use of OCI Native Credentials, include the `--all` argument when running the script to generate OCI Native Credentials and the Auth Token/Swift credential.

About the Create Credential Cloud Shell Script

Run the Create Credential Cloud Shell script in the Cloud Shell developer tool to generate credential scripts to run in your Autonomous Database.

Use the Create Credential script, `adb-create-cred.sh`, to create new or to reuse existing OCI Native Credentials, including an RSA key pair with a fingerprint. The credentials are provided to the user in the form of scripts, `oci_native_credential.sql` and `oci_native_credential.json`, that can be run in your Autonomous Database. Existing credentials are backed up if new credentials are created.

Use the `adb-create-cred.sh` script to run the generated credential scripts in your Autonomous Database or exit the program and run the scripts in your database with a SQL or JSON compatible tool or utility of your choice. See Example: Create OCI Native Credentials and Example: Create OCI Native Credentials and Run in Autonomous Database for more information and examples.

Optionally, `adb-create-cred.sh` is used to create Auth Token/Swift credentials. Oracle recommends the use of OCI Native Credentials. However, if you want to create an Auth Token/ Swift credential, it is supported by this shell script. See Create Auth Token Used for Swift Credential for details.

> **Note:**
>
> If for you do not have access to the Cloud Shell developer tool, create OCI Native Credentials without using the `adb-create-cred.sh` script. See Create Oracle Cloud Infrastructure Native Credentials for details.

Example: Create OCI Native Credentials

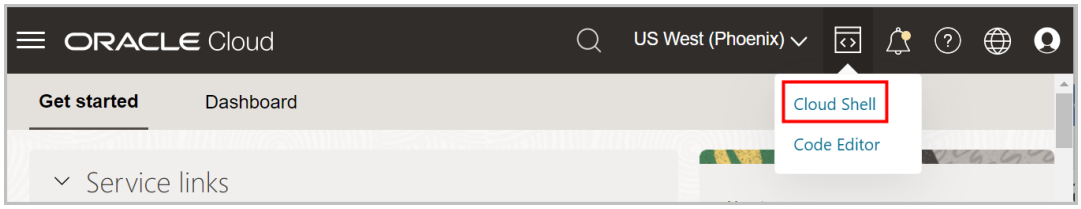
This example uses Cloud Shell to run the Create Credential script to create OCI Native Credential scripts.

Run `adb-create-cred.sh` to generate credential scripts, `oci_native_cred.sql` and `oci_native_cred.json`. The `adb-create-cred.sh` script searches for existing credentials, if found you are asked if you want to reuse them or if you want new credentials created. Depending on your decision, the generated credential scripts include new or existing credentials. Download the OCI Native Credential script or copy it to run it directly in your database using any SQL or JSON tool or utility.

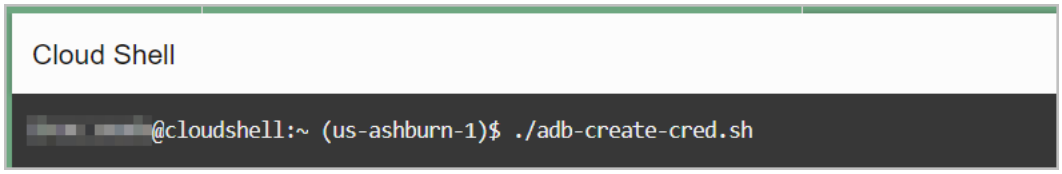
> **Note:**
>
> For a list of arguments supported by the Create Credential script, enter `adb-create-cred.sh --help`.

In this example, OCI Native Credential scripts are generated for your tenancy without running them in your database. See Example: Create OCI Native Credentials and Run in Autonomous Database, for an example of running the OCI Native Credential script in your database.

1. Sign into your Tenancy, select the **Developer Tools** icon and click **Cloud Shell** from the drop down list.



2. Run the **`adb-create-cred.sh`** script.



3. Enter **y** to reuse existing credentials in the generated OCI Native Credential script.

```
################################################################################
# This initial step ensures your OCI API Keys and Fingerprint are setup properly.  #
################################################################################

You have an existing API Key and Fingerprint! Do you want to try to reuse them? (y/n): y
```

If you choose to create new credentials, and decide after overwriting the credentials that you want to use the overwritten credentials, go back to the directory where the backup is and reuse them. The backup file name or folder has a suffix with this syntax: `_bkp_YYYYMMDD_abc`. For example, this backup file was created on June 06, 2024: `_bkp_20240603_woT`.

4. Enter **n** to exit the script without running the OCI Native Credential script in your database.

```
Your API Key was NOT removed, it was found in your OCI Profile. Let's use it!

Your Native Credential SQL file [$HOME/oci_native_credential.sql] was created.

Your Native Credential JSON file [$HOME/oci_native_credential.json] was created.

################################################################################

Proceed to run the Credential Scripts on your Autonomous Database? (y/n):
```

5. The program exits and displays the name of the OCI Native Credential and the command to view it.

```
Exiting...

To view PL/SQL script for creating credentials,

        execute: cat $HOME/oci_native_credential.sql

[redacted]@cloudshell:~ (us-ashburn-1)$
```

Download the OCI Native Credential script or copy it and run it directly in your database using any SQL tool.

Example: Create OCI Native Credentials and Run in Autonomous Database

This example uses the Create Credential script to create an OCI Native Credential script that is run in your Autonomous Database.

Before running the script, consider the following:

- (Optional) Provide the database region, compartment and database name, to avoid having the script search for the database. Although these options are not required, they can save processing time, especially on tenancies spanning a multitude of compartments and Autonomous Databases.

- The script assumes the database is in your home region. If it is in a different region, you must pass in the region name argument, `--region`, when running the script.
- If you don't provide the compartment (--compartment), or database name (--database), the script searches for possible candidates and prompts you to make a selection from lists of possible compartments and databases.
- To list available options, enter `-h` or `--help`.

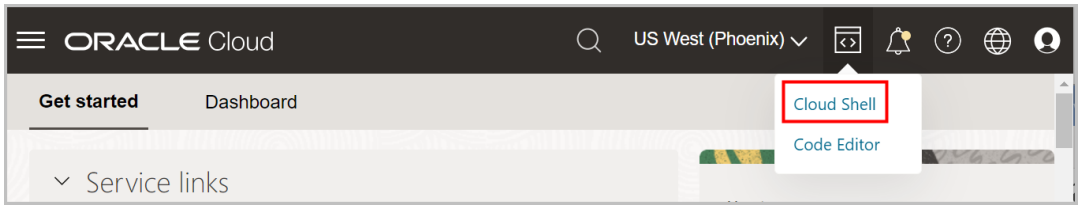
In the following example, the Create OCI Native Credential script, `adb-create-cred.sh`, is used to generate a credential script with existing credentials and run the script in a specified database.

The database region, compartment, and name are passed as arguments to the script. The `--database` option requires both `--region` and `--compartment`. If you specify only `--database` without these, then the option is ignored.

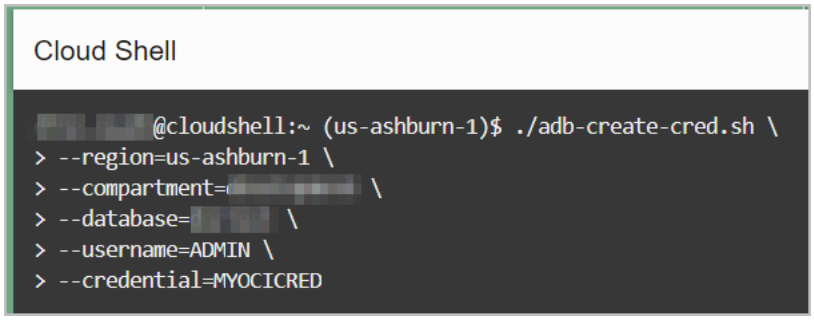
The username used to connect to the database is also provided along with a name for the created credential.

> **Note:**
>
> For a list of arguments supported by the Create Credential script, enter `adb-create-cred.sh --help`.

1. Sign into your Tenancy, select the **Developer Tools** icon and click **Cloud Shell** from the drop down list.



2. Run the **`adb-create-cred.sh`** script, including arguments used to locate the database, connect to the database and a user specified credential name.



3. Existing API keys and fingerprint were found. Enter **y** to reuse them in the generated credential scripts.

```
You specified the Region [us-ashburn-1]

You specified the Compartment [ ]

You named the OCI Native Credential [MYOCICRED]

You specified the Database [ ]

#################################################################################
# This initial step ensures your OCI API Keys and Fingerprint are setup properly.  #
#################################################################################

You have an existing API Key and Fingerprint! Do you want to try to reuse them? (y/n):
```

If you choose to create new credentials, and decide after overwriting the credentials that you want to use the overwritten credentials, you can go back to the directory where the backup is and reuse them. The backup file name or folder has a suffix with this syntax: `_bkp_YYYYMMDD_abc`. For example, this backup file was created on June 06, 2024: `_bkp_20240603_woT`.

4. Enter **y** to run the created credential script, `oci_native_credential.sql`, in the database you specified with the region, compartment and database arguments.

```
Your API Key was NOT removed, it was found in your OCI Profile. Let's use it!

Your Native Credential SQL file [$HOME/oci_native_credential.sql] was created.

Your Native Credential JSON file [$HOME/oci_native_credential.json] was created.

#################################################################################

Proceed to run the Credential Scripts on your Autonomous Database? (y/n): y
```

5. A Wallet File was not found, so it was set up. If a Wallet File was found, you are asked if you want to reuse it or set up a new one. Enter the ***password*** for the ADMIN username you provided. The username and password are used to connect to the database.

```
You do not have a Wallet File. Let's set it up for you.

Fixing path in [ _wallet/sqlnet.ora] to that of your Wallet File.

#################################################################################
# We setup your credentials, generated SQL scripts to add the credentials, we   #
# also setup your Autonomous Database connection, and are now connecting you to #
# your Autonomous Database so you can execute the credential scripts.           #
#################################################################################


You specified the Database Username [ADMIN]

Enter Autonomous Database password:
```

> **Note:**
>
> If you have an Autonomous Database private endpoint in a Virtual Cloud Network, the `adb-create-cred.sh` script generates the SQL and JSON scripts. However, it does not perform the steps required to access the Autonomous Database private end point. Instead, it prompts you to execute `cat ~/oci_native_credential.sql` to copy the SQL into whatever SQL tool you have access to via a Bastion or Jump Host.

6. The login was successful and the `MYOCICRED` credential is created in the specified database. The existing credentials are dropped and the new credentials are created. Enter **n** to not run credentials in another database. The script exists and provides the path for the created `oci_native_credential.sql` script.

```
Login successful.

OCI Native Credential [MYOCICRED] creation successful.

###############################################################################
# Please be aware that it may take up 15 minutes to propagate your credentials to  #
# OCI Services. During this period of time your new credentials will not work.     #
###############################################################################

Proceed to run the Credential Scripts on another Autonomous Database? (y/n): n

Exiting...

To view PL/SQL script for creating credentials,

        execute: cat $HOME/oci_native_credential.sql

[redacted]@cloudshell:~ (us-ashburn-1)$
```

> **Note:**
>
> If the connection to the Autonomous Database is unsuccessful, you can run `cat ~/oci_native_credential.sql`, and copy the SQL to run it directly in your database using any SQL tool.

Create Auth Token Used for Swift Credential

If you want to create an Auth Token/Swift credential, though Oracle recommends the use of OCI Native Credentials, include the `--all` argument when running the script to generate OCI Native Credentials and the Auth Token/Swift credential.

For example:

```
adb-create-cred.sh --all
```

When this flag is used, the script asks if you want to include an Auth Token. By answering **y**, your Auth Token key is generated and uploaded to your OCI profile, and the `oci_auth_token_credential.sql` and `auth_token.tok` scripts are created.

There are 2 options you can use to view the Auth Token key:

- Run the `oci_auth_token_credential.sql` script from the Cloud Shell, to create the Auth Token key in your database. The Auth Token key is the value of the `password` parameter for `DBMS_CLOUD_CREATE_CREDENTIAL`.
- View `auth_token.tok` from the Cloud Shell. Your Auth Token is the value of `token`.

List Credentials

`DBMS_CLOUD` provides the ability to store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`. You can list credentials from the view `ALL_CREDENTIALS`.

For example, to list credentials, run the following command:

```
SELECT credential_name, username, comments FROM all_credentials;

CREDENTIAL_NAME                                                 USERNAME
--------------------------------------------------------------  --------------------
COMMENTS
--------------------------------------------------------------  --------------------
ADB_TOKEN                                                       user_name@example.com
{"comments":"Created via DBMS_CLOUD.create_credential"}
DEF_CRED_NAME                                                   user_name@example.com
{"comments":"Created via DBMS_CLOUD.create_credential"}
```

See ALL_CREDENTIALS for more information.

Delete Credentials

`DBMS_CLOUD` provides the ability to store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`. You can remove credentials with `DBMS_CLOUD.DROP_CREDENTIAL`.

For example, to remove the credential named `DEF_CRED_NAME`, run the following command:

```
BEGIN
   DBMS_CLOUD.DROP_CREDENTIAL('DEF_CRED_NAME');
END;
```

For more information about the `DBMS_CLOUD` procedures and parameters, see DBMS_CLOUD Subprograms and REST APIs.

Use Vault Secret Credentials

Describes using vault secret credentials, where the credentials secret (password) is stored as a secret in a vault. You can then use vault secret credentials to access cloud resources or to access other databases (use anywhere that username/password type credentials are required).
You can create vault secret credentials with secrets stored in any of the supported vaults:

- Oracle Cloud Infrastructure Vault
- Azure Key Vault
- AWS Secrets Manager
- GCP Secret Manager

For example, some possible uses cases for vault secret credentials:

- You can avoid duplicating secrets (passwords) when you access cloud resources from an Autonomous Database instance. In this case, you store secrets in a vault and Autonomous Database accesses the vault. This allows you to rotate secrets without updating the credentials you create to access cloud resources.
- You can use vault secret credentials with database links. In this case, you can create routines that access another database and you don't need to expose passwords in your code.
- You can use vault secret credentials with DBMS_DATAPUMP APIs.

> **Note:**
>
> Operations that use Oracle Data Pump support vault secret credentials (for example `impdp` and `expdp`). Oracle Data Pump support with vault secret credentials is limited to Oracle Cloud Infrastructure Swift URIs and Oracle Cloud Infrastructure Native URIs. See DBMS_CLOUD URI Formats for more information.

Topics

- Use Vault Secret Credentials with Oracle Cloud Infrastructure Vault
  Describes using vault secret credentials, where the secret (password) is stored as a secret in Oracle Cloud Infrastructure Vault.
- Use Vault Secret Credential with Azure Key Vault
  Describes using vault secret credentials, where the credential secret (password) is stored in Azure Key Vault.
- Use Vault Secret Credential with AWS Secrets Manager
  Describes using vault secret credentials, where the credential secret (password) is stored in AWS Secrets Manager.
- Use Vault Secret Credential with GCP Secret Manager
  Describes using vault secret credentials, where the credentials secret (password) are stored as a secret in GCP Secret Manager.
- Refresh Vault Secret Credentials
- Update Vault Secret Credentials

Use Vault Secret Credentials with Oracle Cloud Infrastructure Vault

Describes using vault secret credentials, where the secret (password) is stored as a secret in Oracle Cloud Infrastructure Vault.

You can use vault secret credentials to access cloud resources, to access other databases with database links, or use anywhere that username/password type credentials are required.

- Prerequisites to Create Vault Secret Credentials with Oracle Cloud Infrastructure Vault

- Create Vault Secret Credentials with Oracle Cloud Infrastructure Vault

Prerequisites to Create Vault Secret Credentials with Oracle Cloud Infrastructure Vault

Describes the required prerequisite steps to use vault secret credentials with Oracle Cloud Infrastructure Vault secrets.

To create vault secret credentials where the secret is stored in Oracle Cloud Infrastructure Vault, first perform the required prerequisites.

1. Create a vault and create a secret in the vault with Oracle Cloud Infrastructure Vault.

   For more information, see the instructions for creating a vault and a secret, Managing Vaults and Overview of Key Management.

2. Set up a dynamic group to provide access to the secret in the Oracle Cloud Infrastructure Vault.

   Create a dynamic group for the Autonomous Database instance where you want to create a vault secret credential:

   a. In the Oracle Cloud Infrastructure console click **Identity & Security**.

   b. Under **Identity** click **Domains** and select an identity domain (or create a new identity domain).

   c. Under Identity domain, click **Dynamic groups**.

   d. Click **Create dynamic group** and enter a **Name**, a **Description**, and a rule.

      - **Create Dynamic Group for an existing database:**

        You can specify that an Autonomous Database instance is part of the dynamic group. The dynamic group in the following example includes only the Autonomous Database whose OCID is specified in the `resource.id` parameter:

        ```
        resource.id = 'your_Autonomous_Database_instance_OCID'
        ```

      - **Create a Dynamic Group for a database that has not been provisioned yet:**

        When you are creating the dynamic group before you provision or clone an Autonomous Database instance, the OCID for the new database is not yet available. For this case, create a dynamic group that specifies the resources in a given compartment:

        ```
        resource.compartment.id = 'your_Compartment_OCID'
        ```

   e. Click **Create**.

3. Write policy statements for the dynamic group to enable access to Oracle Cloud Infrastructure resources (secrets).

   a. In the Oracle Cloud Infrastructure Console click **Identity and Security** and click **Policies**.

   b. To write policies for the dynamic group you created in the previous step, click **Create Policy**, and enter a **Name** and a **Description**.

   c. Use the **show manual editor** option of **Policy Builder** to create a policy.

For example, to allow access to the dynamic group to read a specific secret in a compartment:

```
Allow dynamic-group dynamic_group_name to read secret-bundles in
compartment compartment_name
   where target.secret.id='secret_OCID'
```

For example, to allow access to the dynamic group to read all secrets in a compartment:

```
Allow dynamic-group dynamic_group_name to read secret-bundles in
compartment compartment_name
```

See Details for the Vault Service for more information.

d. Select the group or dynamic group and select the location.

e. Click **Create**.

Create Vault Secret Credentials with Oracle Cloud Infrastructure Vault

Describes the steps to use an Oracle Cloud Infrastructure Vault secret with credentials.

This allows you to store a secret in Oracle Cloud Infrastructure Vault and use the secret with the credential you create to access cloud resources or to access other databases.

To create vault secret credentials where the secret is stored in Oracle Cloud Infrastructure Vault:

1. Enable resource principal authentication to provide access to a secret in the Oracle Cloud Infrastructure Vault.

   See Enable Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

2. Create a dynamic group and define policies to allow your Autonomous Database to access secrets in an Oracle Cloud Infrastructure Vault.

   See Prerequisites to Create Vault Secret Credentials with Oracle Cloud Infrastructure Vault for more information.

3. Use `DBMS_CLOUD.CREATE_CREDENTIAL` to create a vault secret credential.

   For example:

```
BEGIN DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name   => 'OCI_SECRET_CRED',
    params            => JSON_OBJECT(
        'username'   value 'SCOTT',
        'secret_id'  value
'ocid1.vaultsecret.oc1.iad.example..aaaaaaaauq5ok5nq3bf2vwetkpqsoa' ));
END;
/
```

   Where:

   - `username`: is the username of the original credential. It can be the username of any type of username/password credential such as the username of an OCI Swift user, the username required to access a database with a database link, and so on.

- `secret_id`: is the vault secret ID. For example, when you store the password *mysecret* in a secret in Oracle Cloud Infrastructure Vault, the `secret_id` value is the vault secret OCID.

To create a vault secret credential you must have `EXECUTE` privilege on the `DBMS_CLOUD` package.

See CREATE_CREDENTIAL Procedure for more information.

4. Use the credential to access a cloud resource.

   For example:

```
SELECT count(*) FROM DBMS_CLOUD.LIST_OBJECTS(
    'OCI_SECRET_CRED',
    'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-
string/b/bucketname/o/' );
```

> **Note:**
>
> Every 12 hours the secret (password) is refreshed from the content in the Oracle Cloud Infrastructure Vault. If you change the secret value in the Oracle Cloud Infrastructure Vault, it can take up to 12 hours for the Autonomous Database instance to pick up the latest secret value.

Run `DBMS_CLOUD.REFRESH_VAULT_CREDENTIAL` to immediately refresh a vault secret credential. This procedure gets the latest version of the vault secret from Oracle Cloud Infrastructure Vault. See REFRESH_VAULT_CREDENTIAL Procedure for more information.

Use Vault Secret Credential with Azure Key Vault

Describes using vault secret credentials, where the credential secret (password) is stored in Azure Key Vault.

You can use vault secret credentials to access cloud resources, to access other databases with database links, or use anywhere that username/password type credentials are required.

- Prerequisites to Create Vault Secret Credential with Azure Key Vault
  Describes the required prerequisites to use vault secret credentials with Azure Key Vault.
- Create Vault Secret Credential with Azure Key Vault
  Describes the steps to use an Azure Key Vault with vault secret credentials.

Prerequisites to Create Vault Secret Credential with Azure Key Vault

Describes the required prerequisites to use vault secret credentials with Azure Key Vault.

To create vault secret credentials where the secret is stored in Azure Key Vault, first perform the required prerequisites.

1. Enable Azure service principal authentication to provide access to the key (password) in the Azure Key Vault.

   See Enable Azure Service Principal for more information.

2. Create an Azure Key Vault and create a secret (password) in the vault.

   See About Azure Key Vault for more information.

3. Set up and enable Azure Service Principal to provide access to the secret in the Azure Key Vault.

   In the Azure portal you must grant read access for a service principal to access the secret.

   a. In the Azure portal, navigate to the "Key Vault" resource that contains the secret you created.

   b. Select "Access policies", then select **Create**.

   c. Under Permissions, select the **Get** permission in the **Secret permissions** section.

   d. Under Principal, enter the name of the service principal in the search field and select the appropriate result.

   e. Click **Next**.

   f. Under Review + create, review the access policy changes and click **Create** to save the access policy.

   g. Back on the Access policies page, verify that your access policy is listed.

   See Assign a Key Vault access policy for more information.

Create Vault Secret Credential with Azure Key Vault

Describes the steps to use an Azure Key Vault with vault secret credentials.

This allows you to store a secret in Azure Key Vault and use the secret with the credentials you create to access cloud resources or to access other databases.

To create vault secret credentials where the secret is stored in Azure Key Vault:

1. Create the Azure Key Vault, the secret, and the access policies to allow your Autonomous Database to access secrets in an Azure Key Vault.

   See Prerequisites to Create Vault Secret Credential with Azure Key Vault for more information.

2. Use `DBMS_CLOUD.CREATE_CREDENTIAL` to create a vault secret credential.

   For example:

```
BEGIN DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name  => 'AZURE_SECRET_CRED',
    params  => JSON_OBJECT(
      'username'          value 'azure_user',
      'secret_id'         value 'sales-secret',
      'azure_vault_name'  value 'azure_keyvault_name' ));
END;
/
```

   Where:

   - `username`: is the username of the original credential. It can be the username of any type of username/password credential.
   - `secret_id`: is the secret name.
   - `azure_vault_name`: is the name of the vault where the secret is located.

   To create a vault secret credential you must have `EXECUTE` privilege on the `DBMS_CLOUD` package.

   See CREATE_CREDENTIAL Procedure for more information.

3. Use the credential to access a cloud resource.

   For example:

```
SELECT count(*) FROM DBMS_CLOUD.LIST_OBJECTS(
     'AZURE_SECRET_CRED',
     'https://adb_user.blob.core.windows.net/adb/' );
```

> **Note:**
>
> Every 12 hours the secret (password) is refreshed from the content in the Azure Key Vault. If you change the secret value in the Azure Key Vault, it can take up to 12 hours for the Autonomous Database instance to pick up the latest secret value.

Run `DBMS_CLOUD.REFRESH_VAULT_CREDENTIAL` to immediately refresh a vault secret credential. This procedure gets the latest version of the vault secret from Azure Key Vault. See REFRESH_VAULT_CREDENTIAL Procedure for more information.

Use Vault Secret Credential with AWS Secrets Manager

Describes using vault secret credentials, where the credential secret (password) is stored in AWS Secrets Manager.

You can use vault secret credentials to access cloud resources, to access other databases with database links, or use anywhere that username/password type credentials are required.

- Prerequisites to Create Vault Secret Credential with AWS Secrets Manager
  Describes the required prerequisites to use vault secret credentials with AWS Secrets Manager.
- Create Vault Secret Credential with AWS Secrets Manager
  Describes the steps to use an AWS Secrets Manager secret with credentials.

Prerequisites to Create Vault Secret Credential with AWS Secrets Manager

Describes the required prerequisites to use vault secret credentials with AWS Secrets Manager.

To create vault secret credentials where the secret is stored in AWS Secrets Manager, first perform the required prerequisites.

1. Create a secret with AWS Secrets Manager and copy the AWS secret ARN.

   See AWS Secrets Manager for more information.

2. Perform AWS management prerequisites to use Amazon Resource Names (ARNs).

   See Perform AWS Management Prerequisites to Use Amazon Resource Names (ARNs) for more information.

3. Enable AWS principal authentication to provide access to the secret in AWS Secrets Manager.

   For example, on the Autonomous Database instance run:

```
BEGIN

DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(
```

```
          provider => 'AWS',
          params =>
             JSON_OBJECT(
                'aws_role_arn' value 'arn:aws:iam::123456:role/
AWS_ROLE_ARN'));
END;
/
```

See ENABLE_PRINCIPAL_AUTH Procedure and About Using Amazon Resource Names (ARNs) to Access AWS Resources for more information.

4. Set up AWS to provide permissions for Amazon Resource Names (ARNs) to access the secret in AWS Secrets Manager.

   On the AWS console, you must grant read access to the secret to the principal authentication credential.

   a. In the AWS console, navigate to the IAM and select **Roles** in Access Management.
   b. Select the role.
   c. In the Permission tab, click **Add permissions** → **create inline policy**.
   d. In the Service section, select **secret manager** as the service.
   e. In the Action section, select **Read** access level.
   f. In the Resources section, click **Add ARN**, specify ARN for secret and click **Add** → click **Review the policy** → give a policy name → click **create policy**.
   g. Back to the Permission tab, verify that the inline policy is attached.

   See Use Amazon Resource Names (ARNs) to Access AWS Resources and Permissions policy examples for AWS Secrets Manager for more information.

Create Vault Secret Credential with AWS Secrets Manager

Describes the steps to use an AWS Secrets Manager secret with credentials.

This allows you to store a secret in AWS Secrets Manager and use the secret with the credentials you create to access cloud resources or to access other databases.

To create vault secret credentials where the secret is stored in AWS Secrets Manager:

1. Create a secret in AWS Secrets Manager and create an inline policy to allow your Autonomous Database to access secrets in AWS Secrets Manager.

   See Prerequisites to Create Vault Secret Credential with AWS Secrets Manager for more information.

2. Use `DBMS_CLOUD.CREATE_CREDENTIAL` to create a vault secret credential to access the AWS Secrets Manager secret.

   For example:

```
BEGIN DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name      => 'AWS_SECRET_CRED',
    params               => JSON_OBJECT(
         'username'   value 'access_key',
         'secret_id'  value 'arn:aws:secretsmanager:region:account-
ID:secret:secret_name' ));
```

```
END;
/
```

Where:

- `username`: is the username of the original credential. It can be the username of any type of username/password credential.
- `secret_id`: is the vault secret AWS ARN.

To create a vault secret credential you must have `EXECUTE` privilege on the `DBMS_CLOUD` package.

See CREATE_CREDENTIAL Procedure for more information.

3. Use the credential to access a cloud resource.

   For example:

```
SELECT count(*) FROM DBMS_CLOUD.LIST_OBJECTS(
     'AWS_SECRET_CRED',
     'https://s3-us-west-2.amazonaws.com/adb/' );
```

> **Note:**
>
> Every 12 hours the secret (password) is refreshed from the content in the AWS Secrets Manager. If you change the secret value in the AWS Secrets Manager, it can take up to 12 hours for the Autonomous Database instance to pick up the latest secret value.

Run `DBMS_CLOUD.REFRESH_VAULT_CREDENTIAL` to immediately refresh a vault secret credential. This procedure gets the latest version of the vault secret from AWS Secrets Manager. See REFRESH_VAULT_CREDENTIAL Procedure for more information.

Use Vault Secret Credential with GCP Secret Manager

Describes using vault secret credentials, where the credentials secret (password) are stored as a secret in GCP Secret Manager.

You can use vault secret credentials to access cloud resources, to access other databases with database links, or use anywhere that username/password type credentials are required.

- Prerequisites to Create Vault Secret Credential with GCP Secret Manager
  Describes the required prerequisites to use vault secret credentials with GCP Secret Manager.
- Create Vault Secret Credential with GCP Secret Manager
  Describes the steps to use an GCP Secret Manager secret to store secrets for use with the credentials you use to access cloud resources.

Prerequisites to Create Vault Secret Credential with GCP Secret Manager

Describes the required prerequisites to use vault secret credentials with GCP Secret Manager.

To create vault secret credentials where the secret is stored in GCP Secret Manager, first perform the required prerequisites.

1. Create a secret in GCP Secret Manager.

   See Secret Manager and Create and access a secret using Secret Manager for more information.

2. Enable Google Service Account authentication to provide access to GCP Secret Manager.

   On the Google Cloud console you must grant read access to the secret to the principal authentication credential.

   a. Go to the **Secret Manager** page in the Google Cloud console.

   b. On the **Secret Manager** page, click the checkbox next to the name of the secret.

   c. If it is not already open, click **Show Info Panel** to open the panel.

   d. In the info panel, click **Add Principal**.

   e. In the **New principals** text area, enter the service account name to add.

   f. In the **Select a role** dropdown, choose **Secret Manager** and then **Secret Manager Secret Accessor**.

   See Enable Google Service Account and Find the GCP Service Account Name and Manage access to secrets for more information.

Create Vault Secret Credential with GCP Secret Manager

Describes the steps to use an GCP Secret Manager secret to store secrets for use with the credentials you use to access cloud resources.

This allows you to store a secret in GCP Secret Manager and use the secret with the credentials you create to access cloud resources or to access other databases.

To create vault secret credentials where the secret is stored in GCP Secret Manager:

1. Create a secret manager secret accessor to allow your Autonomous Database principal to access secrets in GCP Secret Manager.

   See Prerequisites to Create Vault Secret Credential with GCP Secret Manager for more information.

2. Enable Google service account based authentication to provide access to the secret in the GCP Secret Manager.

   See Enable Google Service Account and Find the GCP Service Account Name for more information.

3. Use `DBMS_CLOUD.CREATE_CREDENTIAL` to create a vault secret credential to access the GCP Secret Manager secret.

   For example:

```
BEGIN DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name      => 'GCP_SECRET_CRED',
    params               => JSON_OBJECT(
          'username'   value 'gcp_user1',
          'secret_id'  value 'my-secret',
          'gcp_project_id' value 'my-sample-project-191923' ));
END;
/
```

   Where:

- `username`: is the username of the original credential. It can be the username of any type of username/password credential.
- `secret_id`: is the secret name. When you store the password *mysecret* in the vault, use the secret name as the value of the `secret_id` parameter.
- `gcp_project_id`: is the ID of the project where the secret is located.

See CREATE_CREDENTIAL Procedure for more information.

4. Use the credential to access a cloud resource.

   For example:

```
SELECT count(*) FROM DBMS_CLOUD.LIST_OBJECTS(
           'GCP_SECRET_CRED',
           'https://bucketname.storage.googleapis.com/' );
```

> **Note:**
>
> Every 12 hours the secret (password) is refreshed from the content in the GCP Secret Manager. If you change the secret value in the GCP Secret Manager, it can take up to 12 hours for the Autonomous Database instance to pick up the latest secret value.

Run `DBMS_CLOUD.REFRESH_VAULT_CREDENTIAL` to immediately refresh a vault secret credential. This procedure gets the latest version of the vault secret from GCP Secret Manager. See REFRESH_VAULT_CREDENTIAL Procedure for more information.

Refresh Vault Secret Credentials

Describes how vault secret credentials can be refreshed from the latest value in the vault.

To refresh vault secret credentials, you have two choices:

- Wait for automatic refresh: Every 12 hours the secret (password) is refreshed from the content in the vault. If you change the secret value in the vault, it can take up to 12 hours for the Autonomous Database instance to pick up the latest secret value.
- Run `DBMS_CLOUD.REFRESH_VAULT_CREDENTIAL` to immediately refresh a vault secret credential. This procedure gets the latest version of the vault secret from Oracle Cloud Infrastructure Vault.

  For example:

```
BEGIN
  DBMS_CLOUD.REFRESH_VAULT_CREDENTIAL(
     credential_name => 'VAULT_SECRET_CRED');
END;
/
```

  See REFRESH_VAULT_CREDENTIAL Procedure for more information.

Update Vault Secret Credentials

Describes the steps to update vault secret credential attributes.

To update vault secret credential attributes:

1. Use `DBMS_CLOUD.UPDATE_CREDENTIAL` to update a vault secret credential attribute.

   For example:

   ```
   BEGIN DBMS_CLOUD.UPDATE_CREDENTIAL(
       credential_name   => 'VAULT_SECRET_CRED',
       attribute         => 'secret_id',
       value             => 'new_secret_id'
     );
   END;
   /
   ```

   To update a vault secret credential you must have `EXECUTE` privilege on the `DBMS_CLOUD` package.

   See UPDATE_CREDENTIAL Procedure for more information.

2. Use the updated credential to access a cloud resource.

   For example:

   ```
   SELECT count(*) FROM DBMS_CLOUD.LIST_OBJECTS(
       'VAULT_SECRET_CRED',
       'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-
   string/b/bucketname/o/' );
   ```

Configure Policies and Roles to Access Resources

You may use the following methods to access cloud resources securely without storing user credentials: Oracle Cloud Infrastructure Resource Principals, AWS Amazon Resource Names (ARN)s, Azure service principal, or Google service accounts.

- Use Resource Principal to Access Oracle Cloud Infrastructure Resources
  You can use an Oracle Cloud Infrastructure resource principal with Autonomous Database. You or your tenancy administrator define the Oracle Cloud Infrastructure policies and a dynamic group that allows you to access Oracle Cloud Infrastructure resources with a resource principal. You do not need to create a credential object and Autonomous Database creates and secures the resource principal credentials you use to access the specified Oracle Cloud Infrastructure resources.
- Use Amazon Resource Names (ARNs) to Access AWS Resources
  You can use Amazon Resource Names (ARNs) to access AWS resources with Autonomous Database.
- Use Azure Service Principal to Access Azure Resources
  You can use an Azure service principal with Autonomous Database to access Azure resources without having to create and store your own credential objects in the database.
- Use Google Service Account to Access Google Cloud Platform Resources
  You can use a Google service account to access Google Cloud Platform (GCP) resources from an Autonomous Database instance.

Use Resource Principal to Access Oracle Cloud Infrastructure Resources

You can use an Oracle Cloud Infrastructure resource principal with Autonomous Database. You or your tenancy administrator define the Oracle Cloud Infrastructure policies and a dynamic group that allows you to access Oracle Cloud Infrastructure resources with a resource principal. You do not need to create a credential object and Autonomous Database creates and secures the resource principal credentials you use to access the specified Oracle Cloud Infrastructure resources.

- About Using Resource Principal to Access Oracle Cloud Infrastructure Resources
  You can use a resource principal to authenticate and access Oracle Cloud Infrastructure resources.
- Perform Prerequisites to Use Resource Principal with Autonomous Database
  Prior to making a call to an Oracle Cloud Infrastructure resource using a resource principal, an Oracle Cloud Infrastructure tenancy administrator must create Oracle Cloud Infrastructure policies, dynamic groups, and rules that define the resource principal privileges.
- Enable Resource Principal to Access Oracle Cloud Infrastructure Resources
  Perform the following steps to enable resource principal on Autonomous Database.
- Disable Resource Principal on Autonomous Database
  Shows the steps to disable resource principal for all Autonomous Database users or for a specified user.
- Use Resource Principal with DBMS_CLOUD
  When you specify resource principal credentials with `DBMS_CLOUD` calls, the database authenticates the Oracle Cloud Infrastructure requests for you and the database provides the credentials to access Oracle Cloud Infrastructure resources.

About Using Resource Principal to Access Oracle Cloud Infrastructure Resources

You can use a resource principal to authenticate and access Oracle Cloud Infrastructure resources.

A resource principal consists of a temporary session token and secure credentials that enable the database to authenticate itself to other Oracle Cloud Infrastructure services. Using a resource principal to access services, the token stored with the credentials on Autonomous Database is only valid for the resources to which the dynamic group has been granted access.

To use Resource Principal, you or your tenancy administrator define the Oracle Cloud Infrastructure policies and a dynamic group that allows you to access Oracle Cloud Infrastructure resources with a resource principal. You do not need to create a credential object and Autonomous Database creates and secures the resource principal credentials you use to access the specified Oracle Cloud Infrastructure resources.

For example, while using Autonomous Database you might want to use Oracle Cloud Infrastructure resources to do the following:

- Access data from an Object Storage bucket, perform some operation on the data, and then write the modified data back to the Object Storage bucket.
- Access your vaults, keys, or secrets.
- List work requests or list work request errors.

When you are working with the database, you authenticate and access the database as a database user. An Autonomous Database user does not have an Oracle Cloud Infrastructure Identity and Access Management (IAM) identity, so as an Autonomous Database user you

cannot use your database credentials to access Oracle Cloud Infrastructure services. Without a resource principal you must obtain credentials to access Oracle Cloud Infrastructure resources and create a credential object to access a resource from Autonomous Database.

A resource principal enables resources to be authorized to perform actions on Oracle Cloud Infrastructure services. Each resource has its own identity, and the resource authenticates using the certificates that are added to it. These certificates are automatically created, assigned to resources, and rotated, avoiding the need for you to create and manage your own credentials to access the resource.

Autonomous Database lets you use a resource principal to authenticate to Oracle Cloud Infrastructure APIs using the following interfaces:

- `DBMS_CLOUD` procedures and functions that take a credential argument
- Oracle Cloud Infrastructure PL/SQL SDK APIs

When you authenticate using a resource principal, Autonomous Database provides a secure method to access Oracle Cloud Infrastructure resources.

There are several steps required to set up a resource principal on Autonomous Database:

- You must create an define Oracle Cloud Infrastructure Infrastructure Identity and Access Management (IAM) policies. See Perform Prerequisites to Use Resource Principal with Autonomous Database for more information.
- You must enable the resource principal for the ADMIN user, and optionally enable the resource principal for a database user. See Enable Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

When you authenticate using a resource principal, you do not need to create and manage credentials to access Oracle Cloud Infrastructure resources. Autonomous Database makes the resource principal available to you and secures the resource principal for you.

Perform Prerequisites to Use Resource Principal with Autonomous Database

Prior to making a call to an Oracle Cloud Infrastructure resource using a resource principal, an Oracle Cloud Infrastructure tenancy administrator must create Oracle Cloud Infrastructure policies, dynamic groups, and rules that define the resource principal privileges.

Perform the following steps before you use a resource principal with Autonomous Database:

1. Create an Oracle Cloud Infrastructure dynamic group.
   a. In the Oracle Cloud Infrastructure console click **Identity and Security** and click **Dynamic Groups**
   b. Click **Create Dynamic Group** and enter a **Name**, a **Description**, and a rule or use the **Rule Builder** to add a rule.
   c. Click **Create**.

   Resources that meet the rule criteria are members of the dynamic group. When you define a rule for a dynamic group, consider what resource is going to be given access to other resources.

   For example, consider the following examples:

   - Allow a specific Autonomous Database instance to access a resource.

     The Autonomous Database is specified in the `resource.id` parameter with an OCID:

     ```
     resource.id = 'your_Autonomous_Database_instance_OCID'
     ```

You need to add a Domain Name when the tenancy uses an Identity Domain. For example:

```
resource.id = 'identity_domain_name/
your_Autonomous_Database_instance_OCID'
```

- Allow all Autonomous Databases in a compartment.

  The Autonomous Databases are specified in the `resource.type` parameter and the compartment is identified by a specified OCID in the `resource.compartment.id` parameter:

  ```
  ALL {resource.type = 'autonomousdatabase', resource.compartment.id =
  'your_Compartment_OCID'}
  ```

- Allow all resources in the compartment

  The resource type identified by the OCID, specified in the `resource.compartment.id` parameter:

  ```
  ALL
          {resource.compartment.id='your_Compartment_OCID'}
  ```

See Managing Dynamic Groups for more information on creating a dynamic group and creating rules to add resources to the group.

2. Write policy statements for the dynamic group to enable access to Oracle Cloud Infrastructure resources.
   a. In the Oracle Cloud Infrastructure console click **Identity and Security** and click **Policies**.
   b. To write policies for a dynamic group, click **Create Policy**, and enter a **Name** and a **Description**.
   c. Use the **Policy Builder** to create a policy.

      For example to create a policy to allow access to Oracle Cloud Infrastructure Object Store to manage buckets and objects in the Object Store in a tenancy:

      ```
      Allow dynamic-group Example5 to manage buckets in tenancy
      Allow dynamic-group Example5 to manage objects in tenancy
      ```

      You need to add a Domain Name when the tenancy uses an Identity Domain. For example:

      ```
      Allow dynamic-group identity_domain_name/Example5 to manage buckets in
      tenancy
      Allow dynamic-group identity_domain_name/Example5 to manage objects in
      tenancy
      ```

   d. Click **Create**.

> **Note:**
>
> The resource principal token is cached for two hours. Therefore, if you change the policy or the dynamic group, you have to wait for two hours to see the effect of your changes.

See How Policies Work and Managing Policies for more information on policies.

Enable Resource Principal to Access Oracle Cloud Infrastructure Resources

Perform the following steps to enable resource principal on Autonomous Database.

As a prerequisite, configure dynamic groups and policies. See Perform Prerequisites to Use Resource Principal with Autonomous Database for more information.

To enable a resource principal on Autonomous Database:

1. As the ADMIN user, enable resource principal for the Autonomous Database instance.

   For example:

   ```
   EXEC DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL();

   PL/SQL procedure successfully completed.
   ```

   See ENABLE_RESOURCE_PRINCIPAL Procedure for more information.

   This creates the credential `OCI$RESOURCE_PRINCIPAL`.

2. (Optional) This step is only required if you want to grant access to the resource principal credential to a database user other than the ADMIN user. As the ADMIN user, enable resource principal for a specified database user.

   For example:

   ```
   EXEC DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL(username => 'adb_user');

   PL/SQL procedure successfully completed.
   ```

   This grants the user `adb_user` access to the credential `OCI$RESOURCE_PRINCIPAL`.

   If you want the specified user to have privileges to enable resource principal for other users, set the `grant_option` parameter to `TRUE`.

   For example:

   ```
   BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL(
        username => 'adb_user',
        grant_option => TRUE);
   END;
   /
   ```

After you run this command, `adb_user` can enable resource principal for another user. For example, if you connect as `adb_user`, you can run the following command:

```
EXEC DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL(username => 'adb_user2');
```

See ENABLE_RESOURCE_PRINCIPAL Procedure for more information.

3. Verify that the resource principal credential is enabled.

   For example, as the ADMIN user query the view `DBA_CREDENTIALS`:

```
SELECT owner, credential_name FROM dba_credentials
        WHERE credential_name = 'OCI$RESOURCE_PRINCIPAL' AND owner =
'ADMIN';

OWNER  CREDENTIAL_NAME
-----  ----------------------
ADMIN  OCI$RESOURCE_PRINCIPAL
```

   For example, as a non-ADMIN user query the view `ALL_TAB_PRIVS`:

```
SELECT grantee, table_name, grantor FROM ALL_TAB_PRIVS
   WHERE grantee = 'ADB_USER'
        AND table_name = 'OCI$RESOURCE_PRINCIPAL'
        AND table_schema = 'ADMIN';

GRANTEE   TABLE_NAME               GRANTOR
--------- -----------------------  -------------
ADB_USER  OCI$RESOURCE_PRINCIPAL   ADMIN
```

Enabling the resource principal on an Autonomous Database instance is one-time operation. You do not need to enable the resource principal again, unless you run `DBMS_CLOUD_ADMIN.DISABLE_RESOURCE_PRINCIPAL` to disable the resource principal.

Disable Resource Principal on Autonomous Database

Shows the steps to disable resource principal for all Autonomous Database users or for a specified user.

1. To disable resource principal for all users, as the ADMIN user, run the following command:

```
EXEC DBMS_CLOUD_ADMIN.DISABLE_RESOURCE_PRINCIPAL();
```

   This removes the credential `OCI$RESOURCE_PRINCIPAL`.

2. Verify that the resource principal credential is disabled.

   For example:

```
SELECT owner, credential_name FROM dba_credentials
        WHERE credential_name = 'OCI$RESOURCE_PRINCIPAL' AND owner =
'ADMIN';

No rows selected
```

To remove access to the resource principal credential for a specified database user, include the `username` parameter. This denies the specified user access to the `OCI$RESOURCE_PRINCIPAL` credential.

For example:

```
EXEC DBMS_CLOUD_ADMIN.DISABLE_RESOURCE_PRINCIPAL(username => 'ADB_USER');
```

See DISABLE_RESOURCE_PRINCIPAL Procedure for more information.

Use Resource Principal with DBMS_CLOUD

When you specify resource principal credentials with `DBMS_CLOUD` calls, the database authenticates the Oracle Cloud Infrastructure requests for you and the database provides the credentials to access Oracle Cloud Infrastructure resources.

If you have not already done so, perform the required prerequisite steps:

- Access to Oracle Cloud Infrastructure resources depends on the dynamic group rules and the policies you set in Oracle Cloud Infrastructure policies and dynamic groups. See Perform Prerequisites to Use Resource Principal with Autonomous Database for more information.
- After you define the dynamic group and policies, enable the ADMIN schema or another schema to use a resource principal. See Enable Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

To use a `DBMS_CLOUD` procedure with resource principal credentials:

- Use a `DBMS_CLOUD` procedure or function and specify `OCI$RESOURCE_PRINCIPAL` as the credential name.

  For example, you can access Oracle Cloud Infrastructure Object Storage using a resource principal:

```
CREATE TABLE CHANNELS
   (channel_id CHAR(1),
    channel_desc VARCHAR2(20),
    channel_class VARCHAR2(20)
   );
/

BEGIN
 DBMS_CLOUD.COPY_DATA(
    table_name =>'CHANNELS',
    credential_name =>'OCI$RESOURCE_PRINCIPAL',
    file_uri_list =>'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/channels.txt',
    format => json_object('delimiter' value ',')
 );
END;
/
```

If you compare the steps required to access Object Storage as shown in Create Credentials and Copy Data into an Existing Table, notice that Step 1, creating credentials is not required when you use resource principal because you are using the system defined `OCI$RESOURCE_PRINCIPAL` credential.

Use Amazon Resource Names (ARNs) to Access AWS Resources

You can use Amazon Resource Names (ARNs) to access AWS resources with Autonomous Database.

- About Using Amazon Resource Names (ARNs) to Access AWS Resources
  When you use ARN role based authentication with Autonomous Database, you can securely access AWS resources without creating and saving credentials based on long-term AWS IAM access keys.
- Perform AWS Management Prerequisites to Use Amazon Resource Names (ARNs)
  Using the AWS Management Console or using the APIs, create an AWS user, role, policies, and trust relationship. You perform these steps before you use with `DBMS_CLOUD.CREATE_CREDENTIAL` to create a credential with an ARN parameter on Autonomous Database.
- Perform Autonomous Database Prerequisites to Use Amazon ARNs
  Prior to using an AWS resource with `DBMS_CLOUD.CREATE_CREDENTIAL` with an ARN parameter, the ADMIN user must enable ARN on the Autonomous Database instance.
- Create Credentials with ARN Parameters to Access AWS Resources
  After ARN usage is enabled for the Autonomous Database instance and the ARN is configured by the AWS administrator, on Autonomous Database you can create a credential object with ARN parameters.
- Update Credentials with ARN Parameters for AWS Resources
  The ARN credentials you use on Autonomous Database work with the AWS token service that enables you to use temporary role based credentials to access to AWS resources from Autonomous Database.

About Using Amazon Resource Names (ARNs) to Access AWS Resources

When you use ARN role based authentication with Autonomous Database, you can securely access AWS resources without creating and saving credentials based on long-term AWS IAM access keys.

For example, you may want to load data from an AWS S3 bucket into your Autonomous Database, perform some operation on the data, and then write the modified data back to the S3 bucket. You can do this without using an ARN if you have AWS user credentials to access the S3 bucket. However, using role-based ARNs to access AWS resources from Autonomous Database has the following benefits:

- You can create role-based access, with different policies for different users or schemas that need access to AWS resources from an Autonomous Database instance. This allows you to set a policy to limit access to AWS resources by role. For example, setting a policy limiting to read-only access, by role, to an S3 bucket.
- ARN based credentials provide better security as you do not need to provide long-term AWS user credentials in code to access AWS resources. Autonomous Database manages the temporary credentials generated from the AWS Assume Role Operation.

Steps to Configure ARN Usage with Autonomous Database

Before creating a credential using an ARN in Autonomous Database, in AWS, your account administrator must define a policy that allows you to access AWS resources, such as an S3 bucket. By default, ARN credential services are not enabled on Autonomous Database. The ADMIN user enables ARN credentials for the necessary user which allows them to create and use ARN credentials on the Autonomous Database instance.

In AWS, the role ARN is the identifier for the provided access and can be viewed on the AWS console. For added security, when the AWS administrator configures the role, policies, and trust relationship for the AWS account, they must also configure an External ID in the role's trust relationship.

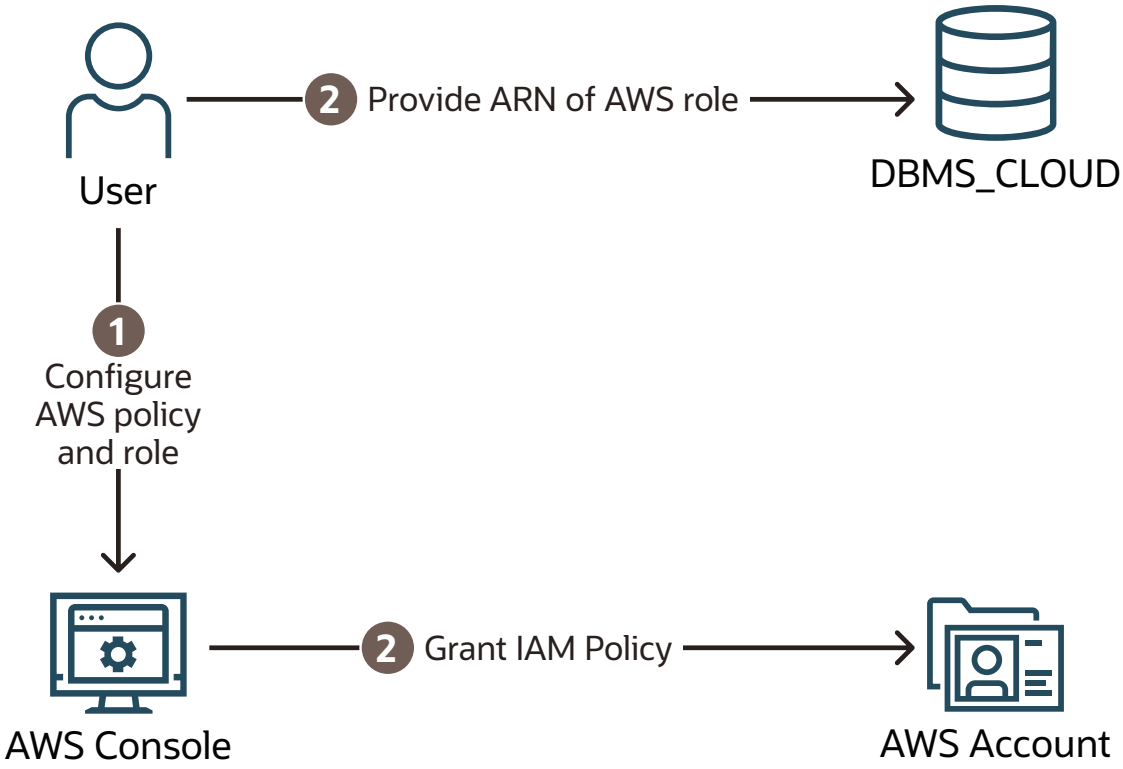
The External ID provides additional protection for assuming roles. The AWS administrator configures configure the External ID as one of the following, based on the Autonomous Database instance:

- The compartment OCID
- The database OCID
- The tenancy OCID

On AWS, the role can only be assumed by trusted users that are identified by the External ID included in the request URL, where the supplied External ID in the request matches the External ID configured in the role's trust relationship.

> **Note:**
>
> Setting the External ID is required for security.

The following figure outlines the configuration steps:



Steps to Use ARNs with DBMS_CLOUD

Each AWS resource has its own identity, and the resource authenticates with the Autonomous Database instance using a `DBMS_CLOUD` credential that you create with parameters that identify the ARN. Autonomous Database creates and secures the principal credentials you use to access AWS resources.

To create a credential with ARN parameters to access AWS resources:

1. Perform prerequisite steps in the AWS Account: In your AWS account, from the AWS Management Console or using the CLI, create the roles and policies for the ARN that you use with Autonomous Database and update the trust relationship for the role. The Oracle user ARN is configured when the trust relationship for the role is updated.

   See Perform AWS Management Prerequisites to Use Amazon Resource Names (ARNs) for more information.

2. Perform prerequisite steps on Autonomous Database: On Autonomous Database you must enable the ADMIN user or another user to use credentials with ARN parameters to access AWS resources.

   See Perform Autonomous Database Prerequisites to Use Amazon ARNs for more information.

3. Create credentials with `DBMS_CLOUD.CREATE_CREDENTIAL` and supply the parameters that identify an AWS role. Using the credential object, Autonomous Database can access AWS resources as specified in the policies defined for the role in the AWS account.

   See Create Credentials with ARN Parameters to Access AWS Resources for details on these steps.

4. Use the credential object you created in the previous step with a `DBMS_CLOUD` procedure or function that takes a credential parameter, such as `DBMS_CLOUD.COPY_DATA` or `DBMS_CLOUD.LIST_OBJECTS`.

Perform AWS Management Prerequisites to Use Amazon Resource Names (ARNs)

Using the AWS Management Console or using the APIs, create an AWS user, role, policies, and trust relationship. You perform these steps before you use with `DBMS_CLOUD.CREATE_CREDENTIAL` to create a credential with an ARN parameter on Autonomous Database.

To use an ARN to access AWS resources your AWS administrator defines the policies and a principal that allows you to access AWS resources. For example, while using Autonomous Database you might want to access data from an S3 bucket, perform some operation on the data, and then write the modified data back to the S3 bucket.

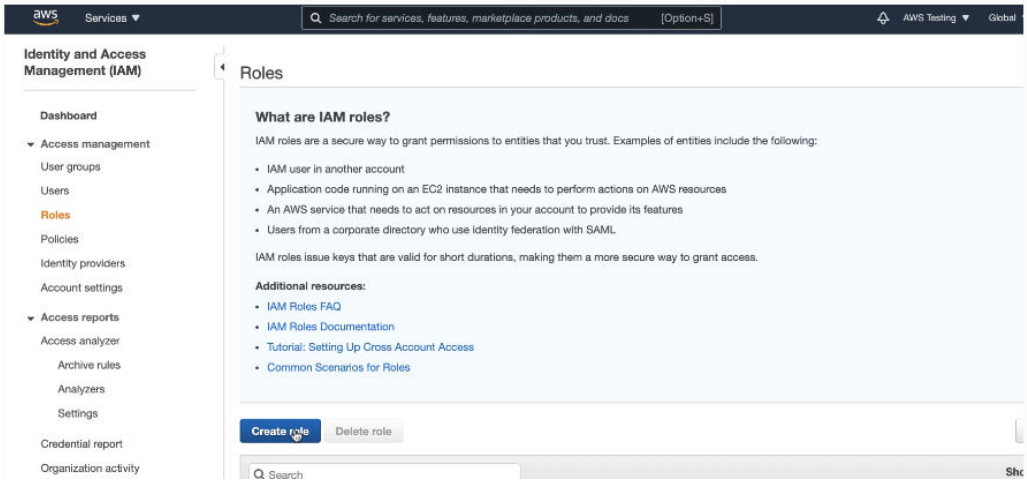
> **Note:**
>
> Depending on your existing AWS configuration and the External ID you use, you do not need to create a new role and policy for each Autonomous Database instance. If you already have an AWS role containing the necessary policy to access a resource, for example to access S3 cloud storage, you can modify the trust relationship to include the details in Step 3. Likewise, if you already have a role with the necessary trust relationship, you can use that role to access all of your databases in an OCI compartment or tenancy if you use an external ID that specifies the compartment OCID or tenancy OCID.

From the AWS Management Console or using the APIs, an AWS administrator performs the following steps:

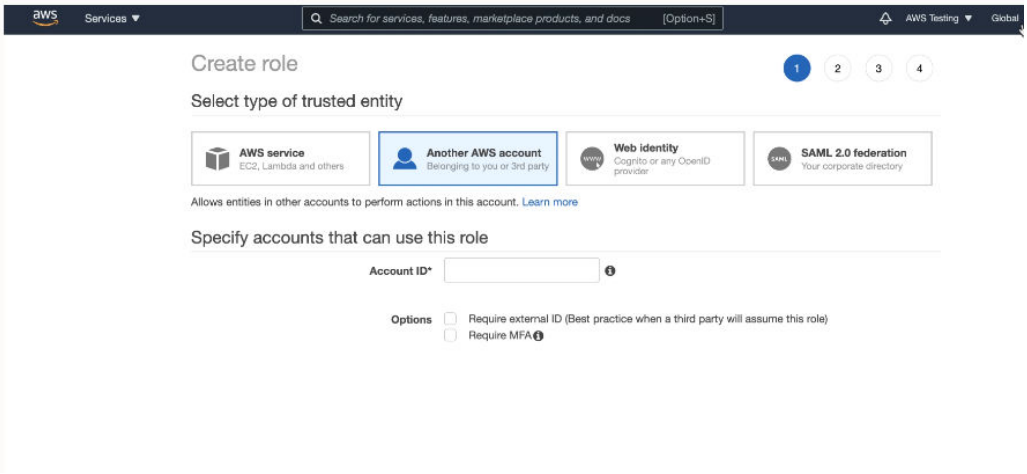
1. Create a policy. In the policy you specify permissions for accessing AWS resources such as S3 buckets.

   See Creating an IAM policy to access Amazon S3 resources for more information.

2. Create a role and attach the policy to the role.
   a. Access the AWS Management Console and choose Identity and Access Management (IAM).
   b. Click **Create role**.



   c. Select **Another AWS account**.



   d. Enter your **Account ID**.

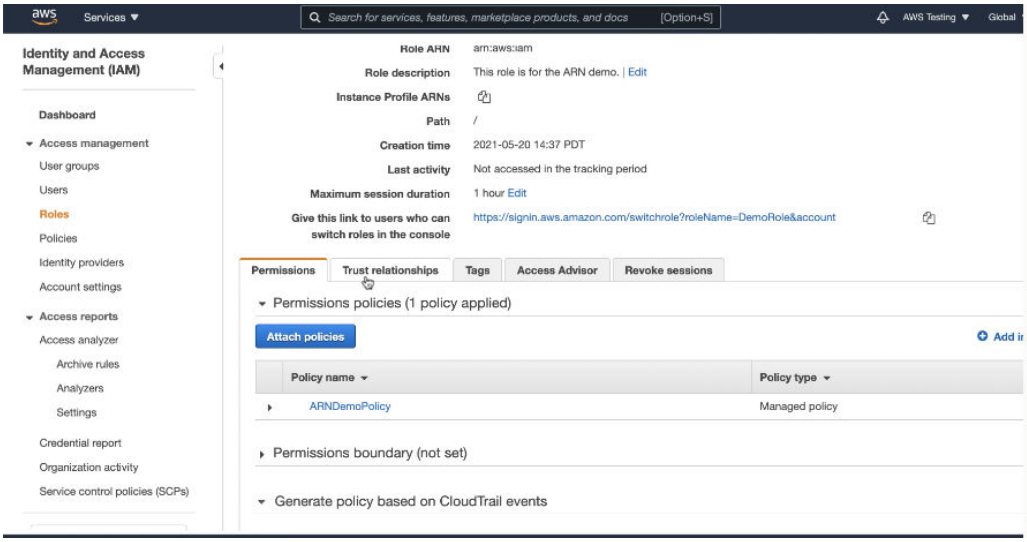
      You use this as a temporary value. Later you replace this with the Account ID you use to access AWS resources.

   e. In the **Options** area select **Require external ID** and enter a temporary external ID, such as `0000`. Later you replace this external ID with a valid value.
   f. Click **Next Permissions** to attach the Policies you created in Step 1 or other policies you want to apply to the role.
   g. Click **Next Tags** and apply or create tags as needed for the role.
   h. Click **Next Review** and add a **Role Name** and **Role Description**.
   i. Click **Create Role**.

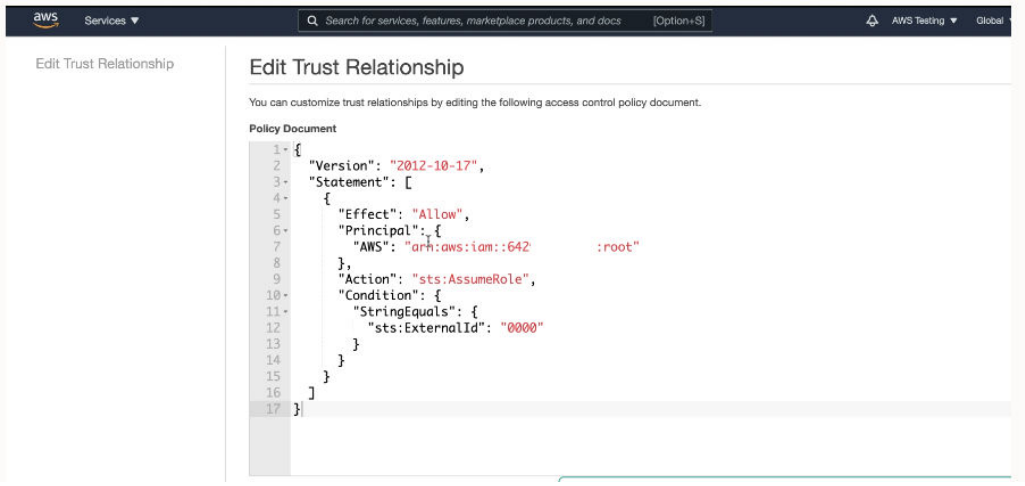
You use the role's ARN with `DBMS_CLOUD.CREATE_CREDENTIAL` to create credential objects with ARN parameters to access AWS resources.

See Creating a role to delegate permissions to an IAM user for more information.

3. Specify a Trust Relationship for the role.
   a. From the Roles list, under **Role name**, select the role you created.
   b. On the roles Summary page for the selected role, select the **Trust relationships** tab.



   c. In the trust relationship, click **Edit trust relationship**.



   d. Edit the trust relationship to specify the Principal parameter **AWS**.

      This AWS user ARN is available in the CLOUD_INTEGRATIONS view. See Perform Autonomous Database Prerequisites to Use Amazon ARNs for more information.

   e. Edit the trust relationship to specify the External ID.

      On Autonomous Database when you create an AWS ARN credential with `DBMS_CLOUD.CREATE_CREDENTIAL` or when you enable AWS ARN with `DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH`, by default the `external_id_type` parameter value is `database_ocid`. Optionally you can set the `external_id_type` value to one of the supported values: `database_ocid`, `compartment_ocid`, or `tenant_ocid`.

      When you use the database OCID as the External ID, the policy's trust relationship only trusts the Autonomous Database instance specified with the OCID. If you use a compartment OCID, the policy's trust relationship trusts all the Autonomous Database instances in the compartment and you can use the same role ARN to grant access to AWS resources to any Autonomous Database in the specified compartment. Likewise,

if you use the tenancy OCID, you can use the same role ARN to grant access to AWS resources to any Autonomous Database in the specified tenancy.

Previously in Step 2 you set the trust relationship External ID to the temporary value `0000`.

On AWS you configure the trust relationship External ID value to match one of the following:

- When the `external_id_type` type is `database_ocid`, on AWS you configure the role's trust relationship External ID to be the Database OCID.

  The Database OCID is available by running the following query:

  ```
  SELECT cloud_identity FROM v$pdbs;
  ```

  See Obtain Tenancy Details for more information.

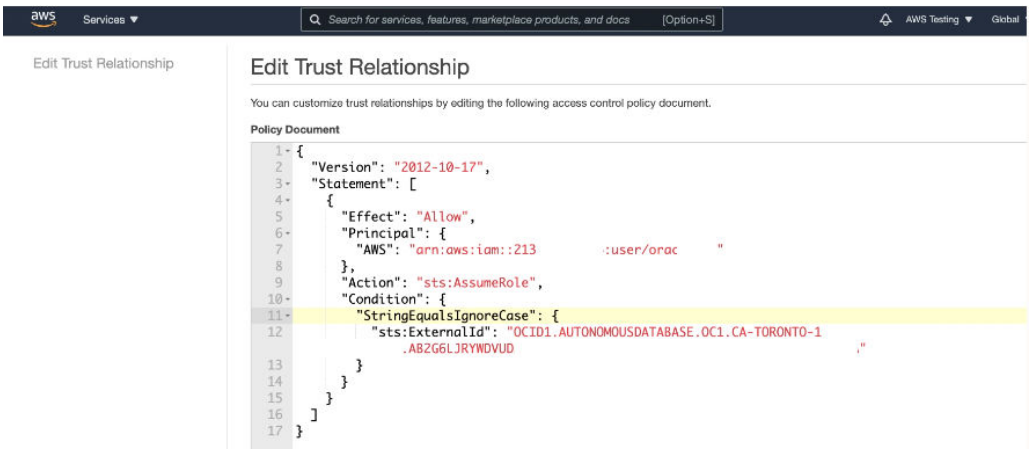
- When the `external_id_type` type is `compartment_ocid`, on AWS you configure the role's trust relationship External ID to be the Compartment OCID.

  The Compartment OCID is available on the Compartment details page from the Oracle Cloud Infrastructure Console. To find the Compartment details page, from the Oracle Cloud Infrastructure left navigation menu click **Identity & and Security** and then select **Compartments**. Select the compartment that contains the Autonomous Database instance to see the Compartment ID.

- When the `external_id_type` type is `tenant_ocid`, on AWS you configure the role's trust relationship External ID to be the Tenancy OCID.

  The Tenancy OCID is available on the Tenancy details page from the Oracle Cloud Infrastructure Console. To find the Tenancy details page, from the Oracle Cloud Infrastructure left navigation menu click **Governance & Administration** and then select **Tenancy Details**. The Tenancy Information tab shows the Tenancy OCID.

- When you set the value for ExternalID, by default the OCID value must be in upper case. If you want to supply the OCID in lower case, set the condition "`StringEqualsIgnoreCase`" instead of "`StringEquals`" in the JSON when you edit the trust relationship.



See How to use trust policies with IAM role and How to use an external ID when granting access to your AWS resources to a third party for more information.

After the ARN role configuration is finished, you can enable ARN on the instance. See Perform Autonomous Database Prerequisites to Use Amazon ARNs for more information.

Perform Autonomous Database Prerequisites to Use Amazon ARNs

Prior to using an AWS resource with `DBMS_CLOUD.CREATE_CREDENTIAL` with an ARN parameter, the ADMIN user must enable ARN on the Autonomous Database instance.

By default, ARN credential services are not enabled on Autonomous Database. The ADMIN user runs the procedure `DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH` to enable the ADMIN user or other users to create credentials with ARN parameters.

1. Enable the use of ARN credentials on the Autonomous Database instance.

```
BEGIN
    DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(
        username => 'adb_user',
        params => JSON_OBJECT(
                'aws_role_arn' value 'arn:aws:iam::123456:role/
AWS_ROLE_ARN'));
END;
/
```

If you want the specified user to have privileges to enable ARN credentials for other users, set the `params` parameter `grant_option` to `TRUE`.

For example:

```
BEGIN
DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(
     username => 'adb_user',
     params => JSON_OBJECT(
                    'aws_role_arn' value 'arn:aws:iam::123456:role/
AWS_ROLE_ARN',
                    'grant_option' value TRUE ));
END;
/
```

After you run this command, `adb_user` has privileges to enable ARN credentials for other users.

For example, if you connect as `adb_user`, you can run the following command:

```
BEGIN
    DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(
        username => 'adb_user2');
END;
/
```

See ENABLE_PRINCIPAL_AUTH Procedure for more information.

2. Query the `CLOUD_INTEGRATIONS` view to obtain Oracle's AWS user ARN.

```
SELECT param_value FROM CLOUD_INTEGRATIONS
        WHERE param_name = 'aws_user_arn';

PARAM_VALUE
```

```
-----------------------------------------
arn:aws:iam::account-ID:user/username
```

The view `CLOUD_INTEGRATIONS` is available to the ADMIN user or to a user with `DWROLE` privileges.

The AWS administrator uses the `aws_user_arn` value when configuring the AWS role's trust relationship with the role and policies on the AWS system. Providing this value grants permission on the AWS side for `DBMS_CLOUD` to access AWS resources.

After you enable ARN on the Autonomous Database instance by running `DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH`, the credential named `AWS$ARN` is available to use with any `DBMS_CLOUD` API that takes a credential as the input. Except for the credential named `AWS$ARN`, you can also create additional credentials with ARN parameters to access AWS resources. See Create Credentials with ARN Parameters to Access AWS Resources for more information.

Create Credentials with ARN Parameters to Access AWS Resources

After ARN usage is enabled for the Autonomous Database instance and the ARN is configured by the AWS administrator, on Autonomous Database you can create a credential object with ARN parameters.

Autonomous Database creates and secures the principal credentials you use to access the Amazon resources when you supply the credential object with `DBMS_CLOUD` procedures and functions.

To use Amazon resources with Autonomous Database, do the following:

1. Create credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL` with the `params` parameter to specify the ARN value. For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_ARN',
    params =>
        JSON_OBJECT('aws_role_arn' value 'arn:aws:iam::123456:role/
AWS_ROLE_ARN',
                    'external_id_type' value 'database_ocid')
  );
END;
/
```

This operation creates the credentials in the database in an encrypted format. You can use any name for the credential name.

For detailed information about the parameters, see CREATE_CREDENTIAL Procedure.

2. Use a `DBMS_CLOUD` procedure to access an Amazon resource with the ARN credentials.

For example, use `DBMS_CLOUD.LIST_OBJECTS`.

```
SELECT object_name FROM DBMS_CLOUD.LIST_OBJECTS(
           credential_name => 'DEF_CRED_ARN',
           location_uri    => 'https://my-bucket.s3.us-
west-2.amazonaws.com/');
```

Update Credentials with ARN Parameters for AWS Resources

The ARN credentials you use on Autonomous Database work with the AWS token service that enables you to use temporary role based credentials to access to AWS resources from Autonomous Database.

When an AWS Administrator revokes the policies, roles, or trust relationship, you need to either update the credentials or create new credentials to access the AWS resources.

Perform the following steps to update credentials:

1. Use `DBMS_CLOUD.UPDATE_CREDENTIAL` to update an ARN based credential to supply a new ARN value.

```
BEGIN
  DBMS_CLOUD.UPDATE_CREDENTIAL(
     credential_name => 'DEF_CRED_ARN',
     attribute => 'aws_role_arn',
     value => 'new_ARN_value');
END;
/
```

   This updates the `aws_role_arn` attribute to the new value *`new_ARN_value`* for the credential named DEF_CRED_ARN.

2. Use `DBMS_CLOUD.UPDATE_CREDENTIAL` to update an ARN based credential to update the attribute `external_id_type` value.

```
BEGIN
  DBMS_CLOUD.UPDATE_CREDENTIAL(
     credential_name => 'DEF_CRED_ARN',
     attribute => 'external_id_type',
     value => 'compartment_ocid');
END;
/
```

   This updates the `external_id_type` attribute value to the value `compartment_ocid`.

See UPDATE_CREDENTIAL Procedure and CREATE_CREDENTIAL Procedure for more information.

Use Azure Service Principal to Access Azure Resources

You can use an Azure service principal with Autonomous Database to access Azure resources without having to create and store your own credential objects in the database.

- Enable Azure Service Principal
  Enable Azure service principal authentication to allow Autonomous Database to access Azure services without providing long-term credentials.
- Provide Azure Application Consent and Assign Roles
  To access Azure resources from Autonomous Database with Azure service principal authentication you must consent the Azure application and assign roles to allow access to your Azure resources.

- Use Azure Service Principal with DBMS_CLOUD
  When you make `DBMS_CLOUD` calls to access Azure resources and specify the credential name as `AZURE$PA`, the authentication on the Azure side happens using the Azure service principal.
- Disable Azure Service Principal
  To disable access to Azure resources from Autonomous Database with Azure service principal, use `DBMS_CLOUD_ADMIN.DISABLE_PRINCIPAL_AUTH`.
- Notes for Azure Service Principal
  Notes for using Azure service principal.

Enable Azure Service Principal

Enable Azure service principal authentication to allow Autonomous Database to access Azure services without providing long-term credentials.

> **Note:**
>
> To use Autonomous Database with Azure service principal authentication you need a Microsoft Azure account. See Microsoft Azure for details.

To enable Azure service principal authentication on Autonomous Database:

1. Obtain your Microsoft Azure Active Directory tenant ID.

   See How to find your Azure Active Directory tenant ID for more information.

2. Enable Azure service principal with `DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH`.

   For example:

```
BEGIN
    DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(
        provider => 'AZURE',
        username => 'adb_user',
        params   => JSON_OBJECT('azure_tenantid' value 'azure_tenantID'));
END;
/
```

   This enables Azure service principal authentication and creates an Azure application on Autonomous Database.

   If you want the specified user to have privileges to enable Azure service principal for other users, set the `params` parameter `grant_option` to `TRUE`.

   For example:

```
BEGIN
    DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(
        provider => 'AZURE',
        username => 'adb_user',
        params   => JSON_OBJECT('grant_option' value TRUE,
                                'azure_tenantid' value 'azure_tenantID'));
END;
/
```

After you run this command, `adb_user` can enable Azure service principal for another user. For example, if you connect as `adb_user`, you can run the following command:

```
BEGIN
    DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(
        provider => 'AZURE',
        username => 'adb_user2');
END;
/
```

See ENABLE_PRINCIPAL_AUTH Procedure for more information.

Provide Azure Application Consent and Assign Roles

To access Azure resources from Autonomous Database with Azure service principal authentication you must consent the Azure application and assign roles to allow access to your Azure resources.

To provide Azure application consent and assign roles, perform the following steps:

1. On Autonomous Database query `CLOUD_INTEGRATIONS`.

   For example:

   ```
   SELECT param_name, param_value FROM CLOUD_INTEGRATIONS;

   PARAM_NAME        PARAM_VALUE
   ---------------
   --------------------------------------------------------------------------------
   ----------------------------------------------------------------
   azure_tenantid    29981886-6fb3-44e3-82ab-d870b0e8e7eb
   azure_consent_url https://login.microsoftonline.com/f8cdef31-91255a/oauth2/
   v2.0/authorize?
   client_id=d66f1b5-1250d5445c0b&response_type=code&scope=User.read
   azure_app_name
   ADBS_APP_OCID1.AUTONOMOUSDATABASE.REGION1.SEA.ANZWKLJSZLYNB3AAWLYL3JVC4ICEX
   LB3ZG6WTCX735JSSY2NRHOBU4DZOOVA
   ```

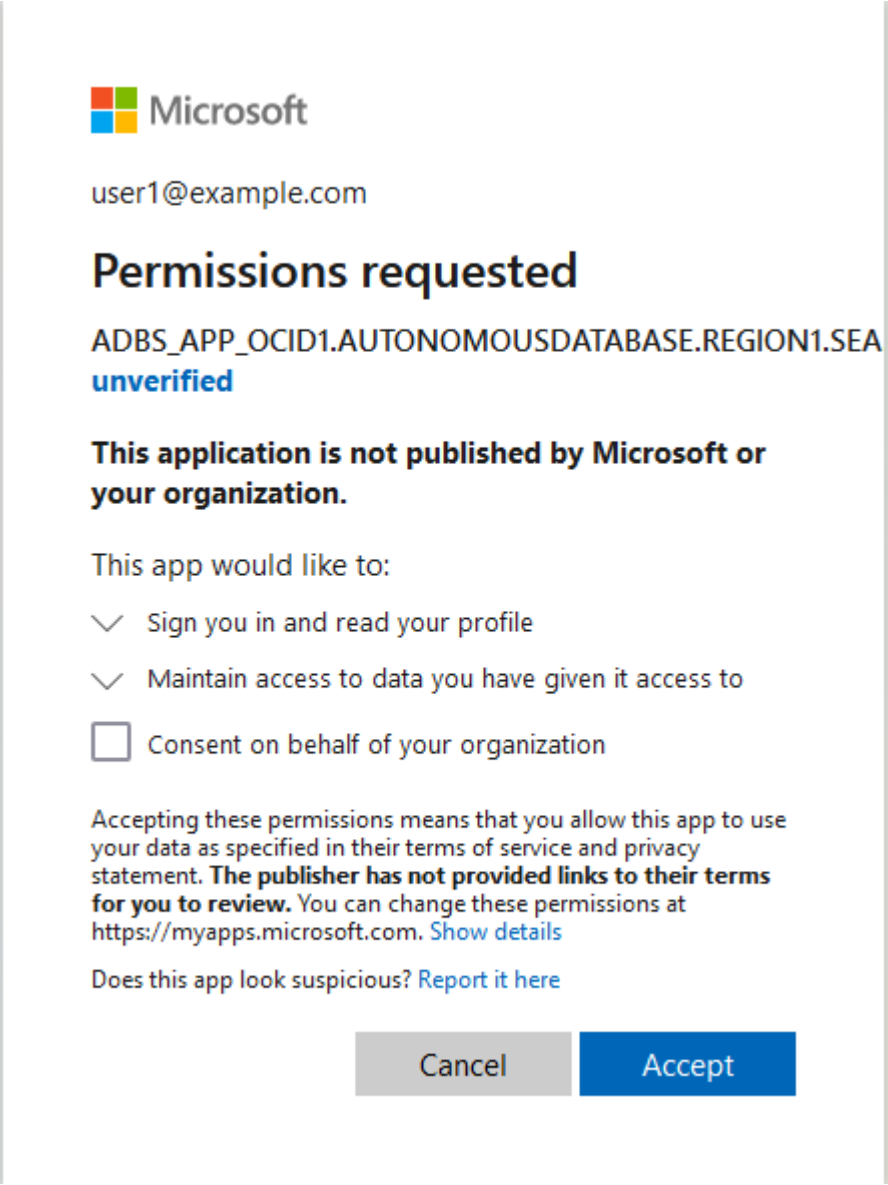
   The view `CLOUD_INTEGRATIONS` is available to the `ADMIN` user or to a user with `DWROLE` role.

2. In a browser, open the Azure consent URL, specified by the `azure_consent_url` parameter.

   For example, copy and enter the URL in your browser:

   ```
   https://login.microsoftonline.com/f8cdef31-91255a/oauth2/v2.0/authorize?
   client_id=d66f1b5-1250d5445c0b&response_type=code&scope=User.read
   ```

   The **Permissions requested** page opens and shows a consent request, similar to the following:



3. To provide consent click **Accept**.
4. On the Microsoft Azure console, assign the roles you want to grant to allow access to the specified Azure resources.

   For example, if you want to access Azure Blob Storage from Autonomous Database, assign roles so that the Azure application (the service principal) has access to Azure Blob Storage.

> **Note:**
>
> To work with Azure Blob Storage, you need an Azure storage account. If you do not have an Azure storage account, create a storage account. See Create a storage account for more information.

   a. On the Microsoft Azure console, under **Azure services**, select **Storage accounts**.

b. Under **Storage accounts**, click the storage account you want to grant service principal access to.

c. On the left, click **Access Control (IAM)**.

d. From the top area, click **+ Add → Add role assignment**.

e. In the search area enter text to narrow the list of roles that you see. For example, enter **Storage Blob** to show the available roles that contain **Storage Blob**.

f. Select one or more roles as appropriate for the access you want to grant. For example, select **Storage Blob Data Contributor**.

g. Click **Next**.

h. In the **Add role assignment**, under **Members** click **+ Select members**.

i. Under **Select members**, in the select field, enter the `azure_app_name` listed in Step 1 (the `param_value` column of the `CLOUD_INTEGRATIONS` view).

j. Select the application.

For example, click `ADBS_APP_OCID1.AUTONOMOUSDATABASE.REGION1.SEA.ANZWKLJSZLYNB3AAWLYL3JVC4ICEXLB3ZG6WTCX735JSSY2NRHOBU4DZOOVA`

k. Click **Select**.

l. Click **Review + assign**.

5. Click **Review + Assign** again.

After assigning a role you need to wait, as role assignments may take up to five minutes to propagate in Azure.

This example shows steps to grant roles for accessing Azure Blob Storage. If you want to provide access for other Azure services you need to perform equivalent steps for the additional Azure services to allow the Azure application (the service principal) to access the Azure service.

Use Azure Service Principal with DBMS_CLOUD

When you make `DBMS_CLOUD` calls to access Azure resources and specify the credential name as `AZURE$PA`, the authentication on the Azure side happens using the Azure service principal.

If you have not already done so, perform the required prerequisite steps:

- Enable the ADMIN schema or another schema to use Azure service principal authentication. See Enable Azure Service Principal for more information.
- Consent the application and perform the Azure role assignment grants. See Provide Azure Application Consent and Assign Roles for more information.

To use a `DBMS_CLOUD` procedure or function with Azure service principal, specify `AZURE$PA` as the credential name. For example, you can access Azure Blob Storage using Azure service principal credentials as follows:

```
SELECT * FROM DBMS_CLOUD.LIST_OBJECTS('AZURE$PA', 'https://
treedata.blob.core.windows.net/treetypes/');

OBJECT_NAME BYTES CHECKSUM                CREATED
LAST_MODIFIED
----------- ----- ----------------------- --------------------
--------------------
```

```
trees.txt      58 aCB1qMOPVobDLCXG+2fcvg== 2022-04-07T23:03:01Z
2022-04-07T23:03:01Z
```

If you compare the steps required to access object storage, as shown in Create Credentials and Copy Data into an Existing Table, notice that Step 1, creating credentials is not required because you are using an Azure service principal called `AZURE$PA`.

Disable Azure Service Principal

To disable access to Azure resources from Autonomous Database with Azure service principal, use `DBMS_CLOUD_ADMIN.DISABLE_PRINCIPAL_AUTH`.

To disable Azure service principal on Autonomous Database:

```
BEGIN
    DBMS_CLOUD_ADMIN.DISABLE_PRINCIPAL_AUTH(
        provider => 'AZURE',
        username => 'adb_user');
END;
/
```

When the `provider` value is `AZURE` and the `username` is a user other than the `ADMIN` user, the procedure revokes the privileges from the specified user. In this case, the `ADMIN` user and other users can continue to use `ADMIN.AZURE$PA` and the application that is created for the Autonomous Database instance remains on the instance.

When the `provider` value is `AZURE` and the `username` is `ADMIN`, the procedure disables Azure service principal based authentication and deletes the Azure service principal application on the Autonomous Database instance. In this case, if you want to enable Azure service principal you must perform all the steps required to use Azure service principal again, including the following:

- Enable the `ADMIN` schema or another schema to use Azure service principal authentication. See Enable Azure Service Principal for more information.
- Consent the application and perform the Azure role assignment grants. See Provide Azure Application Consent and Assign Roles for more information.

See DISABLE_PRINCIPAL_AUTH Procedure for more information.

Notes for Azure Service Principal

Notes for using Azure service principal.

- **Cloning an Autonomous Database instance with Azure service principal**: When you clone an instance with Azure service principal enabled, the Azure service principal configuration is not carried over to the clone. Perform the steps to enable Azure service principal on the clone if you want to enable Azure service principal on a cloned instance.

Use Google Service Account to Access Google Cloud Platform Resources

You can use a Google service account to access Google Cloud Platform (GCP) resources from an Autonomous Database instance.

- About Using a Google Service Account to Access Google Cloud Resources
  When you use Google service account based authentication with Autonomous Database, an application can securely access Google Cloud Platform (GCP) resources without

creating and saving credentials based on long-term IAM access keys for the GCP resources.

- Enable Google Service Account and Find the GCP Service Account Name
  Prior to using a Google Cloud Platform (GCP) resource with a Google service account you need to enable GCP access for your Autonomous Database instance.
- Assign Roles to the Google Service Account and Provide Access for GCP Resources
  To use Google Cloud Platform (GCP) resources from an Autonomous Database instance, you or a Google Cloud Administrator must assign roles and privileges to the Google service account that your application accesses. In addition to assigning roles for the Google service account, for any GCP resources you want to use a Google Cloud administrator must add Google IAM principals.
- Use Google Service Account with DBMS_CLOUD
  When you make `DBMS_CLOUD` calls to access Google Cloud Platform (GCP) resources and specify the credential name as `GCP$PA`, the authentication on the Google Cloud Platform side happens using a Google service account.
- Disable Google Service Account
  To disable Google service account access to Google Cloud Platform (GCP) resources, use `DBMS_CLOUD_ADMIN.DISABLE_PRINCIPAL_AUTH`.
- Google Service Account Notes
  Notes for using Google service account.

About Using a Google Service Account to Access Google Cloud Resources

When you use Google service account based authentication with Autonomous Database, an application can securely access Google Cloud Platform (GCP) resources without creating and saving credentials based on long-term IAM access keys for the GCP resources.

A Google service account is a special kind of GCP account used by an application. You can use a Google service account to make authorized GCP REST API calls from an application (after the service account is given access permissions through IAM role configuration). When an application makes calls with GCP service account based authentication, the initial call generates a temporary access token through OAuth2.0. The OAuth2.0 access token is valid for one hour. Subsequent requests within the hour use the OAuth2.0 access token to make authorized GCP REST API calls.

For example, you may want to load data from Google Cloud Storage into your Autonomous Database, perform some operation on the data, and then write the modified data back to Google Cloud Storage. You can do this without using a service account if you have GCP user credentials to access Google Cloud Storage. However, using a role-based Google service account to access GCP resources from Autonomous Database has the following benefits:

- You can create role-based access, with different policies for different users or schemas that need access to GCP resources from an Autonomous Database instance. This allows you to set a policy to limit access to resources by role. For example, setting a policy that is limited to read-only access, by role, to a Google Cloud Storage bucket.
- Google service account based credentials provide better security, as you do not need to provide long-term user credentials in code when your application accesses GCP resources. Autonomous Database manages the temporary credentials for the Google service account and does not need to store GCP resource user credentials in the database.

See Service accounts for information on Google service accounts.

Enable Google Service Account and Find the GCP Service Account Name

Prior to using a Google Cloud Platform (GCP) resource with a Google service account you need to enable GCP access for your Autonomous Database instance.

1. Enable Google service account authentication with `DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH`.

   For example, to enable Google service account authentication for the `ADMIN` user:

   ```
   BEGIN
       DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(
           provider => 'GCP' );
   END;
   /
   ```

   Enable Google service account authentication for a non-`ADMIN` user, `adb_user` as follows:

   ```
   BEGIN
       DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(
           provider => 'GCP',
           username => 'adb_user');
   END;
   /
   ```

   If you want the specified user to have privileges to enable Google service account authentication for other users, set the `params` parameter `grant_option` to `TRUE`.

   ```
   BEGIN
       DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(
           provider => 'GCP',
           username => 'adb_user',
           params   => JSON_OBJECT('grant_option' value TRUE));
   END;
   /
   ```

   After you run `DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH` with `grant_option` set to `TRUE`, `adb_user` can enable Google service account authentication for another user. For example, if you connect as `adb_user`, you can run the following command to enable GCP service account access for `adb_user2`:

   ```
   BEGIN
       DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(
           provider => 'GCP',
           username => 'adb_user2');
   END;
   /
   ```

2. When `DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH` runs it creates a Google service account. Query `CLOUD_INTEGRATIONS` to obtain the service account details for your Autonomous Database instance.

```
SELECT * FROM CLOUD_INTEGRATIONS WHERE param_name = 'gcp_service_account';

PARAM_NAME           PARAM_VALUE
-------------------------------------------------------------------------
gcp_service_account  GCP-SA-22222-32222@gcp-example.iam.gserviceaccount.com
```

3. Note the `gcp_service_account` parameter value as you must supply this value when you configure GCP resources.

See ENABLE_PRINCIPAL_AUTH Procedure for more information.

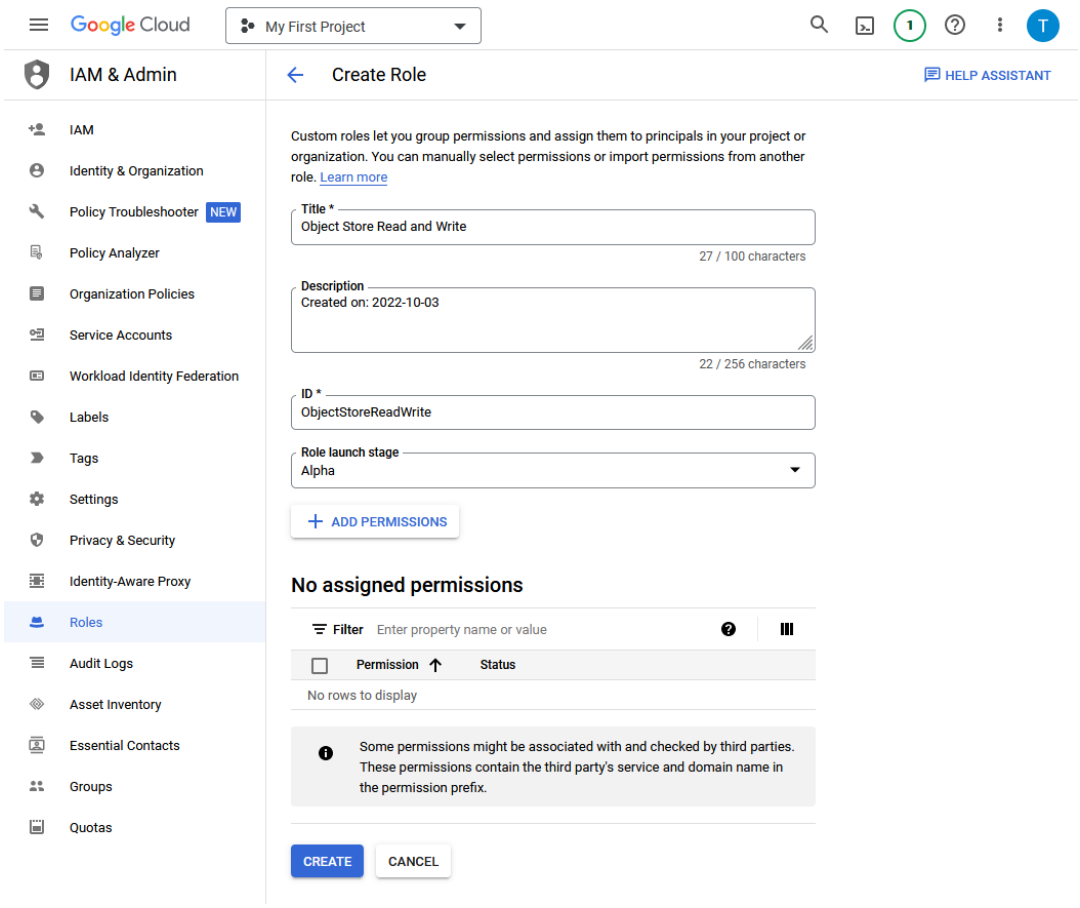
Assign Roles to the Google Service Account and Provide Access for GCP Resources

To use Google Cloud Platform (GCP) resources from an Autonomous Database instance, you or a Google Cloud Administrator must assign roles and privileges to the Google service account that your application accesses. In addition to assigning roles for the Google service account, for any GCP resources you want to use a Google Cloud administrator must add Google IAM principals.

As a prerequisite, first enable the Google service account on your Autonomous Database instance. See Enable Google Service Account and Find the GCP Service Account Name for more information.

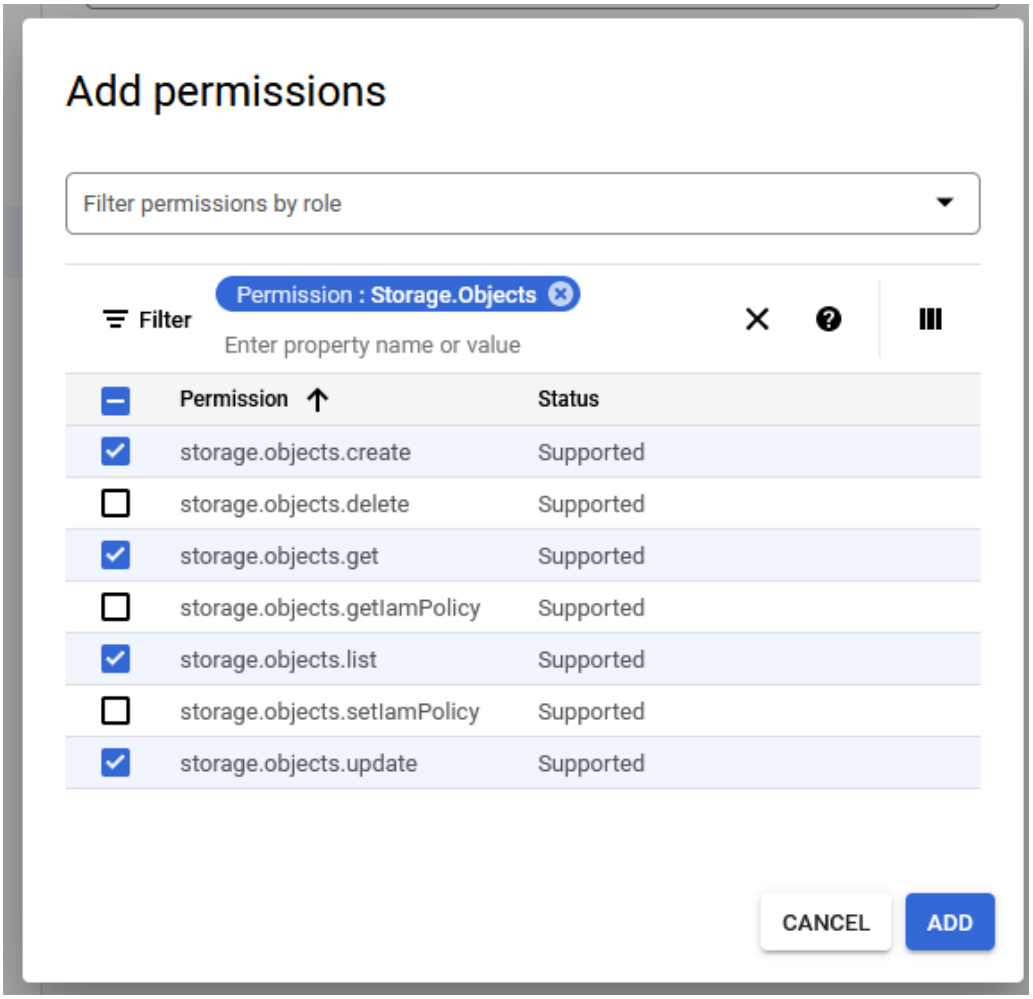
1. Open the Google Cloud Console for your account.
2. Create roles with the specified permissions.
   a. From the navigation menu, select **IAM & Admin**.
   b. From the IAM & Admin navigator, select **Roles**.
   c. On the Roles page, click More Actions and select **+ CREATE ROLE**.

   For example, you can create a role **Object Store Read and Write** to control the use of an Object Store bucket.



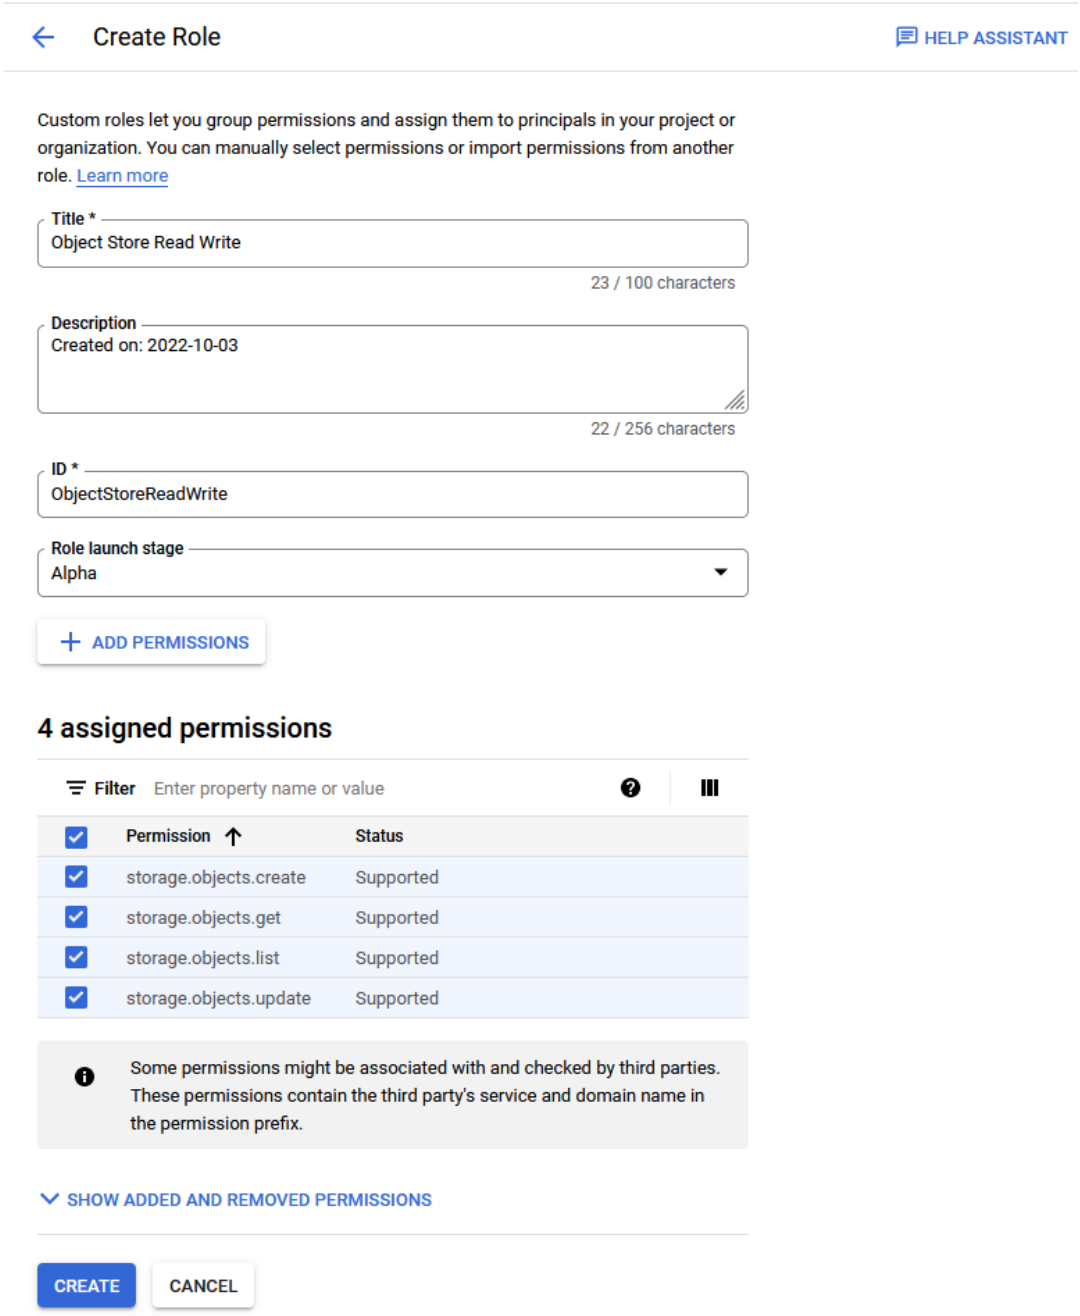
3. On the Create Role page, click **+ ADD PERMISSIONS**.
   a. Select filters to limit the list of permissions.

      For example, enter the filter **Permission: Storage.Objects** to show only the Object Store permissions.



b. On the Add permissions dialog, click **ADD**.

4. On the Create Role page, click **CREATE**.



5. Add roles and principals for the resource you want to access.

   For example, if you want to access Google Cloud Storage using the role you just created, **Object Store Read Write**:

   a. From the navigator, select **Cloud Storage** and select **Buckets**.

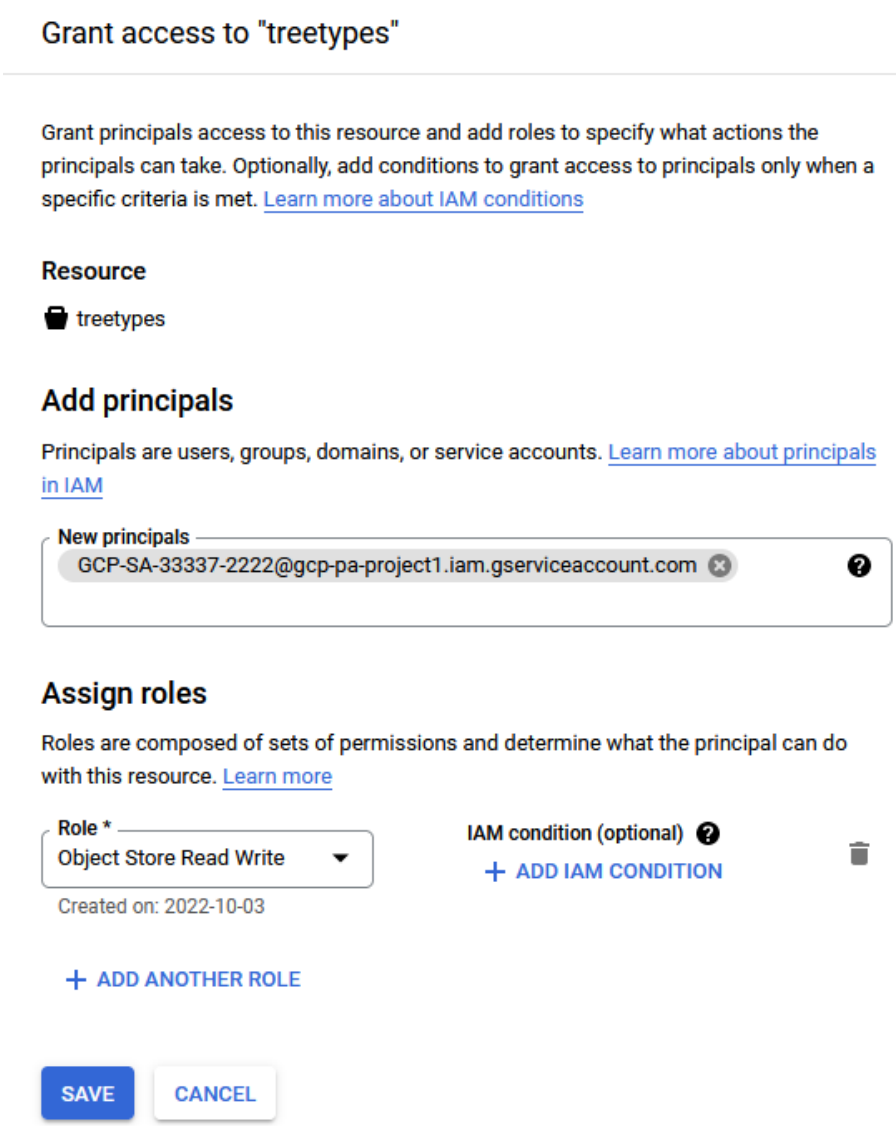
   b. Select the bucket you want to use, and click **PERMISSIONS**.

   c. Click **+ ADD PRINCIPAL**.

6. On the Grant access to "*bucketname*" dialog, add roles and a principals for the selected resource.

   a. Under **Add principals** add the value of the `gcp_service_account` parameter from your Autonomous Database instance.

b. On the Grant access to "*bucketname*" dialog, enter roles under **Assign Roles** and then click **SAVE**.



After you complete these steps the roles and principals are assigned. This allows your application running on the Autonomous Database instance to access the GCP resource with a Google service account.

Use Google Service Account with DBMS_CLOUD

When you make `DBMS_CLOUD` calls to access Google Cloud Platform (GCP) resources and specify the credential name as `GCP$PA`, the authentication on the Google Cloud Platform side happens using a Google service account.

If you have not already done so, perform the prerequisite steps:

- Enable the ADMIN schema or another schema to use Google service account authentication. See Enable Google Service Account and Find the GCP Service Account Name for more information.

- Perform the Google Cloud Platform role assignments for the resources you want to access. See Assign Roles to the Google Service Account and Provide Access for GCP Resources for more information.

To use a `DBMS_CLOUD` procedure or function with Google service account authentication:

1. Use `GCP$PA` as the credential name.
2. Construct the URI to access the GCP resource using virtual hosted-style:

   `https://BUCKET_NAME.storage.googleapis.com/OBJECT_NAME`

   For example, you can access Google Cloud Storage using Google service account credentials as follows:

```
SELECT * FROM DBMS_CLOUD.LIST_OBJECTS('GCP$PA', 'https://
treetypes.storage.googleapis.com/' );

OBJECT_NAME BYTES CHECKSUM                          CREATED LAST_MODIFIED
----------- ----- -------------------------------- -------
------------------------
trees.txt      58 682075a8c38f5686c32c25c6fb67dcbe
2022-10-05T20:03:55.253Z
```

See the following for more information:

- See Request endpoints for more information on GCP virtual hosted-style requests.
- See LIST_OBJECTS Function.

Disable Google Service Account

To disable Google service account access to Google Cloud Platform (GCP) resources, use `DBMS_CLOUD_ADMIN.DISABLE_PRINCIPAL_AUTH`.

When the `provider` value is `GCP` and the `username` is a user other than the `ADMIN` user, the procedure revokes the privileges from the specified user. In this case, the `ADMIN` user and other users can continue to use `GCP$PA`.

For example, to revoke privileges for `adb_user`:

```
BEGIN
    DBMS_CLOUD_ADMIN.DISABLE_PRINCIPAL_AUTH(
        provider => 'GCP',
        username => 'adb_user');
END;
/
```

When the `provider` value is `GCP` and the `username` is `ADMIN`, the procedure disables Google service account access on the Autonomous Database instance. The default value for `username` is `ADMIN`.

For example:

```
BEGIN
    DBMS_CLOUD_ADMIN.DISABLE_PRINCIPAL_AUTH(
        provider => 'GCP' );
```

```
END;
/
```

See DISABLE_PRINCIPAL_AUTH Procedure for more information.

Google Service Account Notes

Notes for using Google service account.

- **Google Cloud Platform (GCP) character restriction**: `DBMS_CLOUD` does not support a URI containing an "_" to access a Google Cloud Storage bucket name. If your Google Cloud Storage bucket name contains an "_", you might see the following error:

  ```
  SELECT * FROM DBMS_CLOUD.LIST_OBJECTS('GCP$PA', 'https://
  app_bucket.storage.googleapis.com/');

  ORA-20006: Unsupported object store URI - https://
  app_bucket.storage.googleapis.com/
  ORA-06512: at "C##CLOUD$SERVICE.DBMS_CLOUD", line 1306
  ```

- **Cloning an Autonomous Database instance with a Google service account**: When you clone an instance with a Google service account enabled, the Google service account configuration is not carried over to the clone. Perform the steps to enable the Google service account on the clone if you want to enable Google service account on a cloned instance.

Load Data into Autonomous Database

Describes packages and tools to load data with Autonomous Database.

- About Data Loading
  Autonomous Database provides many different options to load data into your database.
- Load Data with Oracle Database Actions
  Oracle Database Actions provides a web-based interface with development tools, data tool, administration, and monitoring features and lets you load or access data from local files, from cloud storage, or from remote databases.
- Load Data from Files in the Cloud
  The PL/SQL package `DBMS_CLOUD` provides support for loading data from text, ORC, Parquet, and Avro files in the Cloud to your tables in Autonomous Database. In addition, using `DBMS_CLOUD` you can load Data Pump dump files in the Cloud to your tables in Autonomous Database.
- Import Data Using Oracle Data Pump on Autonomous Database
  Oracle Data Pump offers very fast bulk data and metadata movement between Oracle databases and Autonomous Databases.
- Load Data from Local Files with Oracle Database Actions
  Using Oracle Database Actions, from the Worksheet page, you can load data from local files into an existing table.
- Load Data or Query Data from Files in a Directory
  You can use `DBMS_CLOUD` procedures to load data from files in a directory, including directories created on attached network file systems. You can also use these procedures to create external tables that you can use to query data.

- Load Data from Local Files Using SQL*Loader
  Instead of using SQL*Loader Oracle recommends loading data from the Cloud Object Storage for better performance and enhanced functionality.
- Using Data Pipelines for Continuous Load and Export
  Data pipelines allow you to repeatedly load data from object store or export data to object store. Load pipelines provide continuous incremental data loading from external sources (as data arrives on object store it is loaded to a database table). Export pipelines provide continuous incremental data exporting to object store (as new data appears in a database table it is exported to object store).
- The Data Load Page
  Use the Data Load page to make more data available to your Oracle Autonomous Database. You can load data from files or databases, from links to external databases or cloud storage files, or from a live feed of data from cloud storage.
- Augment Data using Table AI Assist Tool
  The Data Studio suite of tools now features an AI-powered Table Assist Tool to effortlessly enhance your local data tables with the power of AI. You can use natural language prompts to augment and correct your information with this tool.
- Transform Data Using Data Studio
  Data Transforms is an easy-to-use graphical user interface that you can use to design graphical data transformations for data integration with Autonomous Database.
- Data Marketplace
  Data Marketplace provides an innovative platform to access and share datasets effortlessly with Autonomous Database.
- View Data Pump Jobs and Import Data with Data Pump
  Use the Data Pump page to view Data Pump jobs and use the wizard to quickly create and run Data Pump import jobs.

About Data Loading

Autonomous Database provides many different options to load data into your database.

You can load data using any of the following methods:

- You can load data using Oracle Database Actions.
- You can load data using Oracle Database tools and Oracle or other 3rd party data integration tools.
- On transaction processing systems you traditionally ingest data through routine transactions or with DML operations.
- You can load data using `DBMS_CLOUD` procedures.

In general you load data from files local to your client computer or from files stored in a cloud-based object store, or in connected file systems. To load data from files in the cloud, use either Oracle Database Actions or use the Autonomous Database PL/SQL package `DBMS_CLOUD` to load files from the cloud, in connected file systems, or the local file system. `DBMS_CLOUD` provides parallel execution procedures for bulk file upload, download, copy, and transfer activities, which streamlines the user experience and delivers optimal performance for bulk file operations.

For the fastest data loading experience Oracle recommends uploading the source files to a cloud-based object store, such as Oracle Cloud Infrastructure Object Storage, before loading the data into your database. Oracle provides support for loading files that are located locally in

your data center, but when using this method of data loading you should factor in the transmission speeds across the Internet which may be significantly slower.

For more information on Oracle Cloud Infrastructure Object Storage, see Putting Data into Object Storage and Overview of Object Storage.

> **Note:**
>
> If you are not using `ADMIN` user, ensure the user has the necessary privileges for the operations the user needs to perform. See Manage User Privileges on Autonomous Database - Connecting with a Client Tool for more information.

Load Data with Oracle Database Actions

Oracle Database Actions provides a web-based interface with development tools, data tool, administration, and monitoring features and lets you load or access data from local files, from cloud storage, or from remote databases.

On the Database Actions Data Load page you can choose to load data from a file on your local device, from cloud storage, or from a database. You can also choose to explore the data in your Oracle Autonomous Database. See The Data Load Page for detailed information and the steps for loading data using Database Actions.

See Connect with Built-In Oracle Database Actions for information on accessing Oracle Database Actions.

Load Data from Files in the Cloud

The PL/SQL package `DBMS_CLOUD` provides support for loading data from text, ORC, Parquet, and Avro files in the Cloud to your tables in Autonomous Database. In addition, using `DBMS_CLOUD` you can load Data Pump dump files in the Cloud to your tables in Autonomous Database.

The package `DBMS_CLOUD` supports loading files from the following cloud services:

- Oracle Cloud Infrastructure Object Storage
- Azure Blob Storage or Azure Data Lake Storage
- Amazon S3
- Amazon S3-Compatible, including: Wasabi Hot Cloud Storage
- GitHub Repository
- Google Cloud Storage

See DBMS_CLOUD URI Formats for more information.

- Create Credentials and Copy Data into an Existing Table
  For data loading from files in the Cloud, you need to first store your object storage credentials in your Autonomous Database and then use the procedure `DBMS_CLOUD.COPY_DATA` to load data.
- Create Credentials and Load Data Pump Dump Files into an Existing Table
  For data loading you can also use Oracle Data Pump dump files as source files.

- Load JSON on Autonomous Database
  The PL/SQL procedure `DBMS_CLOUD.COPY_COLLECTION` provides support for loading JSON documents into SODA collections. The procedure `DBMS_CLOUD.COPY_DATA` provides support for loading JSON data into an existing table in Autonomous Database.
- Monitor and Troubleshoot Loads
  All data load operations done using the PL/SQL package `DBMS_CLOUD` are logged in the tables `dba_load_operations` and `user_load_operations`:
- Examples for Loading Data into Autonomous Database
  Provides examples for loading data in various formats and from different Cloud Storage providers into Autonomous Database.

Create Credentials and Copy Data into an Existing Table

For data loading from files in the Cloud, you need to first store your object storage credentials in your Autonomous Database and then use the procedure `DBMS_CLOUD.COPY_DATA` to load data.

The source file in this example, `channels.txt`, has the following data:

```
S,Direct Sales,Direct
T,Tele Sales,Direct
C,Catalog,Indirect
I,Internet,Indirect
P,Partners,Others
```

1. Store your object store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`. For example:

```
SET DEFINE OFF
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'adb_user@example.com',
    password => 'password'
  );
END;
/
```

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials you can then use the same credential name for all data loads.

   For detailed information about the parameters, see CREATE_CREDENTIAL Procedure.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

> **Note:**
>
> Some tools like SQL*Plus and SQL Developer use the ampersand character (`&`) as a special character. If you have the ampersand character in your password use the `SET DEFINE OFF` command in those tools as shown in the example to disable the special character and get the credential created properly.

2. Load data into an existing table using the procedure `DBMS_CLOUD.COPY_DATA`. For example:

```
CREATE TABLE CHANNELS
   (channel_id CHAR(1),
    channel_desc VARCHAR2(20),
    channel_class VARCHAR2(20)
   );
/

BEGIN
 DBMS_CLOUD.COPY_DATA(
    table_name =>'CHANNELS',
    credential_name =>'DEF_CRED_NAME',
    file_uri_list =>'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/channels.txt',
    format => json_object('delimiter' value ',')
 );
END;
/
```

The parameters are:

- `table_name`: is the target table's name.
- `credential_name`: is the name of the credential created in the previous step. The `credential_name` parameter must conform to Oracle object naming conventions. See Database Object Naming Rules for more information.
- `file_uri_list`: is a comma delimited list of the source files you want to load.
- `format`: defines the options you can specify to describe the format of the source file, including whether the file is of type text, ORC, Parquet, or Avro.

  If the data in your source files is encrypted, decrypt the data by specifying the `format` parameter with the `encryption` option. See Decrypt Data While Importing from Object Storage for more information on decrypting data.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

> **Note:**
>
> Autonomous Database supports a variety of source file formats, including compressed data formats. See DBMS_CLOUD Package Format Options and the `DBMS_CLOUD compression` format option to see the supported compression types.

For detailed information about the parameters, see COPY_DATA Procedure and COPY_DATA Procedure for Avro, ORC, or Parquet Files.

Create Credentials and Load Data Pump Dump Files into an Existing Table

For data loading you can also use Oracle Data Pump dump files as source files.

The source files for this load type must be exported from the source system using the `ORACLE_DATAPUMP` access driver in External Tables. See Unloading and Loading Data with the ORACLE_DATAPUMP Access Driver for details on exporting using the `ORACLE_DATAPUMP` access driver.

To load the data you first move the dump files that were exported using the `ORACLE_DATAPUMP` access driver to your Object Store and then use `DBMS_CLOUD.COPY_DATA` to load the dump files to an existing table in your database.

The source files in this example are the Oracle Data Pump dump files, `exp01.dmp` and `exp02.dmp`.

1. Store your object store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`. For example:

```
SET DEFINE OFF
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'adb_user@example.com',
    password => 'password'
  );
END;
/
```

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials you can then use the same credential name for all data loads.

   For detailed information about the parameters, see CREATE_CREDENTIAL Procedure.

   > **Note:**
   >
   > Some tools like SQL*Plus and SQL Developer use the ampersand character (`&`) as a special character. If you have the ampersand character in your password use the `SET DEFINE OFF` command in those tools as shown in the example to disable the special character and get the credential created properly.

2. Load data into an existing table using the procedure `DBMS_CLOUD.COPY_DATA`. For example:

```
CREATE TABLE CHANNELS
   (channel_id CHAR(1),
    channel_desc VARCHAR2(20),
    channel_class VARCHAR2(20)
   );
/
```

```
BEGIN
 DBMS_CLOUD.COPY_DATA(
    table_name =>'CHANNELS',
    credential_name =>'DEF_CRED_NAME',
    file_uri_list =>'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/exp01.dmp,
                     https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/exp02.dmp',
    format => json_object('type' value 'datapump')
 );
END;
/
```

The parameters are:

- `table_name`: is the target table's name.
- `credential_name`: is the name of the credential created in the previous step. The `credential_name` parameter must conform to Oracle object naming conventions. See Database Object Naming Rules for more information.
- `file_uri_list`: is a comma delimited list of the Data Pump dump files you want to load.
- `format`: defines the options you can specify to describe the format of the source file. When you specify the `type` as `'datapump'`, the only other valid format parameter is `'rejectlimit'`.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

For detailed information about the parameters, see COPY_DATA Procedure and COPY_DATA Procedure for Avro, ORC, or Parquet Files.

Load JSON on Autonomous Database

The PL/SQL procedure `DBMS_CLOUD.COPY_COLLECTION` provides support for loading JSON documents into SODA collections. The procedure `DBMS_CLOUD.COPY_DATA` provides support for loading JSON data into an existing table in Autonomous Database.

- About Loading JSON Documents
  You load SODA collections into Autonomous Database using the PL/SQL procedure `DBMS_CLOUD.COPY_COLLECTION` and you load JSON data into a table using `DBMS_CLOUD.COPY_DATA`.
- Load a JSON File of Line-Delimited Documents into a Collection
  For loading data from collections in the Cloud, you must first store your object storage credentials in your Autonomous Database and then use the procedure `DBMS_CLOUD.COPY_COLLECTION` to load documents into a collection.
- Load an Array of JSON Documents into a Collection
  To load data from collections in the Cloud, you first store your object storage credentials in your Autonomous Database and then use PL/SQL procedure `DBMS_CLOUD.COPY_COLLECTION` to load documents into a collection. This topic explains how to load documents to your database from a JSON array in a file.

- Monitor and Troubleshoot COPY_COLLECTION Loads
  All data load operations you perform using the PL/SQL package `DBMS_CLOUD` are logged in the tables `dba_load_operations` and `user_load_operations`. Use these tables to monitor loading with `DBMS_CLOUD.COPY_COLLECTION`.
- Textual JSON Objects That Represent Extended Scalar Values
  Native binary JSON data (OSON format) extends the JSON language by adding scalar types, such as date, that correspond to SQL types and are not part of the JSON standard. Oracle Database also supports the use of textual JSON *objects* that *represent* JSON scalar values, including such nonstandard values.
- Create Credentials and Copy JSON Data into an Existing Table
  Use `DBMS_CLOUD.COPY_DATA` to load JSON data in the cloud into a table.

About Loading JSON Documents

You load SODA collections into Autonomous Database using the PL/SQL procedure `DBMS_CLOUD.COPY_COLLECTION` and you load JSON data into a table using `DBMS_CLOUD.COPY_DATA`.

- `DBMS_CLOUD.COPY_COLLECTION` supports the following typical document loading procedures:
  - Loading line-delimited JSON into a collection. See Load a JSON File of Line-Delimited Documents into a Collection for this procedure.
  - Loading an array of JSON documents into a collection. See Load an Array of JSON Documents into a Collection for this procedure.
- `DBMS_CLOUD.COPY_DATA` supports the following for loading from JSON data in Object Store:
  - Create Credentials and Copy JSON Data into an Existing Table

Load a JSON File of Line-Delimited Documents into a Collection

For loading data from collections in the Cloud, you must first store your object storage credentials in your Autonomous Database and then use the procedure `DBMS_CLOUD.COPY_COLLECTION` to load documents into a collection.

This example loads JSON values from a line-delimited file and uses the JSON file `myCollection.json`. Each value, each line, is loaded into a collection on your database as a single document.

Here's an example of such a file. It has three lines, with one object per line. Each of those objects gets loaded as a separate JSON document.

```
{ "name" : "apple", "count": 20 }
{ "name" : "orange", "count": 42 }
{ "name" : "pear", "count": 10 }
```

Before loading the data from `myCollection.json` into your database, copy the file to your object store:

- Create a bucket in the object store. For example, create an Oracle Cloud Infrastructure Object Storage bucket from the Oracle Cloud Infrastructure Object Storage link, and then in your selected compartment click **Create Bucket**, or use a command such as the following OCI CLI command to create a bucket:

  ```
  oci os bucket create --name fruit_bucket -c <compartment id>
  ```

- Copy the JSON file to your object store bucket. For example use the following OCI CLI command to copy the JSON file to the `fruit_bucket` on Oracle Cloud Infrastructure Object Storage:

```
oci os object put --bucket-name fruit_bucket \
                  --file "myCollection.json"
```

Load the JSON file from object store into a collection named `fruit` on your database as follows:

1. Store your object store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`, as shown in the following example:

```
SET DEFINE OFF
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'adb_user@example.com',
    password => 'password'
  );
END;
/
```

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials, you can use the same credential name for loading all documents.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

   See CREATE_CREDENTIAL Procedure for detailed information about the parameters.

   > **Note:**
   >
   > Some tools like SQL*Plus and SQL Developer use the ampersand character (`&`) as a special character. If you have the ampersand character in your password, then use the `SET DEFINE OFF` command in those tools as shown in the example to disable the special character, and get the credential created properly.

2. Load the data into a collection using the procedure `DBMS_CLOUD.COPY_COLLECTION`.

```
BEGIN
 DBMS_CLOUD.COPY_COLLECTION(
    collection_name => 'fruit',
    credential_name => 'DEF_CRED_NAME',
    file_uri_list   =>
      'https://objectstorage.us-ashburn-1.oraclecloud.com/n/namespace-string/b/
fruit_bucket/o/myCollection.json',
    format          =>
      JSON_OBJECT('recorddelimiter' value '''\n''')  );
```

```
END;
/
```

The parameters are:

- `collection_name`: is the name of the target collection.
- `credential_name`: is the name of the credential created in the previous step. The `credential_name` parameter must conform to Oracle object naming conventions. See Database Object Naming Rules for more information.
- `file_uri_list`: is a comma delimited list of the source files that you want to load.
- `format`: defines the options that you can specify to describe the format of the source file. The format options `characterset, compression, encryption, ignoreblanklines, jsonpath, maxdocsize, recorddelimiter, rejectlimit, type, unpackarrays` are supported while loading JSON data. Any other formats specified will result in an error.

  If the data in your source files is encrypted, decrypt the data by specifying the `format` parameter with the `encryption` option. See Decrypt Data While Importing from Object Storage for more information on decrypting data.

  See DBMS_CLOUD Package Format Options for more information.

Where *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and `fruit_bucket` is the bucket name. See Understanding Object Storage Namespaces and Overview of Object Storage for more information.

For detailed information about the parameters, see COPY_COLLECTION Procedure.

The collection `fruit` on your database now contains one document for each line in the file `myCollection.json`.

Load an Array of JSON Documents into a Collection

To load data from collections in the Cloud, you first store your object storage credentials in your Autonomous Database and then use PL/SQL procedure `DBMS_CLOUD.COPY_COLLECTION` to load documents into a collection. This topic explains how to load documents to your database from a JSON array in a file.

> **Note:**
>
> You can also load documents from a JSON array in a file into a collection using SODA for REST. See Load Purchase-Order Sample Data Using SODA for REST.

This example uses the JSON file `fruit_array.json`. The following shows the contents of the file `fruit_array.json`:

```
[{"name" : "apple", "count": 20 },
 {"name" : "orange", "count": 42 },
 {"name" : "pear", "count": 10 }]
```

Before loading data into Autonomous Database, copy the data to your object store as follows:

- Create a bucket in the object store. For example, create an Oracle Cloud Infrastructure Object Store bucket from the Oracle Cloud Infrastructure Object Storage link, in your selected Compartment, by clicking **Create Bucket**, or use a command line tool such as the following OCI CLI command:

  ```
  oci os bucket create -name json_bucket -c <compartment id>
  ```

- Copy the JSON file to the object store. For example, the following OCI CLI command copies the JSON file `fruit_array.json` to the object store:

  ```
  oci os object put --bucket-name json_bucket --file "fruit_array.json"
  ```

Load the JSON file from object store into a SODA collection named `fruit2` on your database:

1. Store your object store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`, as shown in the following example:

   ```
   SET DEFINE OFF
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'DEF_CRED_NAME',
       username => 'adb_user@example.com',
       password => 'password'
     );
   END;
   /
   ```

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials, you can use the same credential name for loading all documents.

   See CREATE_CREDENTIAL Procedure for detailed information about the parameters.

   > **Note:**
   >
   > Some tools like SQL*Plus and SQL Developer use the ampersand character (`&`) as a special character. If you have the ampersand character in your password, then use the `SET DEFINE OFF` command in those tools as shown in the example to disable the special character, and get the credential created properly.

2. Load the data into a collection using the procedure `DBMS_CLOUD.COPY_COLLECTION`.

   ```
   BEGIN
     DBMS_CLOUD.COPY_COLLECTION(
       collection_name => 'fruit2',
       credential_name => 'DEF_CRED_NAME',
       file_uri_list => 'https://objectstorage.us-ashburn-1.oraclecloud.com/n/
   namespace-string/b/json/o/fruit_array.json',
       format => '{"recorddelimiter" : "0x''01''", "unpackarrays" : "TRUE",
   "maxdocsize" : "10240000"}'
     );
   ```

```
END;
/
```

In this example you load a single JSON value which occupies the whole file, so there is no need to specify a record delimiter. To indicate that there is no record delimiter, you can use a character that does not occur in the input file. For this example, to indicate that there is no delimiter, the control character 0x01 (SOH) is set to load the JSON documents into a collection,. Thus, you specify a value for the `recorddelimiter` that does not occur in the JSON file. For example, you can use value `"0x''01''"` because this character does not occur directly in JSON text.

When `unpackarrays` parameter for format value is set to `TRUE`, the array of documents is loaded as individual documents rather than as an entire array. The unpacking of array elements is however limited to single level. If there are nested arrays in the documents, those arrays are not unpacked.

The parameters are:

- `collection_name`: is the name of the target collection.
- `credential_name`: is the name of the credential created in the previous step. The `credential_name` parameter must conform to Oracle object naming conventions. See Database Object Naming Rules for more information.
- `file_uri_list`: is a comma delimited list of the source files that you want to load.
- `format`: defines the options that you can specify to describe the format of the source file. The format options `characterset, compression, encryption, ignoreblanklines, jsonpath, maxdocsize, recorddelimiter, rejectlimit, type, unpackarrays` are supported for loading JSON data. Any other formats specified will result in an error.

  If the data in your source files is encrypted, decrypt the data by specifying the `format` parameter with the `encryption` option. See Decrypt Data While Importing from Object Storage for more information on decrypting data.

  See DBMS_CLOUD Package Format Options for more information.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

For detailed information about the parameters, see COPY_COLLECTION Procedure.

The load of `fruit_array.json`, with `DBMS_CLOUD.COPY_COLLECTION` using the format option `unpackarrays` recognizes array values in the source and instead of loading the data as a single document, as it would by default, the data is loaded in the collection `fruit2` with each value in the array as a single document.

Monitor and Troubleshoot COPY_COLLECTION Loads

All data load operations you perform using the PL/SQL package `DBMS_CLOUD` are logged in the tables `dba_load_operations` and `user_load_operations`. Use these tables to monitor loading with `DBMS_CLOUD.COPY_COLLECTION`.

- `dba_load_operations` shows all load operations
- `user_load_operations` shows the load operations in your schema

You can query these tables to see information about ongoing and completed data loads. For example, the following `SELECT` statement with a `WHERE` clause predicate on the `TYPE` column shows load operations of the type `COPY`:

```
SELECT table_name, owner_name, type, status, start_time, update_time,
logfile_table, badfile_table
   FROM user_load_operations WHERE type = 'COPY';
TABLE_NAME  OWNER_NAME   TYPE    STATUS     START_TIME
UPDATE_TIME           LOGFILE_TABLE   BADFILE_TABLE
FRUIT       ADMIN        COPY    COMPLETED  2020-04-23 22:27:37    2020-04-23
22:27:38   ""              ""
FRUIT       ADMIN        COPY    FAILED     2020-04-23 22:28:36    2020-04-23
22:28:37   COPY$2_LOG      COPY$2_BAD
```

The `LOGFILE_TABLE` column shows the name of the table you can query to look at the log of a load operation. For example, the following query shows the log of the load operation with status `FAILED` and timestamp `2020-04-23 22:28:36`:

```
SELECT * FROM COPY$2_LOG;
```

The column `BADFILE_TABLE` shows the name of the table you can query to review information for the rows reporting errors during loading. For example, the following query shows the rejected records for the load operation:

```
SELECT * FROM COPY$2_BAD;
```

Depending on the errors shown in the log and the rows shown in the `BADFILE_TABLE` table, you might be able to correct errors by specifying different format options with `DBMS_CLOUD.COPY_COLLECTION`.

> **Note:**
>
> The `LOGFILE_TABLE` and `BADFILE_TABLE` tables are stored for two days for each load operation and then removed automatically.

See DELETE_ALL_OPERATIONS Procedure for information on clearing the `user_load_operations` table.

Textual JSON Objects That Represent Extended Scalar Values

Native binary JSON data (OSON format) extends the JSON language by adding scalar types, such as date, that correspond to SQL types and are not part of the JSON standard. Oracle Database also supports the use of textual JSON *objects* that *represent* JSON scalar values, including such nonstandard values.

When you create native binary JSON data from textual JSON data that contains such **extended objects**, they can optionally be *replaced* with corresponding (native binary) JSON scalar values.

An example of an extended object is `{"$numberDecimal":31}`. It represents a JSON scalar value of the nonstandard type *decimal number*, and when interpreted as such it is replaced by a decimal number in native binary format.

For example, when you use the JSON data type constructor, `JSON`, if you use keyword `EXTENDED` then recognized extended objects in the textual input are replaced with corresponding scalar values in the native binary JSON result. If you do not include keyword `EXTENDED` then no such replacement occurs; the textual extended JSON objects are simply converted as-is to JSON objects in the native binary format.

In the opposite direction, when you use SQL/JSON function `json_serialize` to serialize binary JSON data as textual JSON data (`VARCHAR2`, `CLOB`, or `BLOB`), you can use keyword `EXTENDED` to replace (native binary) JSON scalar values with corresponding textual extended JSON objects.

> **Note:**
>
> If the database you use is an Oracle Autonomous Database then you can use PL/SQL procedure `DBMS_CLOUD.copy_collection` to create a JSON document collection from a file of JSON data such as that produced by common NoSQL databases, including Oracle NoSQL Database.
>
> If you use **`ejson`** as the value of the `type` parameter of the procedure, then recognized extended JSON objects in the input file are replaced with corresponding scalar values in the resulting native binary JSON collection. In the other direction, you can use function `json_serialize` with keyword `EXTENDED` to replace scalar values with extended JSON objects in the resulting textual JSON data.

These are the two main use cases for extended objects:

- *Exchange* (import/export):
  - Ingest existing JSON data (from somewhere) that contains extended objects.
  - Serialize native binary JSON data as textual JSON data with extended objects, for some use outside the database.
- *Inspection* of native binary JSON data: see what you have by looking at corresponding extended objects.

For exchange purposes, you can ingest JSON data from a file produced by common NoSQL databases, including Oracle NoSQL Database, converting extended objects to native binary JSON scalars. In the other direction, you can export native binary JSON data as textual data, replacing Oracle-specific scalar JSON values with corresponding textual extended JSON objects.

> **Tip:**
>
> As an example of inspection, consider an object such as `{"dob" : "2000-01-02T00:00:00"}` as the result of serializing native JSON data. Is `"2000-01-02T00:00:00"` the result of serializing a native binary value of type date, or is the native binary value just a string? Using `json_serialize` with keyword `EXTENDED` lets you know.

The mapping of extended object fields to scalar JSON types is, in general, many-to-one: more than one kind of extended JSON object can be mapped to a given scalar value. For example, the extended JSON objects `{"$numberDecimal":"31"}` and `{"$numberLong:"31"}` are both

translated as the value 31 of JSON-language scalar type number, and item method `type()` returns `"number"` for each of those JSON scalars.

Item method `type()` reports the JSON-language scalar type of its targeted value (as a JSON string). Some scalar values are distinguishable internally, even when they have the same scalar type. This generally allows function `json_serialize` (with keyword `EXTENDED`) to reconstruct the original extended JSON object. Such scalar values are distinguished internally either by using *different SQL types* to implement them or by *tagging them with the kind of extended JSON object* from which they were derived.

When `json_serialize` reconstructs the original extended JSON object the result is not always *textually* identical to the original, but it is always *semantically* equivalent. For example, `{"$numberDecimal":"31"}` and `{"$numberDecimal":31}` are semantically equivalent, even though the field values differ in type (string and number). They are translated to the same internal value, and each is tagged as being derived from a `$numberDecimal` extended object (same tag). But when serialized, the *result for both* is `{"$numberDecimal":31}`. Oracle always uses the most directly relevant type for the field value, which in this case is the JSON-language value `31`, of scalar type number.

Table 3-1 presents correspondences among the various types used. It maps across (1) types of extended objects used as input, (2) types reported by item method `type()`, (3) SQL types used internally, (4) standard JSON-language types used as output by function `json_serialize`, and (5) types of extended objects output by `json_serialize` when keyword `EXTENDED` is specified.

Table 3-1 Extended JSON Object Type Relations

Extended Object Type (Input)	Oracle JSON Scalar Type (Reported by type())	SQL Scalar Type	Standard JSON Scalar Type (Output)	Extended Object Type (Output)
`$numberDouble` with value a JSON number, a string representing the number, or one of these strings: `"Infinity"`, `"-Infinity"`, `"Inf"`, `"-Inf"`, `"Nan"`[1]	double	`BINARY_DOUBLE`	number	`$numberDouble` with value a JSON number or one of these strings: `"Inf"`, `"-Inf"`, `"Nan"`[2]
`$numberFloat` with value the same as for `$numberDouble`	float	`BINARY_FLOAT`	number	`$numberFloat` with value the same as for `$numberDouble`
`$numberDecimal` with value the same as for `$numberDouble`	number	`NUMBER`	number	`$numberDecimal` with value the same as for `$numberDouble`
`$numberInt` with value a signed 32-bit integer or a string representing the number	number	`NUMBER`	number	`$numberInt` with value the same as for `$numberDouble`
`$numberLong` with value a JSON number or a string representing the number	number	`NUMBER`	number	`$numberLong` with value the same as for `$numberDouble`

Table 3-1 (Cont.) Extended JSON Object Type Relations

Extended Object Type (Input)	Oracle JSON Scalar Type (Reported by type())	SQL Scalar Type	Standard JSON Scalar Type (Output)	Extended Object Type (Output)
`$binary` with value one of these: • a string of base-64 characters • An object with fields `base64` and `subType`, whose values are a string of base-64 characters and the number `0` (arbitrary binary) or `4` (UUID), respectively When the value is a string of base-64 characters, the extended object can also have field `$subtype` with value 0 or 4, expressed as a one-byte integer (0-255) or a 2-character hexadecimal string. representing such an integer	binary	`BLOB` or `RAW`	string Conversion is equivalent to the use of SQL function `rawtohex`.	One of the following: • `$binary` with value a string of base-64 characters • `$rawid` with value a string of 32 hexadecimal characters, if input had a `subType` value of `4` (UUID)
`$oid` with value a string of 24 hexadecimal characters	binary	`RAW(12)`	string Conversion is equivalent to the use of SQL function `rawtohex`.	`$rawid` with value a string of 24 hexadecimal characters
`$rawhex` with value a string with an even number of hexadecimal characters	binary	`RAW`	string Conversion is equivalent to the use of SQL function `rawtohex`.	`$binary` with value a string of base-64 characters, right-padded with = characters
`$rawid` with value a string of 24 or 32 hexadecimal characters	binary	`RAW`	string Conversion is equivalent to the use of SQL function `rawtohex`.	`$rawid`
`$oracleDate` with value an ISO 8601 date string	date	`DATE`	string	`$oracleDate` with value an ISO 8601 date string
`$oracleTimestamp` with value an ISO 8601 timestamp string	timestamp	`TIMESTAMP`	string	`$oracleTimestamp` with value an ISO 8601 timestamp string
`$oracleTimestampTZ` with value an ISO 8601 timestamp string with a numeric time zone offset or with `Z`	timestamp with time zone	`TIMESTAMP WITH TIME ZONE`	string	`$oracleTimestampTZ` with value an ISO 8601 timestamp string with a numeric time zone offset or with `Z`

Table 3-1 (Cont.) Extended JSON Object Type Relations

Extended Object Type (Input)	Oracle JSON Scalar Type (Reported by type())	SQL Scalar Type	Standard JSON Scalar Type (Output)	Extended Object Type (Output)
`$date` with value one of the following: • An integer millisecond count since January 1, 1990 • An ISO 8601 timestamp string • An object with field `numberLong` with value an integer millisecond count since January 1, 1990	timestamp with time zone	`TIMESTAMP WITH TIME ZONE`	string	`$oracleTimestampTZ` with value an ISO 8601 timestamp string with a numeric time zone offset or with `Z`
`$intervalDaySecond` with value an ISO 8601 interval string as specified for SQL function `to_dsinterval`	daysecondInterval	`INTERVAL DAY TO SECOND`	string	`$intervalDaySecond` with value an ISO 8601 interval string as specified for SQL function `to_dsinterval`
`$intervalYearMonth` with value an ISO 8601 interval string as specified for SQL function `to_yminterval`	yearmonthInterval	`INTERVAL YEAR TO MONTH`	string	`$intervalYearMonth` with value an ISO 8601 interval string as specified for SQL function `to_yminterval`
Two fields: • Field `$vector` with value an array whose elements are numbers or the strings `"Nan"`, `"Inf"`, and `"-Inf"` (representing not-a-number and infinite values). • Field `$vectorElementType` with string value either`"float32"` or `"float64"`. These correspond to IEEE 32-bit and IEEE 64-bit numbers, respectively.	vector	`VECTOR`	array of numbers	Two fields: • Field `$vector` with value an array whose elements are numbers or the strings `"Nan"`, `"Inf"`, and `"-Inf"` (representing not-a-number and infinite values). • Field `$vectorElementType` with string value either`"float32"` or `"float64"`.

[1] The string values are interpreted case-insensitively. For example, `"NAN"` `"nan"`, and `"nAn"` are accepted and equivalent, and similarly `"INF"`, `"inFinity"`, and `"iNf"`. Infinitely large (`"Infinity"` or `"Inf"`) and small (`"-Infinity"` or `"-Inf"`) numbers are accepted with either the full word or the abbreviation.

[2] On output, only these string values are used — no full-word *Infinity* or letter-case variants.

> **See Also:**
>
> IEEE Standard for Floating-Point Arithmetic (IEEE 754)

Create Credentials and Copy JSON Data into an Existing Table

Use `DBMS_CLOUD.COPY_DATA` to load JSON data in the cloud into a table.

The source file in this example is a JSON data file.

1. Store your object store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`. For example:

```
SET DEFINE OFF
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'adb_user@example.com',
    password => 'password'
  );
END;
/
```

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials you can then use the same credential name for all data loads.

   For detailed information about the parameters, see CREATE_CREDENTIAL Procedure.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

2. Load JSON data into an existing table using the procedure `DBMS_CLOUD.COPY_DATA`.

   For example:

```
CREATE TABLE WEATHER2
    (WEATHER_STATION_ID VARCHAR2(20),
     WEATHER_STATION_NAME VARCHAR2(50));
/

BEGIN
  DBMS_CLOUD.COPY_DATA(
      table_name      => 'WEATHER2',
      credential_name => 'DEF_CRED_NAME',
      file_uri_list   => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/jsonfiles*',
      format          =>  JSON_OBJECT('type' value 'json', 'columnpath'
value '["$.WEATHER_STATION_ID",
          "$.WEATHER_STATION_NAME"]')
    );
END;
/
```

   The parameters are:

   - `table_name`: is the target table's name.
   - `credential_name`: is the name of the credential created in the previous step.
   - `file_uri_list`: is a comma delimited list of the source files you want to load. You can use wildcards in the file names in your URIs. The character "*" can be used as the wildcard for multiple characters, the character "?" can be used as the wildcard for a single character.

- `format`: for `DBMS_CLOUD.COPY_DATA` with JSON data, the `type` is `json`. Specify other format values to define the options to describe the format of the JSON source file. See DBMS_CLOUD Package Format Options for more information.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

For detailed information about the parameters, see COPY_DATA Procedure.

Monitor and Troubleshoot Loads

All data load operations done using the PL/SQL package `DBMS_CLOUD` are logged in the tables `dba_load_operations` and `user_load_operations`:

- `dba_load_operations`: shows all load operations.
- `user_load_operations`: shows the load operations in your schema.

Query these tables to see information about ongoing and completed data loads. Use a `SELECT` statement with a `WHERE` clause predicate on the `TYPE` column to show load operations with the type `COPY`.

For example:

```
SELECT table_name, owner_name, type, status, start_time, update_time, logfile_table,
badfile_table
   FROM user_load_operations WHERE type = 'COPY';

TABLE_NAME OWNER_NAME TYPE STATUS    START_TIME               UPDATE_TIME
LOGFILE_TABLE BADFILE_TABLE
---------- ---------- ---- --------- ------------------------ ------------------------
------------- -------------
TREEDATA   ADMIN      COPY COMPLETED 2022-10-20T23:15:19.990Z 2022-10-20T23:15:24.238Z
COPY$1_LOG    COPY$1_BAD
```

The `LOGFILE_TABLE` column shows the name of the table you can query to look at the log of a load operation. For example, the following query shows the log for this load operation:

```
SELECT * FROM COPY$1_LOG;
```

The column `BADFILE_TABLE` shows the name of the table you can query to look at the rows with errors during loading. For example, the following query shows the rejected records for the load operation. If there are not any rejected rows in the operation, the query does not show any rejected rows.

```
SELECT * FROM COPY$1_BAD;
```

Depending on the errors shown in the log and the rows shown in the `BADFILE_TABLE` file you can correct the error by specifying the correct format options in `DBMS_CLOUD.COPY_DATA`.

When the format `type` is "`datapump`", any rows rejected up to the specified `rejectlimit` are logged in the log file, but a `BADFILE_TABLE` is not generated.

By default the LOGFILE_TABLE and BADFILE_TABLE files are retained for two days and then automatically removed. You can change the number of retention days with the logretention option for the format parameter.

See DBMS_CLOUD Package Format Options for information on the enablelogs, logdir, logprefix, and logretention options for the format parameter.

See DELETE_ALL_OPERATIONS Procedure for information on clearing the log files.

See Track DBMS_CLOUD Load Operations for information on the dba_load_operations and user_load_operations views.

Monitor and Troubleshoot Bulk File Operations

See Monitor and Troubleshoot Bulk File Loads for information on monitoring and troubleshooting for bulk file operations.

Monitor and Troubleshoot ORC, Parquet, or Avro File Loading

As with other data files, ORC, Parquet, and Avro data loads generate logs that are viewable in the tables dba_load_operations and user_load_operations. Each load operation adds a record to dba[*user*]_load_operations that indicates the table containing the logs.

The log table provides summary information about the load.

> **Note:**
>
> For ORC, Parquet, or Avro files, when the format parameter type is set to the value orc, parquet or avro the BADFILE_TABLE table is always empty.
>
> - PRIMARY KEY constraint errors throw an ORA error.
> - If data for a column encounters a conversion error, for example, the target column is not large enough to hold the converted value, the value for the column is set to NULL. This does not produce a rejected record.

Examples for Loading Data into Autonomous Database

Provides examples for loading data in various formats and from different Cloud Storage providers into Autonomous Database.

- Load Data into Autonomous Database from Oracle Cloud Infrastructure Object Storage
  This example shows you how to load data from Oracle Cloud Infrastructure Object Storage to Autonomous Database using SQL commands.
- Load Data into Autonomous Database from AWS S3
  This example shows you how to load data from Amazon S3 object storage to Autonomous Database.
- Load Data into Autonomous Database from Azure Blob Storage or Azure Data Lake Storage
  This example shows you how to load data from Azure Blob Storage or Azure Data Lake Storage to Autonomous Database.
- Load Data into Autonomous Database from Google Cloud Storage
  This example shows how to load data from Google Cloud Storage to Autonomous Database.

- Load a Fixed-Width File into a New Table
  This provides an example using `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` to load data from a fixed-width source file to an external table.

Load Data into Autonomous Database from Oracle Cloud Infrastructure Object Storage

This example shows you how to load data from Oracle Cloud Infrastructure Object Storage to Autonomous Database using SQL commands.

You have various options to perform data loading into Autonomous Database, such as:

- Using UI options: You can use the Data Studio Load tool user interface to create credentials for the cloud store location, select files containing data, and run data load jobs.

  See Loading Data from Cloud Storage
- Using Rest Data Services APIs: You can use the Data Studio Load tool APIs to create links to the cloud store location and run data load jobs.

  See Using Data Studio Data Load APIs in PL/SQL
- Using SQL commands as explained in this example.

All these methods use the same PL/SQL package `DBMS_CLOUD` for loading data. However, Data Studio provides additional benefits over SQL commands. It not only helps to analyze the source and create table definitions but also performs validation checks.

You require Oracle Cloud Infrastructure access credentials for user account authentication and an object URL for accessing the object in your Oracle Cloud Infrastructure Object Storage bucket.

To load data from Oracle Cloud Infrastructure Object Storage:

1. Create credentials for Oracle Cloud Infrastructure user account in the Autonomous Database.
2. Copy data from Oracle Cloud Infrastructure Object Storage to the database.

Topics

- Prepare for Loading Data from Oracle Cloud Infrastructure
  Verify the prerequisites and prepare for loading data from Oracle Cloud Infrastructure Object Storage.
- Steps for Loading Data from Oracle Cloud Infrastructure
  Run these steps to load data from Oracle Cloud Infrastructure Object Storage to Autonomous Database.

Prepare for Loading Data from Oracle Cloud Infrastructure

Verify the prerequisites and prepare for loading data from Oracle Cloud Infrastructure Object Storage.

Prerequisites

A data file, for example, `oci-data.txt` exists in the Oracle Cloud Infrastructure bucket that you can import. The sample file in this example has the following contents:

```
1,OCI Direct Sales
2,OCI Tele Sales
3,OCI Catalog
```

```
4,OCI Internet
5,OCI Partners
```

On the Oracle Cloud Infrastructure side, log in to your Oracle Cloud Infrastructure account and do the following:

1. Obtain an Auth token for the Oracle Cloud Infrastructure account.

   For more information, see Getting an Auth Token.

2. From the Object Details page, obtain the object URL for the data file stored in the Oracle Cloud Infrastructure Object Storage bucket.

   For more information, see Getting an Object Storage Object's Details.

Steps for Loading Data from Oracle Cloud Infrastructure

Run these steps to load data from Oracle Cloud Infrastructure Object Storage to Autonomous Database.

1. Store the Oracle Cloud Infrastructure account credentials in your Autonomous Database and specify a credential name. This enables the database to authenticate with your Oracle Cloud Infrastructure account and access the items in the Oracle Cloud Infrastructure Object Storage bucket.

   ```
   SET DEFINE OFF
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'OCI_CRED_NAME',
       username => 'username',
       password => 'password'
     );
   END;
   /
   ```

   > **Note:**
   >
   > Here, the *username* is your Oracle Cloud Infrastructure user name and the *password* is your user Auth token.

   For detailed information about the parameters, see CREATE_CREDENTIAL Procedure.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

   Optionally, you can test the access to Oracle Cloud Infrastructure as shown in this example.

   ```
   SELECT * FROM DBMS_CLOUD.LIST_OBJECTS('OCI_CRED_NAME', 'https://
   objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/
   bucketname/o/');
   ```

2. Create a table in your database where you want to load the data.

```
CREATE TABLE myocitable (id NUMBER, name VARCHAR2(64));
```

3. Import data from the Oracle Cloud Infrastructure bucket to your Autonomous Database.

   Specify the table name and the Oracle Cloud Infrastructure credential name followed by the Oracle Cloud Infrastructure object URL.

```
BEGIN
      DBMS_CLOUD.COPY_DATA(
          table_name => 'myocitable',
          credential_name => 'OCI_CRED_NAME',
          file_uri_list => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/oci-data.txt',
          format => json_object('delimiter' value ',')
      );
END;
/
```

   In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

   For detailed information about the parameters, see COPY_DATA Procedure.

You have successfully imported data from Oracle Cloud Infrastructure Object Storage to your Autonomous Database. You can run this statement and verify the data in your table.

```
SELECT * FROM myocitable;


ID  NAME
--  --------------
 1  OCI Direct Sales
 2  OCI Tele Sales
 3  OCI Catalog
 4  OCI Internet
 5  OCI Partners
```

For more information about loading data, see Load Data from Files in the Cloud.

Load Data into Autonomous Database from AWS S3

This example shows you how to load data from Amazon S3 object storage to Autonomous Database.

You have various options to perform data loading into Autonomous Database, such as:

- Using UI options: You can use the Data Studio Load tool user interface to create credentials for the cloud store location, select files containing data, and run data load jobs.

  See Loading Data from Cloud Storage

- Using Rest Data Services APIs: You can use the Data Studio Load tool APIs to create links to the cloud store location and run data load jobs.

  See Using Data Studio Data Load APIs in PL/SQL

- Using SQL commands as explained in this example.

All these methods use the same PL/SQL package `DBMS_CLOUD` for loading data. However, Data Studio provides additional benefits over SQL commands. It not only helps to analyze the source and create table definitions but also performs validation checks.

You require AWS access credentials for user account authentication and an object URL for accessing the object in Amazon S3 bucket.

To load data from an Amazon S3 bucket:

1. Create credentials for AWS user account in the Autonomous Database.
2. Copy data from the Amazon S3 bucket to the database.

Topics

- Prepare for Loading Data from AWS S3
  Verify the prerequisites and prepare for loading data from Amazon S3.
- Steps for Loading Data from AWS S3
  Run these steps to load data from Amazon S3 to Autonomous Database.

Prepare for Loading Data from AWS S3

Verify the prerequisites and prepare for loading data from Amazon S3.

Prerequisites

A data file, for example, `data.txt` exists in the Amazon S3 bucket that you can import. The sample file in this example has the following contents:

```
1,Direct Sales
2,Tele Sales
3,Catalog
4,Internet
5,Partners
```

On the AWS side, log in to your AWS account and do the following:

1. Grant access privileges to the AWS IAM user for the Amazon S3 bucket.

   For more information, see Controlling access to a bucket with user policies.
2. Create an access key for the user.

   For more information, see Managing access keys for IAM users.
3. Obtain an object URL for the data file stored in the Amazon S3 bucket.

   For more information, see Accessing and listing an Amazon S3 bucket.

Steps for Loading Data from AWS S3

Run these steps to load data from Amazon S3 to Autonomous Database.

1. Store the AWS access credentials in your Autonomous Database and specify a credential name. This enables the database to authenticate with your AWS user account and access the items in the S3 bucket.

   ```
   SET DEFINE OFF
   BEGIN
   ```

```
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'AWS_CRED_NAME',
    username => 'username',
    password => 'password'
  );
END;
/
```

> **Note:**
>
> Here, the *username* is your AWS Access key ID and the *password* is your user access key.

For detailed information about the parameters, see CREATE_CREDENTIAL Procedure.

Creating a credential to access AWS resources is not required if you enable Amazon Resource Names (ARNs). See Use Amazon Resource Names (ARNs) to Access AWS Resources for more information.

Optionally, you can test the access to S3 bucket as shown in this example.

```
SELECT * FROM DBMS_CLOUD.LIST_OBJECTS('AWS_CRED_NAME', ' https://aws-
bucket-01.s3.amazonaws.com/');
```

2. Create a table in your database where you want to load the data.

```
CREATE TABLE mytable (id NUMBER, name VARCHAR2(64));
```

3. Import data from the Amazon S3 bucket to your Autonomous Database.

   Specify the table name and the AWS credential name followed by the S3 object URL.

```
BEGIN
      DBMS_CLOUD.COPY_DATA(
          table_name => 'mytable',
          credential_name => 'AWS_CRED_NAME',
          file_uri_list => https://aws-bucket-01.s3.amazonaws.com/
data.txt',
          format => json_object('delimiter' value ',')
      );
END;
/
```

   For detailed information about the parameters, see COPY_DATA Procedure.

You have successfully imported data from Amazon S3 to your Autonomous Database. You can run this statement and verify the data in your table.

```
SELECT * FROM mytable;


ID  NAME
--  --------------
 1  Direct Sales
```

```
2  Tele Sales
3  Catalog
4  Internet
5  Partners
```

For more information about loading data, see Load Data from Files in the Cloud.

Load Data into Autonomous Database from Azure Blob Storage or Azure Data Lake Storage

This example shows you how to load data from Azure Blob Storage or Azure Data Lake Storage to Autonomous Database.

You have various options to perform data loading into Autonomous Database, such as:

- Using UI options: You can use the Data Studio Load tool user interface to create credentials for the cloud store location, select files containing data, and run data load jobs.

  See Loading Data from Cloud Storage
- Using Rest Data Services APIs: You can use the Data Studio Load tool APIs to create links to the cloud store location and run data load jobs.

  See Using Data Studio Data Load APIs in PL/SQL
- Using SQL commands as explained in this example.

All these methods use the same PL/SQL package `DBMS_CLOUD` for loading data. However, Data Studio provides additional benefits over SQL commands. It not only helps to analyze the source and create table definitions but also performs validation checks.

You require Azure access credentials for user account authentication and an object URL for accessing the object in your Azure Storage account container.

To load data from Azure Blob Storage or Azure Data Lake Storage:

1. Create credentials for Azure user account in the Autonomous Database.
2. Copy data from Azure Blob Storage or Azure Data Lake Storage to the database.

Topics

- Prepare for Loading Data from Azure Blob Storage or Azure Data Lake Storage
  Verify the prerequisites and prepare for loading data from Azure Blob Storage or Azure Data Lake Storage.
- Steps for Loading Data from Azure Blob Storage or Azure Data Lake Storage
  Run these steps to load data from Azure Blob Storage or Azure Data Lake Storage to Autonomous Database.

Prepare for Loading Data from Azure Blob Storage or Azure Data Lake Storage

Verify the prerequisites and prepare for loading data from Azure Blob Storage or Azure Data Lake Storage.

Prerequisites

A data file, for example, `azure-data.txt` exists in the Azure Storage account container that you can import. The sample file in this example has the following contents:

```
1,Azure Direct Sales
2,Azure Tele Sales
```

```
3,Azure Catalog
4,Azure Internet
5,Azure Partners
```

On the Azure side, log in to your Azure Storage account and do the following:

1. Grant the required role, for example Storage Blob Data Contributor, to your Azure Storage account.

   For more information, see Assign Azure roles using the Azure portal.

2. Obtain an access key for the Azure Storage account.

   For more information, see View account access keys.

3. Obtain the object URL for the data file stored in the Azure Storage account container.

Steps for Loading Data from Azure Blob Storage or Azure Data Lake Storage

Run these steps to load data from Azure Blob Storage or Azure Data Lake Storage to Autonomous Database.

1. Store the Azure Storage account credentials in your Autonomous Database and specify a credential name. This enables the database to authenticate with your Azure Storage account and access the items in the Azure Storage account container.

   ```
   SET DEFINE OFF
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'AZURE_CRED_NAME',
       username => 'username',
       password => 'password'
     );
   END;
   /
   ```

   > **Note:**
   >
   > Here, the *username* is your Azure Storage account name and the *password* is your user access key.

   For detailed information about the parameters, see CREATE_CREDENTIAL Procedure.

   Creating a credential to access Azure Blob Storage or Azure Data Lake Storage is not required if you enable Azure service principal. See Use Azure Service Principal to Access Azure Resources for more information.

   Optionally, you can test the access to Azure Blob Storage or Azure Data Lake Storage as shown in this example.

   ```
   SELECT * FROM DBMS_CLOUD.LIST_OBJECTS('AZURE_CRED_NAME', 'https://
   dbazure.blob.core.windows.net/my-azure-container/');
   ```

2. Create a table in your database where you want to load the data.

   ```
   CREATE TABLE myazuretable (id NUMBER, name VARCHAR2(64));
   ```

3. Import data from the Azure Blob Storage or Azure Data Lake Storage container to your Autonomous Database.

   Specify the table name and the Azure credential name followed by the Azure Blob Storage or Azure Data Lake Storage object URL.

```
BEGIN
      DBMS_CLOUD.COPY_DATA(
          table_name => 'myazuretable',
          credential_name => 'AZURE_CRED_NAME',
          file_uri_list => 'https://dbazure.blob.core.windows.net/my-azure-
container/azure-data.txt',
          format => json_object('delimiter' value ',')
      );
END;
/
```

   For detailed information about the parameters, see COPY_DATA Procedure.

You have successfully imported data from Azure Blob Storage or Azure Data Lake Storage to your Autonomous Database. You can run this statement and verify the data in your table.

```
SELECT * FROM myazuretable;


ID  NAME
--  --------------
 1  Azure Direct Sales
 2  Azure Tele Sales
 3  Azure Catalog
 4  Azure Internet
 5  Azure Partners
```

For more information about loading data, see Load Data from Files in the Cloud.

Load Data into Autonomous Database from Google Cloud Storage

This example shows how to load data from Google Cloud Storage to Autonomous Database.

You have various options to perform data loading into Autonomous Database, such as:

- Using UI options: You can use the Data Studio Load tool user interface to create credentials for the cloud store location, select files containing data, and run data load jobs.

  See Loading Data from Cloud Storage

- Using Rest Data Services APIs: You can use the Data Studio Load tool APIs to create links to the cloud store location and run data load jobs.

  See Using Data Studio Data Load APIs in PL/SQL

- Using SQL commands as explained in this example.

All these methods use the same PL/SQL package `DBMS_CLOUD` for loading data. However, Data Studio provides additional benefits over SQL commands. It not only helps to analyze the source and create table definitions but also performs validation checks.

You require Google Cloud Platform (GCP) access credentials for user account authentication and an object URL for accessing the object in Google Cloud Storage bucket.

To load data from a Google Cloud Storage bucket:

1. Create credentials for GCP user account in the Autonomous Database.
2. Copy data from the Google Cloud Storage bucket to the database.

Topics

- Prepare for Loading Data from Google Cloud Storage
  Verify the prerequisites and prepare for loading data from Google Cloud Storage.
- Steps for Loading Data from Google Cloud Storage
  Run these steps to load data from Google Cloud Storage to Autonomous Database.

Prepare for Loading Data from Google Cloud Storage

Verify the prerequisites and prepare for loading data from Google Cloud Storage.

Prerequisites

A data file, for example, `gcp-data.txt` exists in the Google Cloud Storage bucket that you can import. The sample file in this example has the following contents:

```
1,GCP Direct Sales
2,GCP Tele Sales
3,GCP Catalog
4,GCP Internet
5,GCP Partners
```

On the Google side, log in to your Google Cloud Platform (GCP) account and do the following:

1. Create an access key for the user.

   For more information, see HMAC keys.
2. Obtain an object URL for the data file stored in the Google Cloud Storage bucket.

   For more information, see Request endpoints.

Steps for Loading Data from Google Cloud Storage

Run these steps to load data from Google Cloud Storage to Autonomous Database.

1. Store the Google Cloud Platform (GCP) access credentials in your Autonomous Database and specify a credential name. This enables the database to authenticate with your GCP user account and access the items in the Google Cloud Storage bucket.

```
SET DEFINE OFF
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'GOOGLE_CRED_NAME',
    username => 'username',
    password => 'password'
  );
END;
/
```

> **Note:**
>
> Here, the *username* is your Google account access key and the *password* is your access key secret.

For detailed information about the parameters, see CREATE_CREDENTIAL Procedure.

Creating a credential to access Google Cloud Platform (GCP) is not required if you enable Google service account. See Use Google Service Account to Access Google Cloud Platform Resources for more information.

Optionally, you can test the access to Google Cloud Storage bucket as shown in this example.

```
SELECT * FROM DBMS_CLOUD.LIST_OBJECTS('GOOGLE_CRED_NAME', 'https://
bucketname.storage.googleapis.com/');
```

2. Create a table in your database where you want to load the data.

```
CREATE TABLE mygoogletable (id NUMBER, name VARCHAR2(64));
```

3. Import data from the Google Cloud Storage bucket to your Autonomous Database.

   Specify the table name and the GCP credential name followed by the Google Cloud Storage object URL.

```
BEGIN
      DBMS_CLOUD.COPY_DATA(
          table_name => 'mygoogletable',
          credential_name => 'GOOGLE_CRED_NAME',
          file_uri_list => 'https://bucketname.storage.googleapis.com/gcp-
data.txt',
          format => json_object('delimiter' value ',')
      );
END;
/
```

   For detailed information about the parameters, see COPY_DATA Procedure.

You have successfully imported data from Google Cloud Storage to your Autonomous Database. You can run this statement and verify the data in your table.

```
SELECT * FROM mygoogletable;


ID  NAME
--  --------------
 1  GCP Direct Sales
 2  GCP Tele Sales
 3  GCP Catalog
 4  GCP Internet
 5  GCP Partners
```

For more information about loading data, see Load Data from Files in the Cloud.

Load a Fixed-Width File into a New Table

This provides an example using `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` to load data from a fixed-width source file to an external table.

For this example, the fixed-width source file has the following data:

```
 0INDEX01    INDEX                 2001272020012720200127VALID
 1INDEX02    INDEX                 2001272020012720200127VALID
 2INDEX03    INDEX                 2001272020012720200127VALID
 3INDEX04    INDEX                 2001272020012720200127VALID
 4TABLE01    TABLE                 2001272020012720200918VALID
 5TABLE02    TABLE                 2001272020012720200918VALID
 6CLUSTER01  CLUSTER               2001272020012720200127VALID
 7INDEX05    INDEX                 2001272020012720200127VALID
 8INDEX06    INDEX                 2001272020012720200127VALID
 9INDEX07    INDEX                 2001272020012720200127VALID
10INDEX08    INDEX                 2001272020012720200127VALID
11TABLE03    TABLE                 2001272020012720200127VALID
12INDEX09    INDEX                 2001272020012720200127VALID
13INDEX10    INDEX                 2001272020012720200127VALID
14TABLE04    TABLE                 2001272020012720200127VALID
15INDEX11    INDEX                 2001272020012720200127VALID
```

1. From the console, select the compartment for your Autonomous Database, and then select the link to your Autonomous Database to open the console.

   > **Note:**
   >
   > These steps are shown using **Database Actions** to execute the PL/SQL code, and query the data. These actions can be performed from any SQL connection, connecting to the Autonomous Database as a user with the proper privileges.

2. On the Autonomous Database Details page select **Database Actions** and in the list click **SQL**.

   As an alternative, select **Database Actions** and click **View all database actions** to access the Database Actions Launchpad. From the **Development** section of the **Database Actions Launchpad**, select **SQL**.

3. Within the SQL Worksheet, enter and execute the following code:

```
BEGIN DBMS_CLOUD.CREATE_EXTERNAL_TABLE(
             table_name      => '<YOUR_TABLE_NAME>'
,            credential_name => '<YOUR_CREDENTIAL_NAME>'
,            file_uri_list   => '<YOUR_ORACLE_OBJECT_STORE_URL>'
,            format          => json_object('trimspaces' value
'rtrim','skipheaders' value '1', 'dateformat' value 'YYYYMMDD')
,            field_list      => 'object_id      (1:3)   char
                               , object_name    (4:14)  char
                               , object_type    (15:39)  char
                               , created_date1  (40:45)  date mask "YYMMDD"
                               , created_date2  (46:53)  date
                               , last_ddl_time  (54:61)  date
                               , status         (62:71)'
,            column_list     => 'object_id      number
```

```
                                        , object_name    varchar2(30)
                                        , object_type    varchar2(25)
                                        , status         varchar2(10)
                                        , created_date1  date
                                        , created_date2  date
                                        , last_ddl_time  date');
END;
/
```

Import Data Using Oracle Data Pump on Autonomous Database

Oracle Data Pump offers very fast bulk data and metadata movement between Oracle databases and Autonomous Databases.

Data Pump Import lets you import data from Data Pump files residing on Oracle Cloud Infrastructure Object Storage, Microsoft Azure, AWS S3, and Oracle Cloud Infrastructure Object Storage Classic. You can save your data to your Cloud Object Store and use Oracle Data Pump to load data to Autonomous Database.

> **Note:**
>
> When a load or import operation results in the following timezone related error, you need to upgrade your timezone file to the latest version available for your database:
>
> ```
> ORA-39405: Oracle Data Pump does not support importing from a source
> database with TSTZ version n+1
> into a target database with TSTZ version n.
> ```
>
> See Manage Time Zone File Updates on Autonomous Database for more information on this timezone related error.

Topics

- Export Your Existing Oracle Database to Import into Autonomous Database
  Use Oracle Data Pump to export your existing Oracle Database as the first step to migrate to Autonomous Database.
- Import Data with Oracle Data Pump Latest Version or an Older Version
  The steps to import data to Autonomous Database with Oracle Data Pump differ depending on the Oracle Data Pump version.
- Access Log Files for Data Pump Import
  The log files for Data Pump Import operations are stored in the directory you specify with the data pump `impdp` *directory* parameter.
- Oracle Data Pump Import and Table Compression
  Provides notes for using Oracle Data Pump import on Autonomous Database.
- Notes for Importing with Oracle Data Pump
  Lists notes for importing with Oracle Data Pump.

Export Your Existing Oracle Database to Import into Autonomous Database

Use Oracle Data Pump to export your existing Oracle Database as the first step to migrate to Autonomous Database.

Oracle recommends using Oracle Data Pump schema mode to migrate your database to Autonomous Database. You can list the schemas you want to export by using the `schemas` parameter.

For a faster migration, export your schemas into multiple Data Pump files and use parallelism. You can specify the dump file name format you want to use with the `dumpfile` parameter. Set the `parallel` parameter to at least the number of CPUs you have in your database.

You can use other Data Pump Export parameters, such as `compression`, depending on your requirements. For more information on Oracle Data Pump Export see *Oracle Database Utilities*.

You have two options to export data with Data Pump, depending on the Data Pump version and whether you include the `CREDENTIAL` parameter:

- Export Your Existing Oracle Database to a Local Database File (Export without Credential Parameter)
  Use Oracle Data Pump to export your existing Oracle Database and specify an optional directory object and file name in `directory-object-name:file-name` format, as the first step to migrate to Autonomous Database.
- Export Your Existing Oracle Database Directly to Cloud Object Storage (Export with Credential Parameter)
  Use Oracle Data Pump to export your existing Oracle Database directly to Cloud Object Storage, as the first step to migrate to Autonomous Database.

Export Your Existing Oracle Database to a Local Database File (Export without Credential Parameter)

Use Oracle Data Pump to export your existing Oracle Database and specify an optional directory object and file name in `directory-object-name:file-name` format, as the first step to migrate to Autonomous Database.

Oracle recommends using the following Data Pump parameters for faster and easier migration to Autonomous Database:

```
exclude=cluster,indextype,db_link
parallel=n
schemas=schema_name
dumpfile=export%l.dmp
```

The `exclude` parameters ensure that these object types are not exported.

With `encryption_pwd_prompt=yes` Oracle Data Pump export prompts for an encryption password to encrypt the dump files.

The following example exports the SH schema from a source Oracle Database for migration to a database with 16 CPUs:

```
expdp sh/sh@orcl \
exclude=cluster,indextype,db_link \
parallel=16 \
schemas=sh \
```

```
dumpfile=export%l.dmp \
encryption_pwd_prompt=yes
```

After you create the dumpfile(s), upload the dumpfile(s) to Cloud Object Storage. For example use `DBMS_CLOUD.PUT_OBJECT` to upload the dumpfiles.

See PUT_OBJECT Procedure for more information.

Notes for Data Pump parameters:

- If during the export with `expdp` you use the `encryption_pwd_prompt=yes` parameter, also use `encryption_pwd_prompt=yes` with your import and input the same password at the `impdp` prompt to decrypt the dump files (remember the password you supply during export). The maximum length of the encryption password is 128 bytes.
- The `dumpfile` parameter supports the `%L` and `%l` wildcards in addition to the legacy `%U` and `%u` wildcards. For example, `dumpfile=export%L.dmp`. Use the `%L` or `%l` wildcard for exports from Oracle Database Release 12.2 and higher. This wildcard expands the dumpfile file name into a 3-digit to 10-digit, variable-width incrementing integer, starting at 100 and ending at 2147483646.

  Use the legacy `%U` or `%u` wildcard for exports from Oracle Database prior to Release 12.2. If you use this option and more than 99 dump files are needed, you must specify multiple dumpfile names, each with the `%U` or `%u` parameter.

Export Your Existing Oracle Database Directly to Cloud Object Storage (Export with Credential Parameter)

Use Oracle Data Pump to export your existing Oracle Database directly to Cloud Object Storage, as the first step to migrate to Autonomous Database.

> **Note:**
>
> This option is supported for Oracle Data Pump Export 19.9 or later versions.

Oracle recommends using the following Data Pump parameters for faster and easier migration to Autonomous Database:

```
exclude=cluster,indextype,db_link
parallel=n
schemas=schema_name
dumpfile=https://objectstorage.example.com/schema_name.dmp
credential=credential_name
```

The `exclude` parameters ensure that these object types are not exported.

The following example exports the SH schema from a source Oracle Database for migration to a database with 16 CPUs:

```
expdp sh/sh@orcl \
exclude=cluster,indextype,db_link \
parallel=16 \
schemas=sh \
dumpfile=https://objectstorage.example.com/%l.dmp \
credential=sales-dept
```

Notes for Data Pump parameters:

- If during the export with `expdp` you use the `encryption_pwd_prompt=yes` parameter, also use `encryption_pwd_prompt=yes` with your import and input the same password at the `impdp` prompt to decrypt the dump files (remember the password you supply during export). The maximum length of the encryption password is 128 bytes.
- The `dumpfile` parameter supports the `%L` and `%l` wildcards in addition to the legacy `%U` and `%u` wildcards. For example, `dumpfile=export%L.dmp`. Use the `%L` or `%l` wildcard for exports from Oracle Database Release 12.2 and higher. This wildcard expands the dumpfile file name into a 3-digit to 10-digit, variable-width incrementing integer, starting at 100 and ending at 2147483646.

  Use the legacy `%U` or `%u` wildcard for exports from Oracle Database prior to Release 12.2. If you use this option and more than 99 dump files are needed, you must specify multiple dumpfile names, each with the `%U` or `%u` parameter.
- Data Pump Export does not support using an Oracle Cloud Infrastructure Object Storage pre-authenticated URL (PAR URL) for the `dumpfile` parameter.
- The `credential` parameter can be one of the following supported credential types:
  - Data Pump export supports Oracle Cloud Infrastructure Auth Token based credentials and Oracle Cloud Infrastructure Signing Key based credentials.
  - Data Pump export supports using a resource principal credential with `expdp`.

  See CREATE_CREDENTIAL Procedure for more information on the credential types for Oracle Cloud Infrastructure Cloud Object Storage.

Import Data with Oracle Data Pump Latest Version or an Older Version

The steps to import data to Autonomous Database with Oracle Data Pump differ depending on the Oracle Data Pump version.

- Import Data Using Oracle Data Pump Version 18.3 or Later
  Oracle recommends using the latest Oracle Data Pump version for importing data from Data Pump files into your Autonomous Database, as it contains enhancements and fixes for a better experience.
- Import Data Using Oracle Data Pump Version 18.3 or Later with OCI Resource Principal
  Oracle Data Pump supports importing data pump files into your Autonomous Database using an Oracle Cloud Infrastructure resource principal as a credential object.
- Import Data Using Oracle Data Pump Version 12.2 and Earlier
  You can import data from Data Pump files into your Autonomous Database using Data Pump client versions 12.2.0.1 and earlier by setting the `default_credential` parameter.
- Use Oracle Data Pump Import with a Pre-Authenticated URL
  Provides the steps to use a pre-authenticated URL (PAR URL) with Oracle Data Pump Import.

Import Data Using Oracle Data Pump Version 18.3 or Later

Oracle recommends using the latest Oracle Data Pump version for importing data from Data Pump files into your Autonomous Database, as it contains enhancements and fixes for a better experience.

Download the latest version of Oracle Instant Client, which includes Oracle Data Pump, for your platform from Oracle Instant Client Downloads. See the installation instructions on the platform install download page for the installation steps required after you download Oracle Instant Client.

In Oracle Data Pump version 18.3 and later, the `credential` argument authenticates Data Pump to the Cloud Object Storage service you are using for your source files. The `dumpfile` argument is a comma delimited list of URLs for your Data Pump files.

In Oracle Data Pump, if your source files reside on Oracle Cloud Infrastructure Object Storage you can use Oracle Cloud Infrastructure native URIs or Swift URIs. See DBMS_CLOUD URI Formats for details on these file URI formats.

To load an Oracle Data Pump dumpfile from Cloud Object Storage, run Oracle Data Pump Import and include the `credential` parameter:

1. Store your Cloud Object Storage credential using `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example, to create Oracle Cloud Infrastructure Auth Token credentials with the default Identity domain:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'adb_user@example.com',
    password => 'password'
  );
END;
/
```

   For OCI Object Storage the `username` parameter value must include the Identity domain and the user name from your profile. You can find the Identity domain associated with a user in the Oracle Cloud Infrastructure Console. For example, with the `oracleidentitycloudservice` Identity domain, this would be:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'oracleidentitycloudservice/adb_user@example.com',
    password => 'password'
  );
END;
/
```

   For more information on Oracle Cloud Infrastructure Auth Token authentication see CREATE_CREDENTIAL Procedure.

   For example, to create Oracle Cloud Infrastructure Signing Key based credentials:

```
BEGIN
   DBMS_CLOUD.CREATE_CREDENTIAL (
       credential_name => 'DEF_CRED_NAME',
       user_ocid       =>
'ocid1.user.oc1..aaaaaaaauq54mi7zdyfhw33ozkwuontjceel7fok5nq3bf2vwetkpqsoa'
,
       tenancy_ocid    =>
'ocid1.tenancy.oc1..aabbbbbbaafcue47pqmrf4vigneebgbcmmoy5r7xvoypicjqqge32ew
nrcyx2a',
       private_key     =>
'MIIEogIBAAKCAQEAtUnxbmrekwgVac6FdWeRzoXvIpA9+0r1.....wtnNpESQQQ0QLGPD8NM//
JEBg=',
```

```
       fingerprint     =>
'f2:db:f9:18:a4:aa:fc:94:f4:f6:6c:39:96:16:aa:27');
END;
/
```

For more information on Oracle Cloud Infrastructure Signing Key based credentials see CREATE_CREDENTIAL Procedure.

Supported credential types:

- Data Pump Import supports Oracle Cloud Infrastructure Auth Token based credentials and Oracle Cloud Infrastructure Signing Key based credentials.

  For more information on the credential types for Oracle Cloud Infrastructure Cloud Object Storage, see CREATE_CREDENTIAL Procedure.

- Data Pump supports using an Oracle Cloud Infrastructure Object Storage pre-authenticated URL for the `dumpfile` parameter. When you use a pre-authenticated URL, providing the `credential` parameter is required and the `credential` value can be `NULL`. See Using Pre-Authenticated Requests for more information.

- Data Pump supports using a resource principal credential with `impdp`. See Import Data Using Oracle Data Pump Version 18.3 or Later with OCI Resource Principal for more information.

2. Run Data Pump Import with the *dumpfile* parameter set to the list of file URLs on your Cloud Object Storage and the *credential* parameter set to the name of the credential you created in the previous step. For example:

```
impdp admin/password@db2022adb_high \
     directory=data_pump_dir \
     credential=def_cred_name \
     dumpfile= https://namespace-string.objectstorage.us-ashburn-1.oci.customer-
oci.com/n/namespace-string/b/bucketname/o/export%l.dmp \
     parallel=16 \
     encryption_pwd_prompt=yes \
     exclude=cluster,indextype,db_link
```

Notes for Data Pump parameters:

- If during the export with `expdp` you used the `encryption_pwd_prompt=yes` parameter then use `encryption_pwd_prompt=yes` and input the same password at the `impdp` prompt that you specified during the export.

- The `dumpfile` parameter supports the `%L` and `%l` wildcards in addition to the legacy `%U` and `%u` wildcards. For example, `dumpfile=export%L.dmp`. Use the `%L` or `%l` wildcard for exports from Oracle Database Release 12.2 and higher. This wildcard expands the dumpfile file name into a 3-digit to 10-digit, variable-width incrementing integer, starting at 100 and ending at 2147483646.

  Use the legacy `%U` or `%u` wildcard for exports from Oracle Database prior to Release 12.2. If you use this option and more than 99 dump files are needed, you must specify multiple dumpfile names, each with the `%U` or `%u` parameter.

- If your bucket and data do not reside in a commercial (OC1) realm, where OCI Dedicated Endpoints are not supported, use the following URI format for `dumpfile`:

  `https://objectstorage.`*`region`*`.oraclecloud.com/n/`*`namespace-string`*`/b/`*`bucketname`*`/o/`*`filename`*

- The `credential` parameter cannot be an Azure service principal, Amazon Resource Name (ARN), or a Google service account. See Configure Policies and Roles to Access Resources for more information on resource principal based authentication.
- When you specify the `dumpfile` value as a pre-authenticated URL, the `credential` parameter is required and the value can be `NULL`.

  See Use Oracle Data Pump Import with a Pre-Authenticated URL for details on the steps required to use a pre-authenticated URL.

For `dumpfile`, this example uses the recommended URI format using OCI Dedicated Endpoints for commercial realm (OC1). The *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Object Storage Dedicated Endpoints, Regions and Availability Domains, and Understanding Object Storage Namespaces for more information.
For the dump file URL format for different Cloud Object Storage services, see DBMS_CLOUD URI Formats.

In this example the following are excluded during the Data Pump Import:

- Clusters
- Indextypes
- Database links

To perform a full import or to import objects that are owned by other users, you need the `DATAPUMP_CLOUD_IMP` role.

For information on disallowed objects in Autonomous Database, see SQL Commands.

See Notes for Importing with Oracle Data Pump for additional notes for using Oracle Data Pump Import.

See Oracle Data Pump Import and Table Compression for details on table compression using Oracle Data Pump import on Autonomous Database.

For detailed information on Oracle Data Pump Import parameters see *Oracle Database Utilities*.

Import Data Using Oracle Data Pump Version 18.3 or Later with OCI Resource Principal

Oracle Data Pump supports importing data pump files into your Autonomous Database using an Oracle Cloud Infrastructure resource principal as a credential object.

If you use Oracle Data Pump `expdp` to export directly to Object Store then you must use the same credential that was used to export when you import with `impdp`. In this case, Oracle Data Pump import does not support Oracle Cloud Infrastructure resource principal credentials. Other methods for uploading are supported for using `impdp` using resource principal credentials. For example, if you upload Oracle Data Pump files on Object Store using `DBMS_CLOUD.PUT_OBJECT`, you can import the files using Oracle Data pump `impdp` using resource principal credentials. Likewise, when you use the Oracle Cloud Infrastructure Console to upload data pump files to Object Store, you can use resource principal credentials to import into an Autonomous Database instance with Oracle Data pump `impdp`.

In Oracle Data Pump, if your source files reside on Oracle Cloud Infrastructure Object Storage you can use Oracle Cloud Infrastructure native URIs or Swift URIs. See DBMS_CLOUD URI Formats for details on these file URI formats.

1. Configure the dynamic groups and policies and enable Oracle Cloud Infrastructure resource principal to access the Object Store location where the data pump files you want to import reside.

See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

2. Run Data Pump Import with the *dumpfile* parameter set to the list of file URLs on your Cloud Object Storage and the *credential* parameter set to `OCI$RESOURCE_PRINCIPAL`.

   For example:

```
impdp admin/password@db2022adb_high \
     directory=data_pump_dir \
     credential= 'OCI$RESOURCE_PRINCIPAL' \
     dumpfile= https://namespace-string.objectstorage.us-ashburn-1.oci.customer-
oci.com/n/namespace-string/b/bucketname/o/export%l.dmp \
     parallel=16 \
     encryption_pwd_prompt=yes \
     exclude=cluster,indextype,db_link
```

   Notes for Data Pump parameters:

   - If during the export with `expdp` you used the `encryption_pwd_prompt=yes` parameter, also use `encryption_pwd_prompt=yes` and input the same password at the `impdp` prompt that you specified during the export.
   - The `dumpfile` parameter supports the `%L` and `%l` wildcards in addition to the legacy `%U` and `%u` wildcards. For example, `dumpfile=export%L.dmp`. Use the `%L` or `%l` wildcard for exports from Oracle Database Release 12.2 and higher. This wildcard expands the dumpfile file name into a 3-digit to 10-digit, variable-width incrementing integer, starting at 100 and ending at 2147483646.

     Use the legacy `%U` or `%u` wildcard for exports from Oracle Database prior to Release 12.2. If you use this option and more than 99 dump files are needed, you must specify multiple dumpfile names, each with the `%U` or `%u` parameter.
   - If your bucket and data do not reside in a commercial (OC1) realm, where OCI Dedicated Endpoints are not supported, use the following URI format for `dumpfile`:

     `https://objectstorage.region.oraclecloud.com/n/namespace-string/b/bucketname/o/filename`
   - When you specify the `dumpfile` value as a pre-authenticated URL, the `credential` parameter is required and the value can be `NULL`.

     See Use Oracle Data Pump Import with a Pre-Authenticated URL for details on the steps required to use a pre-authenticated URL.

   For `dumpfile`, this example uses the recommended URI format using OCI Dedicated Endpoints for commercial realm (OC1). The *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Object Storage Dedicated Endpoints, Regions and Availability Domains, and Understanding Object Storage Namespaces for more information.
   For the dump file URL format for different Cloud Object Storage services, see DBMS_CLOUD URI Formats.

   In this example the following are excluded during the Data Pump Import:

   - Clusters
   - Indextypes
   - Database links

   To perform a full import or to import objects that are owned by other users, you need the `DATAPUMP_CLOUD_IMP` role.

For information on disallowed objects in Autonomous Database, see SQL Commands.

See Notes for Importing with Oracle Data Pump for additional notes for using Oracle Data Pump Import.

For detailed information on Oracle Data Pump Import parameters see *Oracle Database Utilities*.

Import Data Using Oracle Data Pump Version 12.2 and Earlier

You can import data from Data Pump files into your Autonomous Database using Data Pump client versions 12.2.0.1 and earlier by setting the `default_credential` parameter.

Data Pump Import versions 12.2.0.1 and earlier do not have the *credential* parameter. If you are using an older version of Data Pump Import you need to define a default credential property for Autonomous Database and use the `default_credential` keyword in the `dumpfile` parameter.

In Oracle Data Pump, if your source files reside on Oracle Cloud Infrastructure Object Storage you can use the Oracle Cloud Infrastructure native URIs, or Swift URIs. See DBMS_CLOUD URI Formats for details on these file URI formats.

Importing with Older Oracle Data Pump Versions and Setting `default_credential`

1. Store your Cloud Object Storage credential using `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example, to create Oracle Cloud Infrastructure Auth Token credentials:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'adb_user@example.com',
    password => 'password'
  );
END;
/
```

   For more information on Oracle Cloud Infrastructure Auth Token authentication see CREATE_CREDENTIAL Procedure.

   For example, to create Oracle Cloud Infrastructure Signing Key based credentials:

```
BEGIN
   DBMS_CLOUD.CREATE_CREDENTIAL (
       credential_name => 'DEF_CRED_NAME',
       user_ocid       =>
'ocid1.user.oc1..aaaaaaaauq54mi7zdyfhw33ozkwuontjceel7fok5nq3bf2vwetkpqsoa'
,
       tenancy_ocid    =>
'ocid1.tenancy.oc1..aabbbbbbaafcue47pqmrf4vigneebgbcmmoy5r7xvoypicjqqge32ew
nrcyx2a',
       private_key     =>
'MIIEogIBAAKCAQEAtUnxbmrekwgVac6FdWeRzoXvIpA9+0r1.....wtnNpESQQQ0QLGPD8NM//
JEBg=',
       fingerprint     =>
'f2:db:f9:18:a4:aa:fc:94:f4:f6:6c:39:96:16:aa:27');
```

```
END;
/
```

For more information on Oracle Cloud Infrastructure Signing Key based credentials see CREATE_CREDENTIAL Procedure.

Supported credential types:

- Data Pump import supports Oracle Cloud Infrastructure Auth Token based credentials and Oracle Cloud Infrastructure Signing Key based credentials.

  For more information on the credential types for Oracle Cloud Infrastructure Cloud Object Storage, see CREATE_CREDENTIAL Procedure.

- Data Pump supports using an Oracle Cloud Infrastructure Object Storage pre-authenticated URL for the `dumpfile`. When you use a pre-authenticated URL, setting the `DEFAULT_CREDENTIAL` is required and the `DEFAULT_CREDENTIAL` value can be `NULL`. See Using Pre-Authenticated Requests for more information.

- Data Pump supports using a resource principal credential with `impdp`.

2. Set the credential as the default credential for your Autonomous Database, as the ADMIN user. For example:

```
ALTER DATABASE PROPERTY SET DEFAULT_CREDENTIAL = 'ADMIN.DEF_CRED_NAME'
```

The `DEFAULT_CREDENTIAL` can be an OCI Resource Principal. For example:

```
ALTER DATABASE PROPERTY SET DEFAULT_CREDENTIAL = 'OCI$RESOURCE_PRINCIPAL'
```

See Configure Policies and Roles to Access Resources for more information on resource principal based authentication.

> **Note:**
>
> The `DEFAULT_CREDENTIAL` value cannot be an Azure service principal, Amazon Resource Name (ARN), or a Google service account.

The `DEFAULT_CREDENTIAL` value can be set to `NULL` if you are using a pre-authenticated URL.

3. Run Data Pump Import with the `dumpfile` parameter set to the list of file URLs on your Cloud Object Storage, and set the `default_credential` keyword. For example:

```
impdp admin/password@db2022adb_high \
     directory=data_pump_dir \
     dumpfile=default_credential:https://namespace-string.objectstorage.us-
ashburn-1.oci.customer-oci.com/n/namespace-string/b/bucketname/o/export%l.dmp \
     parallel=16 \
     encryption_pwd_prompt=yes \
     exclude=cluster,indextype,db_link
```

Notes for the Data Pump parameters:

- If during the export with `expdp` you used the `encryption_pwd_prompt=yes` parameter then use `encryption_pwd_prompt=yes` and input the same password at the `impdp` prompt that you specified during the export.

- The `dumpfile` parameter supports the `%L` and `%l` wildcards in addition to the legacy `%U` and `%u` wildcards. For example, `dumpfile=export%L.dmp`. Use the `%L` or `%l` wildcard for exports from Oracle Database Release 12.2 and higher. This wildcard expands the dumpfile file name into a 3-digit to 10-digit, variable-width incrementing integer, starting at 100 and ending at 2147483646.

  Use the legacy `%U` or `%u` wildcard for exports from Oracle Database prior to Release 12.2. If you use this option and more than 99 dump files are needed, you must specify multiple dumpfile names, each with the `%U` or `%u` parameter.
- If your bucket and data do not reside in a commercial (OC1) realm, where OCI Dedicated Endpoints are not supported, use the following URI format for `dumpfile`:

  `https://objectstorage.`*`region`*`.oraclecloud.com/n/`*`namespace-string`*`/b/`*`bucketname`*`/o/`*`filename`*
- When you specify the `dumpfile` value as a pre-authenticated URL, setting the `default_credential` parameter is required and the value can be `NULL`.

  See Use Oracle Data Pump Import with a Pre-Authenticated URL for details on the steps required to use a pre-authenticated URL.

For `dumpfile`, this example uses the recommended URI format using OCI Dedicated Endpoints for commercial realm (OC1). The *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Object Storage Dedicated Endpoints, Regions and Availability Domains, and Understanding Object Storage Namespaces for more information.
For the dump file URL format for different Cloud Object Storage services, see DBMS_CLOUD URI Formats.

In this example the following are excluded during the Data Pump Import:

- Clusters
- Indextypes
- Database links

> **Note:**
>
> To perform a full import or to import objects that are owned by other users, you need the `DATAPUMP_CLOUD_IMP` role.

For information on disallowed objects in Autonomous Database, see SQL Commands.

See Notes for Importing with Oracle Data Pump for additional notes for using Oracle Data Pump Import.

See Oracle Data Pump Import and Table Compression for details on table compression using Oracle Data Pump import on Autonomous Database.

For detailed information on Oracle Data Pump Import parameters see *Oracle Database Utilities*.

Use Oracle Data Pump Import with a Pre-Authenticated URL

Provides the steps to use a pre-authenticated URL (PAR URL) with Oracle Data Pump Import.

1. Use Oracle Data Pump `expdp` and export to local storage.

   See Export Your Existing Oracle Database to a Local Database File (Export without Credential Parameter) for more information.

2. Upload the dumpfile(s) to Oracle Cloud Infrastructure Object Storage.

   For example use `DBMS_CLOUD.PUT_OBJECT` to upload the dumpfiles.

3. Create a pre-authenticated URL on Oracle Cloud Infrastructure Object Storage, either on the object store bucket if there are multiple dumpfiles or on the dumpfile, if there is only one file.

   See Object Storage Pre-Authenticated Requests for more information.

4. Import the dumpfile(s) with the `dumpfile` parameter value set to the pre-authenticated URL.

   See Import Data Using Oracle Data Pump Version 18.3 or Later for more information.

Access Log Files for Data Pump Import

The log files for Data Pump Import operations are stored in the directory you specify with the data pump `impdp` *directory* parameter.

To access the log file you need to move the log file to your Cloud Object Storage using the procedure `DBMS_CLOUD.PUT_OBJECT`. For example, the following PL/SQL block moves the file `import.log` to your Cloud Object Storage:

```
BEGIN
  DBMS_CLOUD.PUT_OBJECT(
    credential_name => 'DEF_CRED_NAME',
    object_uri => 'https://objectstorage.us-ashburn-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/import.log',
    directory_name  => 'DATA_PUMP_DIR',
    file_name => 'import.log');
END;
/
```

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

For more information, see DBMS_CLOUD Subprograms and REST APIs.

Oracle Data Pump Import and Table Compression

Provides notes for using Oracle Data Pump import on Autonomous Database.

By default the Oracle Data Pump Import utility imports data with the same compression type as specified for tables on the source database (the database where you exported your data from). If you want to leave compression to Autonomous Database, specify the following parameter when you import your data:

```
TRANSFORM=TABLE_COMPRESSION_CLAUSE:NONE
```

The `TRANSFORM` parameter with this option specifies that Oracle Data Pump Import should ignore the compression type of your source tables. Using this option, Oracle Data Pump imports the tables into Autonomous Database using the default compression type, where the default compression type depends on the Autonomous Database workload type:

- **Data Warehouse**: The default table compression is Hybrid Columnar Compression.

  Oracle recommends using this default if your application primarily uses bulk load operations on your tables, as the loads will compress the data. Query performance on these tables will benefit from compression as queries need to do less IO.

  If you have staging tables replicated from other systems using Oracle GoldenGate or other replication tools, or your application primarily uses row-by-row DML operations on tables, Oracle recommends keeping the tables uncompressed or using Advanced Row Compression.
- **Transaction Processing**: The default table compression is no compression.
- **JSON Database**: The default table compression is no compression.
- **APEX**: The default table compression is no compression.

See TRANSFORM for more information on the Oracle Data Pump Import `TRANSFORM` parameter.

See About Table Compression for more information.

Notes for Importing with Oracle Data Pump

Lists notes for importing with Oracle Data Pump.

- For the best import performance use the `HIGH` database service for your import connection and set the `parallel` parameter to one quarter the number of ECPUs (.25 x **ECPU count**). If you are using OCPU compute model, set the parallel parameter to the number of OCPUs (1 x **OCPU count**).

  See Manage Concurrency and Priorities on Autonomous Database for information on database service names.
- You can also use Data Pump Import to import SODA collections on Autonomous Database. See Import SODA Collection Data Using Oracle Data Pump Version 19.6 or Later for more information.
- When you perform Oracle Data Pump export to Object Storage with a Swift URI, you must use Swift credentials to import with Oracle Data Pump import. See Oracle Cloud Infrastructure Object Storage Swift URI Format for more information on Swift URIs.

  See Export Your Existing Oracle Database Directly to Cloud Object Storage (Export with Credential Parameter) for more information.
- When you perform Oracle Data Pump export to Object Storage with a native URI, you can import using either Swift credentials or Signing Key based Credentials. See Oracle Cloud Infrastructure Object Storage Native URI Format for more information on Native URIs.

  See Export Your Existing Oracle Database Directly to Cloud Object Storage (Export with Credential Parameter) for more information.
- You must use a multiple step process to export with Oracle Data Pump Export if you want to use Oracle Data Pump Import with a PAR URL. See Use Oracle Data Pump Import with a Pre-Authenticated URL for details.

Load Data from Local Files with Oracle Database Actions

Using Oracle Database Actions, from the Worksheet page, you can load data from local files into an existing table.

- Load Data into Existing Autonomous Database Table with Oracle Database Actions
  You can load data into an existing table in Autonomous Database with the Database Actions import from file feature.

Load Data into Existing Autonomous Database Table with Oracle Database Actions

You can load data into an existing table in Autonomous Database with the Database Actions import from file feature.

Before you load data, create the table in Autonomous Database. The file formats that you can upload with the Database Actions upload feature are CSV, XLS, XLSX, TSV and TXT.

To upload data from local files to an existing table with Database Actions, do the following:

1. Access Database Actions from the Oracle Cloud Infrastructure Console or with the Database Actions link provided to you.
2. To import data, in Database Actions, under Development click **SQL**.

   This shows an SQL worksheet.
3. In the **Navigator**, right-click the table where you want to load data.
4. In the menu select **Data loading → Upload Data...**

   For example, select the `SALES` table, right-click, and select **Data loading → Upload Data...**

   This shows the Import data dialog:
5. In the Import data dialog you can either drag and drop files or click **Select files** to show a browser to select the files to import.
6. Complete the mapping for the columns you are importing. There are a number of options for column mapping. Click  (Show/Hide options) icon to show the data import and format options to change column names, skip rows, rows to load, and various other options.

   Click **Apply** to apply the options you select.
7. When you finish selecting format and mapping options, click **Next** to preview the column mapping.

   If there is a problem at this stage, information shows with more details, such as: **2 pending actions**. This means you need to correct or fix the source file data before you import.
8. Click **Next**.
9. Click **Next** to review the column mapping.

   This shows the **Review** page to review the source columns and target columns for the import:
10. Click **Finish**.
11. Click **OK** to confirm the import.

    Depending on the size of the data file you are importing, the import may take some time.

Database Actions provides history to show the status of the import and to allow you to review the results or errors associated with the import operation.

For a detailed summary of the upload process, right-click the table in the **Navigator** tab, select **Data loading**, and then select **Loaded Data**. A summary of the data loaded is displayed in the Loaded data summary dialog.

If any data failed to load, you can view the number of rows in the Failed Rows column. Click the column and a dialog is displayed showing the failed rows.

In the Loaded data summary dialog, you can also search for files loaded by schema name, table name, or file name. To remove the loaded files, click the Delete icon.

See Uploading Data from Local Files for more information on using Database Actions to upload data.

Load Data or Query Data from Files in a Directory

You can use `DBMS_CLOUD` procedures to load data from files in a directory, including directories created on attached network file systems. You can also use these procedures to create external tables that you can use to query data.

The following `DBMS_CLOUD` procedures support loading data into the database from a directory:

- `DBMS_CLOUD.COPY_COLLECTION`
- `DBMS_CLOUD.COPY_DATA`

In addition, the following `DBMS_CLOUD` procedures support creating external tables from data in files in a directory.

- `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE`
- `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`
- `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE`

Depending on the procedure, use either the `file_uri_list` parameter or the `partitioning_clause` parameter to specify files in one or more directories.

To load data from a directory, you do not need to include the `credential_name` parameter, but you need `READ` object privileges on the directory.

For example, you can load data into an existing table using `DBMS_CLOUD.COPY_DATA`:

1. Use an existing directory, create a directory, or attach a network file system for the source files.

   See Create Directory in Autonomous Database for information on creating a local directory on your Autonomous Database instance.

   See Attach Network File System to Autonomous Database for information on attaching a network file system that contains the data you want to load.

2. Load data with a `DBMS_CLOUD` procedure.

   For example:

```
CREATE TABLE CHANNELS
   (channel_id CHAR(1),
    channel_desc VARCHAR2(20),
    channel_class VARCHAR2(20)
```

```
   );
/

BEGIN
 DBMS_CLOUD.COPY_DATA(
    table_name => 'CHANNELS',
    file_uri_list => 'MY_DIR:channels.txt',
    format => json_object('delimiter' value ',')
 );
END;
/
```

The parameters are:

- `table_name`: is the target table's name.
- `file_uri_list`: is a directory and file names specification for the source files you want to load.

  You can specify one directory and one or more file names or use a comma separated list of directories and file names. The format to specify a directory is:`'MY_DIR:filename.ext'`. By default the directory name `MY_DIR` is a database object and is case-insensitive. The file name is case sensitive.

  See Load Data from Directories in Autonomous Database for more details and example of specifying files in a directory. This also shows information on quoting the directory name to make the directory name case sensitive, and for information on using wildcards.
- `format`: defines the options you can specify to describe the format of the source file, including whether the file is of type text, ORC, Parquet, or Avro.

See COPY_DATA Procedure for more information.

See Attach Network File System to Autonomous Database for information on attaching network file systems.

Load Data from Local Files Using SQL*Loader

Instead of using SQL*Loader Oracle recommends loading data from the Cloud Object Storage for better performance and enhanced functionality.

For information on loading from Cloud Object Store, see Load Data from Files in the Cloud.

Depending on your workload type, note the following:

- **Data Warehouse:** If you use SQL*Loader to load data, note that Autonomous Database does not gather optimizer statistics for your load and you need to gather optimizer statistics manually as explained in Manage Optimizer Statistics on Autonomous Database. Autonomous Database gathers optimizer statistics automatically for tables loaded with direct path operations issued in SQL (direct path load operations that bypass the SQL data processing, such as SQL*Loader direct path, do not collect statistics).
- **Transaction Processing** or **JSON Database**: If you use SQL*Loader to load data, note that Autonomous Database does not gather optimizer statistics for your load and you need to gather optimizer statistics manually as explained in in Manage Optimizer Statistics on Autonomous Database or wait for the automatic statistic gathering task to kick in.

For detailed information on SQL*Loader see, Oracle Database Utilities.

Using Data Pipelines for Continuous Load and Export

Data pipelines allow you to repeatedly load data from object store or export data to object store. Load pipelines provide continuous incremental data loading from external sources (as data arrives on object store it is loaded to a database table). Export pipelines provide continuous incremental data exporting to object store (as new data appears in a database table it is exported to object store).

- About Data Pipelines on Autonomous Database
- Create and Configure Pipelines
  You can create one or more load or export pipelines. When you create a pipeline, you use parameters and set pipeline attributes to configure the pipeline.
- Test Pipelines
- Control Pipelines (Start, Stop, Drop, or Reset a Pipeline)
  After you create and test a pipeline, you control a pipeline by starting, stopping, or dropping the pipeline. You can also reset a pipeline.
- Monitor and Troubleshoot Pipelines
  All pipelines that are created are logged in the `DBMS_CLOUD_PIPELINE` views.
- Use Oracle Maintained Pipelines
  Autonomous Database provides built-in pipelines. These pipelines are preconfigured and can be started by the ADMIN user. Oracle Maintained pipelines are owned by ADMIN user.

About Data Pipelines on Autonomous Database

Autonomous Database data pipelines are either load pipelines or export pipelines.

Load pipelines provide continuous incremental data loading from external sources (as data arrives on object store it is loaded to a database table). Export pipelines provide continuous incremental data exporting to object store (as new data appears in a database table it is exported to object store). Pipelines use database scheduler to continuously load or export incremental data.

Autonomous Database data pipelines provide the following:

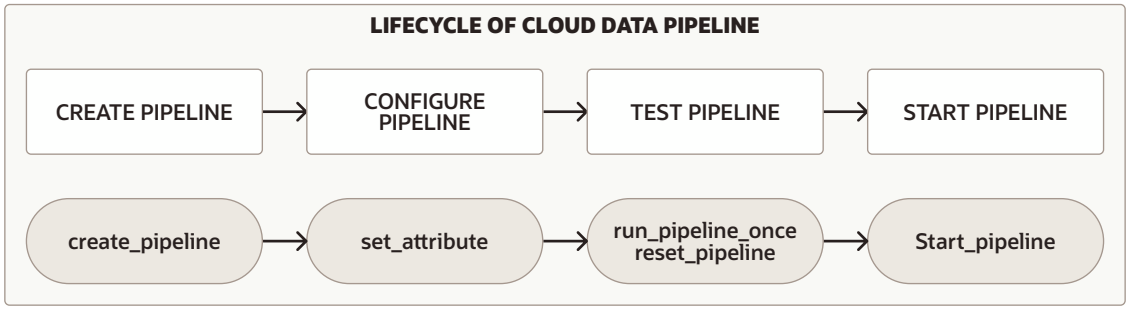
- **Unified Operations**: Pipelines allow you to quickly and easily load or export data and repeat these operations at regular intervals for new data. The `DBMS_CLOUD_PIPELINE` package provides a unified set of PL/SQL procedures for pipeline configuration and for creating and starting a scheduled job for load or export operations.
- **Scheduled Data Processing**: Pipelines monitor their data source and periodically load or export the data as new data arrives.
- **High Performance**: Pipelines scale data transfer operations with the available resources on your Autonomous Database. Pipelines by default use parallelism for all load or export operations, and scale based on the CPU resources available on your Autonomous Database, or based on a configurable priority attribute.
- **Atomicity and Recovery**: Pipelines guarantee atomicity such that files in object store are loaded exactly once for a load pipeline.
- **Monitoring and Troubleshooting**: Pipelines provide detailed log and status tables that allow you to monitor and debug pipeline operations.
- **Multicloud Compatible**: Pipelines on Autonomous Database support easy switching between cloud providers without application changes. Pipelines support all the credential and object store URI formats that Autonomous Database supports (Oracle Cloud

Infrastructure Object Storage, Amazon S3, Azure Blob Storage or Azure Data Lake Storage, Google Cloud Storage, and Amazon S3-Compatible object stores).

- About the Data Pipeline Lifecycle on Autonomous Database
  The `DBMS_CLOUD_PIPELINE` package provides procedures for creating, configuring, testing, and starting a pipeline. The pipeline lifecycle and procedures are the same for both load and export pipelines.
- About Load Pipelines on Autonomous Database
- About Export Pipelines on Autonomous Database
  Use an export pipeline for continuous incremental export of data from the database to object store. An export pipeline periodically identifies candidate data and uploads the data to object store.
- About Oracle Maintained Pipelines
  Autonomous Database provides built-in pipelines for exporting logs to object store. These pipelines are preconfigured and can be started by the ADMIN user.

About the Data Pipeline Lifecycle on Autonomous Database

The `DBMS_CLOUD_PIPELINE` package provides procedures for creating, configuring, testing, and starting a pipeline. The pipeline lifecycle and procedures are the same for both load and export pipelines.



For either pipeline type you perform the following steps to create and use a pipeline:

1. Create and configure the pipeline. See Create and Configure Pipelines for more information.
2. Test a new pipeline. See Test Pipelines for more information.
3. Start a pipeline. See Start a Pipeline for more information.

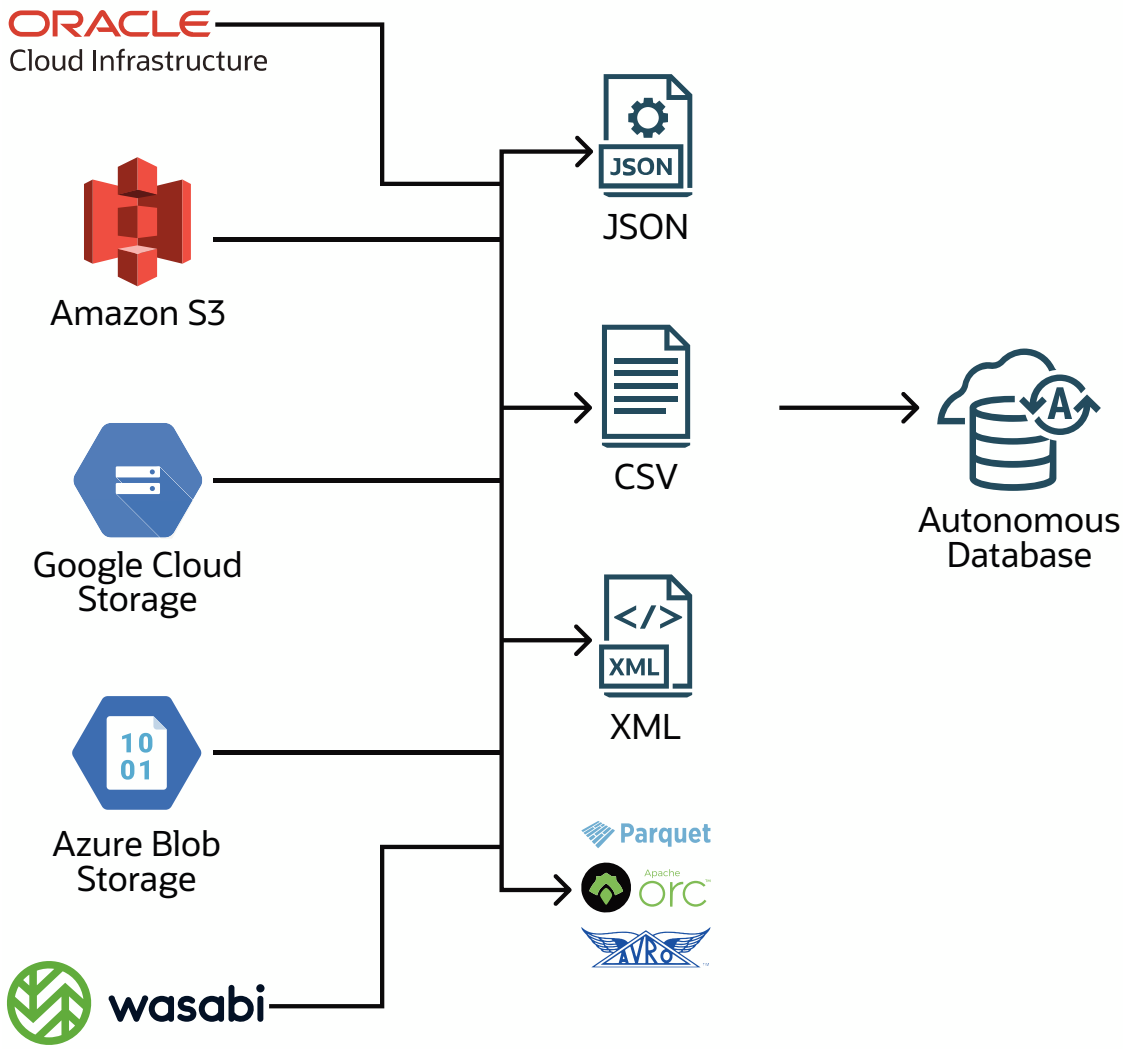
In addition, you can monitor, stop, or drop pipelines:

- While a pipeline is running, either during testing or during regular use after you start the pipeline, you can monitor the pipeline. See Monitor and Troubleshoot Pipelines for more information.
- You can stop a pipeline and later start it again, or drop a pipeline when you are finished using the pipeline. See Stop a Pipeline and Drop a Pipeline for more information.

About Load Pipelines on Autonomous Database

Use a load pipeline for continuous incremental data loading from external files in object store into a database table. A load pipeline periodically identifies new files in object store and loads the new data into the database table.
A load pipeline operates as follows (some of these features are configurable using pipeline attributes):

- Object store files are loaded in parallel into a database table.
  - A load Pipeline uses the object store file name to uniquely identify and load newer files.
  - Once a file in object store has been loaded in the database table, if the file content changes in object store, it will not be loaded again.
  - If the object store file is deleted, it does not impact the data in the database table.
- If failures are encountered, a load pipeline automatically retries the operation. Retries are attempted on every subsequent run of the pipeline's scheduled job.
- In cases where the data in a file does not comply with the database table, it is marked as `FAILED` and can be reviewed to debug and troubleshoot the issue.
  - If any file fails to load, the pipeline does not stop and continues to load the other files.
- Load pipelines support multiple input file formats, including: JSON, CSV, XML, Avro, ORC, and Parquet.



Migration from non-Oracle databases is one possible use case for a load pipeline. When you need to migrate your data from a non-Oracle database to Oracle Autonomous Database, you can extract the data and load it into Autonomous Database (Oracle Data Pump format cannot be used for migrations from non-Oracle databases). By using a generic file format such as CSV to export data from a non-Oracle database, you can save your data to files and upload

the files to object store. Next, create a pipeline to load the data to Autonomous Database. Using a load pipeline to load a large set of CSV files provides important benefits such as fault tolerance, and resume and retry operations. For a migration with a large data set you can create multiple pipelines, one per table for the non-Oracle database files, to load data into Autonomous Database.

About Export Pipelines on Autonomous Database

Use an export pipeline for continuous incremental export of data from the database to object store. An export pipeline periodically identifies candidate data and uploads the data to object store.

There are three export pipeline options (the export options are configurable using pipeline attributes):

- Export incremental results of a query to object store using a date or timestamp column as key for tracking newer data.
- Export incremental data of a table to object store using a date or timestamp column as key for tracking newer data.
- Export data of a table to object store using a query to select data without a reference to a date or timestamp column (so that the pipeline exports all the data that the query selects for each scheduler run).

Export pipelines have the following features (some of these are configurable using pipeline attributes):

- Results are exported in parallel to object store.
- In case of any failures, a subsequent pipeline job repeats the export operation.
- Export pipelines support multiple export file formats, including: CSV, JSON, Parquet, or XML.

About Oracle Maintained Pipelines

Autonomous Database provides built-in pipelines for exporting logs to object store. These pipelines are preconfigured and can be started by the ADMIN user.

The Oracle Maintained pipelines are:

- `ORA$AUDIT_EXPORT`: This pipeline exports the database audit logs to object store in JSON format and runs every 15 minutes after starting the pipeline (based on the `interval` attribute value).
- `ORA$APEX_ACTIVITY_EXPORT`: This pipeline exports the Oracle APEX workspace activity log to object store in JSON format. This pipeline is preconfigured with the SQL query for retrieving APEX activity records and runs every 15 minutes after starting the pipeline (based on the `interval` attribute value).

The Oracle Maintained pipelines are owned by the ADMIN user and attributes of Oracle Maintained Pipelines can be modified by the ADMIN user.

By default the Oracle Maintained Pipelines use `OCI$RESOURCE_PRINCIPAL` as the `credential_name`.

See Use Oracle Maintained Pipelines for more information.

Create and Configure Pipelines

You can create one or more load or export pipelines. When you create a pipeline, you use parameters and set pipeline attributes to configure the pipeline.

The options to create and configure a pipeline follow:

- **Load Pipeline**:
  - See Create and Configure a Pipeline for Loading Data.
- **Export Pipeline**:
  - Export incremental results of a query to object store using a date or timestamp column as key for tracking newer data. See Create and Configure a Pipeline for Export with Timestamp Column.
  - Export incremental data of a table to object store using a date or timestamp column as key for tracking newer data. See Create and Configure a Pipeline for Export with Timestamp Column.
  - Export data of a table to object store using a query to select data without a reference to a date or timestamp column (so that the pipeline exports all the data that the query selects for each scheduler run). See Create and Configure a Pipeline to Export Query Results (Without a Timestamp).
- Create and Configure a Pipeline for Loading Data
  You can create a pipeline to load data from external files in object store to tables in Autonomous Database.
- Create and Configure a Pipeline for Export with Timestamp Column
- Create and Configure a Pipeline to Export Query Results (Without a Timestamp)
  You can create an export pipeline to automatically export data from your Autonomous Database to object store. Using this export pipeline option you specify a SQL query that the pipeline runs periodically to export data to object store. You can use this export option to share the latest data from your Autonomous Database to object store for other applications to consume the data.

Create and Configure a Pipeline for Loading Data

You can create a pipeline to load data from external files in object store to tables in Autonomous Database.

A load pipeline consumes data placed on object store and loads it to a table in Autonomous Database. When you create a load pipeline, the pipeline runs at regular intervals to consume data placed on object store, when new data files arrive the pipeline loads the new data. You can also use a pipeline to reliably copy files, with resume and retry capabilities, from object store to a table on your database.

With a load pipeline, the pipeline package uses `DBMS_CLOUD.COPY_DATA` to load data.

On your Autonomous Database, either use an existing table or create the database table where you are loading data. For example:

```
CREATE TABLE EMPLOYEE
            (name     VARCHAR2(128),
             age      NUMBER,
             salary   NUMBER);
```

1. Create a pipeline to load data from object store.

```
BEGIN
     DBMS_CLOUD_PIPELINE.CREATE_PIPELINE(
        pipeline_name => 'MY_PIPE1',
        pipeline_type => 'LOAD',
        description   => 'Load metrics from object store into a table'
  );
END;
/
```

   See CREATE_PIPELINE Procedure for more information.

2. Create a credential object to access the object store that contains the files you are loading.

   You specify the credential for the pipeline source location with the attribute `credential_name`. If you do not supply a `credential_name` in the next step, the `credential_name` value is set to `NULL`. You can use the default `NULL` value when the `location` attribute is a public or pre-authenticated URL.

   See CREATE_CREDENTIAL Procedure for more information.

3. Set the pipeline attributes, including the required attributes: `location`, `table_name`, and `format`.

```
BEGIN
     DBMS_CLOUD_PIPELINE.SET_ATTRIBUTE(
       pipeline_name => 'MY_PIPE1',
       attributes    => JSON_OBJECT(
            'credential_name' VALUE 'OBJECT_STORE_CRED',
            'location' VALUE 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/',
            'table_name' VALUE 'employee',
            'format' VALUE '{"type":"json", "columnpath":["$.NAME",
"$.AGE", "$.SALARY"]}',
            'priority' VALUE 'HIGH',
            'interval' VALUE '20')
  );
END;
/
```

   The following attributes must be set to run a load pipeline:

   - `location`: Specifies the source file location on object store.
   - `table_name`: Specifies the table in your database where you are loading data. The `location` you specify is for one `table_name` per pipeline.
   - `format`: Describes the format of the data you are loading.

     See DBMS_CLOUD Package Format Options for more information.

   The `credential_name` is the credential you created in the previous step.

   The `priority` value determines the number of files loaded in parallel. A pipeline with a higher priority consumes more database resources and completes each run faster, as compared to running at a lower priority.

The `interval` value specifies the time interval in minutes between consecutive runs of a pipeline job. The default `interval` is 15 minutes.

See DBMS_CLOUD_PIPELINE Attributes for details on the pipeline attributes.

After you create a pipeline you can test the pipeline or start the pipeline:

- Test Pipelines
- Start a Pipeline

As an alternative, to set the format for JSON, you could use the following format:

```
BEGIN
    DBMS_CLOUD_PIPELINE.SET_ATTRIBUTE(
        pipeline_name   => 'MY_PIPE1',
        attribute_name  => 'format',
        attribute_value => JSON_OBJECT('type' value 'json', 'columnpath'
value '["$.NAME", "$.AGE", "$.SALARY"]')
    );
END;
/
```

Create and Configure a Pipeline for Export with Timestamp Column

You can create an export pipeline to automatically export time-series data from your Autonomous Database to object store.

Using this export pipeline option you specify a table or SQL query and a column with a timestamp that the pipeline uses to keep track of the time of the last upload. You can use an export pipeline to share data for consumption by other applications or to save data to object store.

With an export pipeline, the pipeline package uses `DBMS_CLOUD.EXPORT_DATA` to export data.

An export pipeline exports data from your Autonomous Database to object store. When you create an export pipeline, the pipeline runs at regular intervals and places data on object store.

1. Create a pipeline to export data to object store.

   ```
   BEGIN
        DBMS_CLOUD_PIPELINE.CREATE_PIPELINE(
           pipeline_name=>'EXP_PIPE1',
           pipeline_type=>'EXPORT',
           description=>'Export time series metrics to object store');
   END;
   /
   ```

   See CREATE_PIPELINE Procedure for more information.

2. Create a credential object to access the destination object store location where you are exporting data files.

   You specify the credential for the pipeline destination location with the attribute `credential_name`. If you do not supply a `credential_name` in the next step, the `credential_name` value is set to `NULL`. You can use the default `NULL` value when the `location` attribute is a public or pre-authenticated URL.

   See CREATE_CREDENTIAL Procedure for more information.

3. Set the export pipeline attributes.

   When you specify a `table_name` parameter, table rows are exported to object store. When you specify a `query` parameter, the query specifies a `SELECT` statement so that only the required data is exported to object store.

   - Using using a `table_name` parameter:

```
BEGIN
     DBMS_CLOUD_PIPELINE.SET_ATTRIBUTE(
       pipeline_name => 'EXP_PIPE1',
       attributes    => JSON_OBJECT('credential_name' VALUE
'OBJECT_STORE_CRED',
          'location' VALUE 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/',
          'table_name' VALUE 'metric_table',
          'key_column' VALUE 'metric_time',
          'format' VALUE '{"type": "json"}',
          'priority' VALUE 'MEDIUM',
          'interval' VALUE '20')
  );
END;
/
```

   - Using a `query` parameter:

```
BEGIN
     DBMS_CLOUD_PIPELINE.SET_ATTRIBUTE(
       pipeline_name => 'EXP_PIPE1',
       attributes    => JSON_OBJECT('credential_name' VALUE
'OBJECT_STORE_CRED',
           'location' VALUE 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/',
          'query' VALUE 'SELECT * from metrics_table',
          'key_column' VALUE 'metric_time',
          'format' VALUE '{"type": "json"}',
          'priority' VALUE 'MEDIUM',
          'interval' VALUE '20')
  );
END;
/
```

   Where the `credential_name` is the credential you created in the previous step.

   The following attributes must be set to run an export pipeline:

   - `location`: Specifies the destination object store location. The `location` you specify is for one `table_name` per pipeline.
   - `table_name`: Specifies the table in your database that contains the data you are exporting (either the `table_name` parameter or the `query` parameter is required).
   - `query`: Specifies the query to run in your database that provides the data you are exporting (either the `table_name` parameter or the `query` parameter is required).
   - `format`: Describes the format of the data you are exporting.

     See DBMS_CLOUD Package Format Options for EXPORT_DATA for more information.

The `priority` value determines the degree of parallelism for fetching data from the database.

The `interval` value specifies the time interval in minutes between consecutive runs of a pipeline job. The default `interval` is 15 minutes.

See DBMS_CLOUD_PIPELINE Attributes for details on the pipeline attributes.

After you create a pipeline you can test the pipeline or start the pipeline:

- Test Pipelines
- Start a Pipeline

Create and Configure a Pipeline to Export Query Results (Without a Timestamp)

You can create an export pipeline to automatically export data from your Autonomous Database to object store. Using this export pipeline option you specify a SQL query that the pipeline runs periodically to export data to object store. You can use this export option to share the latest data from your Autonomous Database to object store for other applications to consume the data.

An export pipeline exports data from your Autonomous Database to object store. When you create an export pipeline, the pipeline runs at regular intervals and places data on object store.

1. Create a pipeline to export data to object store.

```
BEGIN
     DBMS_CLOUD_PIPELINE.CREATE_PIPELINE(
        pipeline_name=>'EXP_PIPE2',
        pipeline_type=>'EXPORT',
        description=>'Export query results to object store.');
END;
/
```

   See CREATE_PIPELINE Procedure for more information.

2. Create a credential object to access the destination object store location where you are exporting data files.

   You specify the credential for the pipeline destination location with the attribute `credential_name`. If you do not supply a `credential_name` in the next step, the `credential_name` value is set to `NULL`. You can use the default `NULL` value when the `location` attribute is a public or pre-authenticated URL.

   See CREATE_CREDENTIAL Procedure for more information.

3. Set the export pipeline attributes.

```
BEGIN
     DBMS_CLOUD_PIPELINE.SET_ATTRIBUTE(
       pipeline_name => 'EXP_PIPE2',
       attributes    => JSON_OBJECT(
          'credential_name' VALUE 'OBJECT_STORE_CRED',
          'location' VALUE 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/',
          'query' VALUE 'SELECT * FROM table_name',
          'format' VALUE '{"type": "json"}',
          'priority' VALUE 'MEDIUM',
          'interval' VALUE '20')
```

```
  );
END;
/
```

Where the `credential_name` is the credential you created in the previous step.

The following attributes must be set to run an export pipeline:

- `location`: Specifies the destination object store location.
- `query`: Specifies the query to run in your database that provides the data you are exporting.
- `format`: Describes the format of the data you are exporting.

  See DBMS_CLOUD Package Format Options for EXPORT_DATA for more information.

The `priority` value determines the degree of parallelism for fetching data from the database.

The `interval` value specifies the time interval in minutes between consecutive runs of a pipeline job. The default `interval` is 15 minutes.

See DBMS_CLOUD_PIPELINE Attributes for details on the pipeline attributes.

After you create a pipeline you can test the pipeline or start the pipeline:

- Test Pipelines
- Start a Pipeline

Test Pipelines

Use `RUN_PIPELINE_ONCE` to run a pipeline once on-demand without creating a scheduled job.

`RUN_PIPELINE_ONCE` is useful for testing a pipeline before you start the pipeline. After you run a pipeline once to test the pipeline and check that it is working as expected, use `RESET_PIPELINE` to reset the pipeline's state (to the state before you ran `RUN_PIPELINE_ONCE`).

1. Create a pipeline.

   See Create and Configure Pipelines for more information.

2. Run a pipeline once to test the pipeline.

   ```
   BEGIN
       DBMS_CLOUD_PIPELINE.RUN_PIPELINE_ONCE(
           pipeline_name => 'MY_PIPE1'
   );
   END;
   /
   ```

   See RUN_PIPELINE_ONCE Procedure for more information.

3. Perform any required checks to verify the pipeline is operating as expected.

   See Monitor and Troubleshoot Pipelines for more information.

4. Reset the pipeline.

```
BEGIN
   DBMS_CLOUD_PIPELINE.RESET_PIPELINE(
     pipeline_name => 'MY_PIPE1',
     purge_data => TRUE
);
END;
/
```

See RESET_PIPELINE Procedure for more information.

Control Pipelines (Start, Stop, Drop, or Reset a Pipeline)

After you create and test a pipeline, you control a pipeline by starting, stopping, or dropping the pipeline. You can also reset a pipeline.

- Start a Pipeline
- Stop a Pipeline
- Drop a Pipeline
- Reset a Pipeline

Start a Pipeline

After you create a pipeline, you can start the pipeline.

When a pipeline is started the pipeline runs continuously in a scheduled job. The pipeline's scheduled job repeats, either by default every 15 minutes or at the interval you set with the `interval` attribute.

1. Start a pipeline.

```
BEGIN
  DBMS_CLOUD_PIPELINE.START_PIPELINE(
      pipeline_name => 'EMPLOYEE_PIPELINE'
  );
END;
/
```

By default, a pipeline job begins immediately, as soon as the pipeline is started. To start a pipeline job at a later time, specify a valid future date or timestamp using the `start_date` parameter.

See START_PIPELINE Procedure for more information.

2. Verify that the pipeline is started.

For example:

```
SELECT pipeline_name, status from USER_CLOUD_PIPELINES
   WHERE pipeline_name = 'EMPLOYEE_PIPELINE';

PIPELINE_NAME            STATUS
```

```
------------------------ -------
EMPLOYEE_PIPELINE        STARTED
```

Stop a Pipeline

Use `STOP_PIPELINE` to stop a pipeline. When a pipeline is stopped, no future jobs are scheduled for the pipeline.

By default currently running jobs complete when you stop a pipeline. Set the `force` parameter to `TRUE` to terminate any running jobs and stop the pipeline immediately.

1. Stop a pipeline.

```
BEGIN
  DBMS_CLOUD_PIPELINE.STOP_PIPELINE(
      pipeline_name => 'EMPLOYEE_PIPELINE'
  );
END;
/
```

   See STOP_PIPELINE Procedure for more information.

2. Verify that the pipeline is stopped.

```
SELECT pipeline_name, status from USER_CLOUD_PIPELINES
   WHERE pipeline_name = 'EMPLOYEE_PIPELINE';

PIPELINE_NAME            STATUS
------------------------ -------
EMPLOYEE_PIPELINE        STOPPED
```

See STOP_PIPELINE Procedure for more information.

Drop a Pipeline

The procedure `DROP_PIPELINE` drops an existing pipeline.

If a pipeline has been started, it must be stopped before the pipeline can be dropped. See STOP_PIPELINE Procedure for more information.

In order to drop a pipeline that is started, set the `force` parameter to `TRUE` to terminate any running jobs and drop the pipeline immediately

1. Drop a pipeline.

```
BEGIN
  DBMS_CLOUD_PIPELINE.DROP_PIPELINE(
      pipeline_name => 'EMPLOYEE_PIPELINE'
  );
END;
/
```

2. Verify the pipeline is dropped.

```
SELECT pipeline_name, status from USER_CLOUD_PIPELINES
   WHERE pipeline_name = 'EMPLOYEE_PIPELINE';

No rows selected
```

See DROP_PIPELINE Procedure for more information.

Reset a Pipeline

Use the reset pipeline operation to clear the record of the pipeline to the initial state.

> **Note:**
>
> You can optionally use reset pipeline to purge data in the database table associated with a load pipeline or to remove files in object store for an export pipeline. Usually this option is used when you are testing a pipeline during pipeline development.

Reset pipeline operates as follows:

- **Load Pipeline**: For a load pipeline, resetting the pipeline clears the record of the files being loaded by the pipeline. When you call either `START_PIPELINE` or `RUN_PIPELINE_ONCE` after resetting a load pipeline, the pipeline repeats the data load and includes all the files present in the object store location.

  When `purge_data` is set to `TRUE`, the procedure truncates the data in the database table.

- **Export Pipeline**: For an export pipeline, resetting the pipeline clears the last tracked data in the database table. When you call either `START_PIPELINE` or `RUN_PIPELINE_ONCE` after resetting an export pipeline, the pipeline repeats exporting data from the table or query.

  When `purge_data` set to `TRUE`, the procedure deletes existing files in the object store location specified with the `location` attribute.

To reset a pipeline:

1. Stop the pipeline you want to reset.

   A data pipeline must be in stopped state to reset it. See STOP_PIPELINE Procedure for more information.

2. Reset the pipeline.

```
BEGIN
     DBMS_CLOUD_PIPELINE.RESET_PIPELINE(
        pipeline_name => 'EMPLOYEE_PIPELINE',
        purge_data => TRUE);
END;
/
```

   Only use the `purge_data` parameter with value `TRUE` if you want to clear data in your database table, for a load pipeline, or clear files in object store for an export pipeline. Usually this option is used when you are testing a pipeline during pipeline development.

   See RESET_PIPELINE Procedure for more information.

Monitor and Troubleshoot Pipelines

All pipelines that are created are logged in the `DBMS_CLOUD_PIPELINE` views.

View Pipeline Status Information

Check pipeline status and obtain other pipeline information using the `USER_CLOUD_PIPELINES` or `DBA_CLOUD_PIPELINES` views. For example, the following `SELECT` statement with a `WHERE` clause predicate on the `pipeline_name` shows that `MY_TREE_DATA` is a load pipeline and the pipeline is started:

```
SELECT pipeline_name, pipeline_type, status FROM USER_CLOUD_PIPELINES
   WHERE pipeline_name = 'MY_TREE_DATA';


PIPELINE_NAME PIPELINE_TYPE STATUS
------------- ------------- -------
MY_TREE_DATA  LOAD          STARTED
```

See DBMS_CLOUD_PIPELINE Views for more information.

View Pipeline Attributes

The pipeline attributes can be monitored by querying the `USER_CLOUD_PIPELINE_ATTRIBUTES` or `DBA_CLOUD_PIPELINE_ATTRIBUTES` views. Query these views to see pipeline attribute information.

For example:

```
SELECT pipeline_name, attribute_name, attribute_value FROM
user_cloud_pipeline_attributes
     WHERE pipeline_name = 'MY_TREE_DATA';

PIPELINE_NAME ATTRIBUTE_NAME
ATTRIBUTE_VALUE

------------- ---------------
--------------------------------------------------------------------------------
---
MY_TREE_DATA  credential_name
DEF_CRED_OBJ_STORE

MY_TREE_DATA  format          {"type":
"csv"}
MY_TREE_DATA  interval
20

MY_TREE_DATA  location        https://objectstorage.us-
ashburn-1.oraclecloud.com/n/namespace/b/treetypes/o/
MY_TREE_DATA  priority
high

MY_TREE_DATA  table_name      TREES
```

See DBMS_CLOUD_PIPELINE Views for more information.

View Pipeline History

The USER_CLOUD_PIPELINE_HISTORY and DBA_CLOUD_PIPELINE_HISTORY views show the state of running jobs. Use the pipeline history views to help you monitor the health of a pipeline and detect failures in a running pipeline.

For example:

```
SELECT pipeline_id, pipeline_name, status, error_message  FROM
user_cloud_pipeline_history
     WHERE pipeline_name = 'MY_TREE_DATA';

PIPELINE_ID PIPELINE_NAME STATUS    ERROR_MESSAGE
----------- ------------- --------- -------------
          7  MY_TREE_DATA SUCCEEDED
```

See DBMS_CLOUD_PIPELINE Views for more information.

Pipeline Status Table: Additional Monitoring for Load Pipelines

The pipeline status table shows each file name and its status for a load pipeline. The STATUS_TABLE column in DBA_CLOUD_PIPELINES and USER_CLOUD_PIPELINES shows the status table name.

For example, the following SELECT statement with a WHERE clause predicate on the pipeline_name shows the status table name for a pipeline:

```
SELECT pipeline_name, status_table FROM user_cloud_pipelines
   WHERE pipeline_name = 'MY_TREE_DATA';

PIPELINE_NAME STATUS_TABLE
------------- --------------------
MY_TREE_DATA  PIPELINE$9$41_STATUS
```

View the status table to see information about the pipeline, including the following:

- The relevant error number and error message are recorded in the status table if an operation on a specific file fails.
- For completed pipeline operations, the time needed for each operation can be calculated using the reported START_TIME and END_TIME.

For example the following shows that the load operation for two files failed and one completed:

```
SELECT id, name, status, error_code, error_message, sid FROM
PIPELINE$9$41_STATUS;

ID NAME       STATUS    ERROR_CODE ERROR_MESSAGE                    SID
-- ---------- --------- ---------- -------------------------------- -----
 1 trees1.txt FAILED         30653 ORA-30653: reject limit reached  18070
 2 trees2.txt FAILED         30653 ORA-30653: reject limit reached  18070
 3 trees3.txt COMPLETED                                             18070
```

Pipelines for loading data, where the `pipeline_type` is `'LOAD'`, reserve an `ID` that is shown in `USER_LOAD_OPERATIONS` and in `DBA_LOAD_OPERATIONS`. The `ID` value in these views maps to the pipeline's `OPERATION_ID` in `USER_CLOUD_PIPELINES` and `DBA_CLOUD_PIPELINES`.

To obtain more information for a load pipeline, query the pipeline's `OPERATION_ID`:

```
SELECT PIPELINE_NAME, OPERATION_ID FROM USER_CLOUD_PIPELINES
     WHERE PIPELINE_NAME = 'MY_TREE_DATA';

PIPELINE_NAME OPERATION_ID
------------- ------------
MY_TREE_DATA            41
```

Next, query either `USER_LOAD_OPERATIONS` or `DBA_LOAD_OPERATIONS` with a `WHERE` clause predicate on the `ID` column (using the `OPERATION_ID` value).

For example:

```
SELECT ID, TYPE, LOGFILE_TABLE, BADFILE_TABLE, STATUS_TABLE FROM
USER_LOAD_OPERATIONS
     WHERE ID = 41;

ID TYPE     LOGFILE_TABLE     BADFILE_TABLE     STATUS_TABLE
-- -------- ----------------- ----------------- --------------------
41 PIPELINE PIPELINE$9$41_LOG PIPELINE$9$41_BAD PIPELINE$9$41_STATUS
```

This query shows `ID`, `TYPE`, `LOGFILE_TABLE`, `BADFILE_TABLE` if it exists, and the `STATUS_TABLE`. You can view these tables for additional pipeline load information.

Pipeline Status Table Details

Column	Datatype	Description
`ID`	`NUMBER`	Unique number assigned to the pipeline.
`NAME`	`VARCHAR2(4000)`	Name of the pipeline.
`BYTES`	`NUMBER`	Bytes
`CHECKSUM`	`VARCHAR2(128)`	Checksum
`LAST_MODIFIED`	`TIMESTAMP(6) WITH TIME ZONE`	Last modification time for the pipeline.
`STATUS`	`VARCHAR2(30)`	The `STATUS` value is one of: • `COMPLETED`: The file operation completed successfully. • `FAILED`: The file operation failed, a retry may be attempted two times. • `PENDING`: The file operation has not yet started. • `RUNNING`: The file operation is currently in progress. • `SKIPPED`: The file operation skipped.
`ERROR_CODE`	`NUMBER`	Error code
`ERROR_MESSAGE`	`VARCHAR2(4000)`	Error message
`START_TIME`	`TIMESTAMP(6) WITH TIME ZONE`	Start time for the pipeline.

Column	Datatype	Description
END_TIME	TIMESTAMP(6) WITH TIME ZONE	End time for the pipeline.
SID	NUMBER	The session SID andSERIAL# indicate the job session that was running the pipeline load operation.
SERIAL#	NUMBER	The session SID andSERIAL# indicate the job session that was running the pipeline load operation.
ROWS_LOADED	NUMBER	Number of rows loaded.
OPERATION_ID	NUMBER	Reserved for future use.

Pipeline Log File and Bad File Tables

To obtain the log file and bad file names for a load pipeline, query the pipeline's OPERATION_ID. For example:

```
SELECT PIPELINE_NAME, OPERATION_ID FROM USER_CLOUD_PIPELINES
     WHERE PIPELINE_NAME = 'MY_TREE_DATA';

PIPELINE_NAME OPERATION_ID
------------- ------------
MY_TREE_DATA            41
```

Next, query either USER_LOAD_OPERATIONS or DBA_LOAD_OPERATIONS with a WHERE clause predicate on the ID column (using the OPERATION_ID value).

For example:

```
SELECT ID, TYPE, LOGFILE_TABLE, BADFILE_TABLE, STATUS_TABLE FROM
USER_LOAD_OPERATIONS
     WHERE ID = 41;

ID TYPE     LOGFILE_TABLE     BADFILE_TABLE     STATUS_TABLE
-- -------- ----------------- ----------------- --------------------
41 PIPELINE PIPELINE$9$41_LOG PIPELINE$9$41_BAD PIPELINE$9$41_STATUS
```

This query shows ID, TYPE, LOGFILE_TABLE, BADFILE_TABLE if it exists, and the STATUS_TABLE. You can view these tables for additional pipeline load information.

View the pipeline log file table to see a complete log of the pipeline's load operations.

For example:

```
SELECT * FROM PIPELINE$9$41_LOG;
```

View the pipeline bad file table to see details on input format records with errors. The bad file table show information for the rows reporting errors during loading. Depending on the errors shown in the log file table and the rows shown in the pipeline's bad file table, you might be able to correct the errors either by modifying the pipeline format attribute options, or by modifying the data in the file you are loading.

For example:

```
SELECT * FROM PIPELINE$9$41_BAD;
```

See Monitor and Troubleshoot Loads for more information.

Use Oracle Maintained Pipelines

Autonomous Database provides built-in pipelines. These pipelines are preconfigured and can be started by the ADMIN user. Oracle Maintained pipelines are owned by ADMIN user.

The Oracle Maintained pipelines are:

- `ORA$AUDIT_EXPORT`: This pipeline exports the database audit logs to object store in JSON format and runs every 15 minutes after starting the pipeline (based on the `interval` attribute value).
- `ORA$APEX_ACTIVITY_EXPORT`: This pipeline exports the Oracle APEX workspace activity log to object store in JSON format. This pipeline is preconfigured with the SQL query for retrieving APEX activity records and runs every 15 minutes after starting the pipeline (based on the `interval` attribute value).

To configure and start an Oracle Managed pipeline:

1. Determine the Oracle Managed Pipeline you want to use: `ORA$AUDIT_EXPORT` or `ORA$APEX_ACTIVITY_EXPORT`.
2. Set the `credential_name` and `location` attributes.

   For example:

   ```
   BEGIN
     DBMS_CLOUD_PIPELINE.SET_ATTRIBUTE(
       pipeline_name => 'ORA$AUDIT_EXPORT',
       attribute_name => 'credential_name',
       attribute_value => 'DEF_CRED_OBJ_STORE'
     );
     DBMS_CLOUD_PIPELINE.SET_ATTRIBUTE(
       pipeline_name => 'ORA$AUDIT_EXPORT',
       attribute_name => 'location',
       attribute_value => 'https://objectstorage.us-
   phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/'
     );
   END;
   /
   ```

   The log data from database is exported to the object store location you specify.

   See SET_ATTRIBUTE Procedure for more information.

   By default the Oracle Maintained Pipelines use `OCI$RESOURCE_PRINCIPAL` as the `credential_name`. See Enable Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.
3. Optionally, set the `interval`, `format`, or `priority` attributes.

   See SET_ATTRIBUTE Procedure for more information.
4. Start the pipeline.

See START_PIPELINE Procedure for more information.

The Data Load Page

Use the Data Load page to make more data available to your Oracle Autonomous Database. You can load data from files or databases, from links to external databases or cloud storage files, or from a live feed of data from cloud storage.

See The Data Load Page for more information.

- The Data Load Page
  Use the Data Load dashboard page to make more data available to your Oracle Autonomous Database. You can load data from files or databases, from links to external databases or cloud storage files, or from a live feed of data from cloud storage.

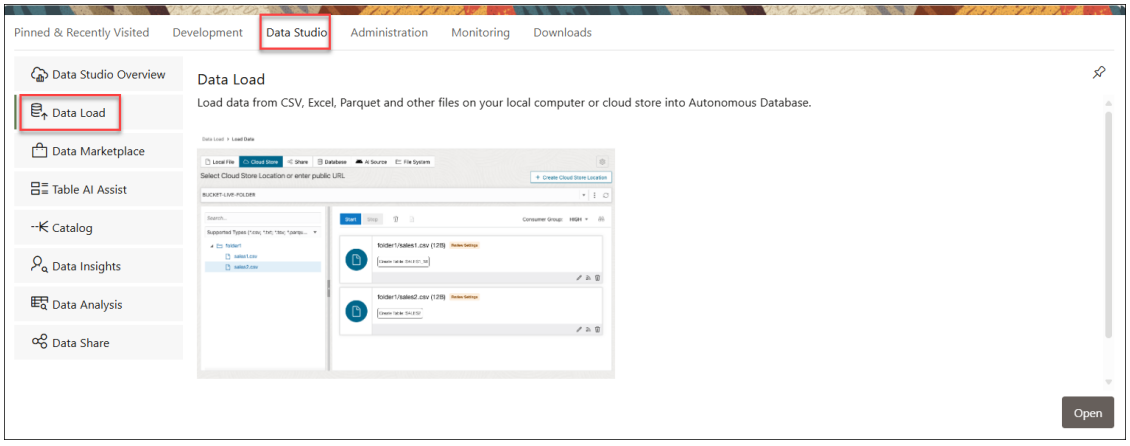
The Data Load Page

Use the Data Load dashboard page to make more data available to your Oracle Autonomous Database. You can load data from files or databases, from links to external databases or cloud storage files, or from a live feed of data from cloud storage.

From the Data Load page, you can also explore the data in your autonomous database and manage your cloud storage locations.

> **Note:**
>
> If you do not see the Data Load card then your database user is missing the required DWROLE role.

To reach the Data Load page, click the **Data Studio** tab in the Database Actions page and select the **Data Load** pane.



You can alternatively click the Selector icon and select **Data Load** from the Data Studio menu in the navigation pane.

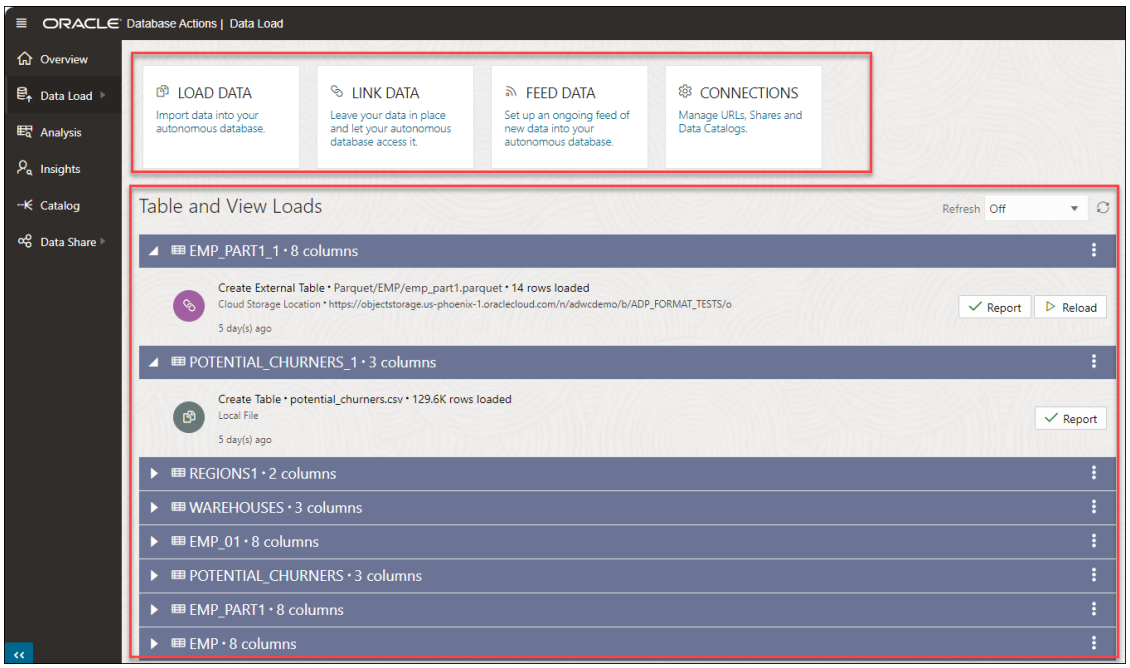
The Data Load Dashboard provides a central location for all data load and live feed tasks. The various icons on the dashboard, provide access to loading data, linking data, setting up an ongoing feed, checking data load jobs, and managing connections.

You can view the following tiles on the top section of the dashboard:

- Load Data
- Link Data
- Feed Data
- Connections

The bottom section of the dashboard displays the tables, views, and life feed you load into the tool.

The following image illustrates the **Data Load Dashboard**.



After drilling down onto a table, in the **Table and View Loads** section, you can view the details of the table load such as the name of the table, the number of rows and columns in the table, and the time you load the table.

Tables which are not active for last 7 days are collapsed.

You can also view the source location of the data you load. Click the location to view objects in the location.

Click the **Actions** icon beside the table name and select **Table** to view the following properties of the table:

- View Details
- Gather Statistics
- Register to Cloud Link
- Publish Listings to Marketplace
- Create Analytic View
- The Table AI Assist Tool
- Export Data to Cloud

- The Data Analysis Tool
- Edit
- **Drop** : Select **Drop** to delete the table.

Select **Report** to view a report of total rows processed successfully and failed for the selected table.

Select **Reload** to reload the source of the data load job. Clicking Reload takes you to the **Link Data** page where you can link your data source to the Autonomous Database.

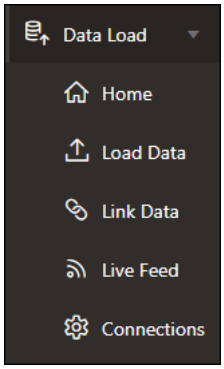
> **Note:**
>
> For tables you load via Cloud Store, you can view **Create Live Feed** icon. See Create Live Feed from Data Load. After the creation of live feed is successful, you will view a Live Feed badge besides the table load header.

In case of errors, you can reload data with suggested fixes.

To load or create links to data or create live table feeds, on the Data Load dashboard, select a combination of an operation and a data source location. The following table lists the operations and the source locations that support those operations.

Operation	Source Location	Description
Load data	Local file Database Cloud storage Directory Share AI Sources	Load data from files on your local device, from remote databases, or from cloud storage into tables in your Oracle Autonomous Database.
Link data	Database Objects in Oracle Cloud Infrastructure (OCI) Catalog	Create external tables or views in your Oracle Autonomous Database that link to data in cloud storage or remote databases. Changes to the source data automatically appear in the target objects.
Feed data	Cloud storage	Set up a feed of data from a cloud storage bucket into a table. Changes to the source data load into the target table as scheduled or on demand.
Connections	Cloud storage Data Catalog links Share Providers	Create a cloud storage location. Subscribe to Share Provider and Register Data Catalog. You can manage cloud storage locations, Data Catalog and Shares.

The left navigation pane of the Data Load Dashboard provides links to the Data Load features and pages. When you select a link from the left navigation pane, the corresponding page appears in the target content area.



Topics:

- Data Studio Settings
  This topic explains the Data Studio settings on the Data Studio tool.
- Managing Connections
  Connections that are established from the Data Studio to the cloud, catalogs and shares are listed on this page.
- Loading Data
  You can load data from files on your local device, from remote databases, or from cloud storage buckets from directories and share providers. The file formats that you can load are CSV, XLS, XLSX, TSV, TXT, XML, JSON, GEOJSON, AVRO, PARQUET, GZ, GZIP, ZIP, PDF, PNG, JPG, JPEG and TIFF.
- Linking Data
  You can link to data in remote databases or in cloud storage buckets.
- Feeding Data
  You can run a live table feed on demand, on a schedule, or as the result of a notification.
- DBMS_LIVE_FEED Package Reference
  This chapter provides information about the packages you use with the Data Load Tool in Data Studio. The `DBMS_LIVE_FEED` topic also covers the procedures included in the `DBMS_LIVE_FEED` package.

Data Studio Settings

This topic explains the Data Studio settings on the Data Studio tool.

This settings on the **Data Studio** navigation menu and the Data studio overview page enables you to use the associated Cloud Service, configure AI profile and configure SMTP.

This helps in defining the data during the load. This setting sets a variety of general preferences like setting default credentials to define the data you load, configure SMTP for data sharing, configure AI profile and configure Marketplace cloud link access.

To set the Data Studio Settings, on the **General** tab of the wizard, specify the email address to configure your profile for Marketplace and Sharing.

On the **Cloud Services** tab of the wizard, under **Cloud Object Store Access**, specify the following field values:

- **OCI Credential**: Selecting an OCI credential enables you to discover the OCI buckets. Click **+** to create OCI Credentials.

- **OCI Compartment**: This drop-down field lists the buckets present in your compartment for loading and linking from Oracle cloud services. If you select any other cloud service provider apart from Oracle, this field displays **None** by default.
- **AWS Credential**: Selecting an AWS credential in Data Studio Settings enables you to discover the S3 buckets. Click **+** to create credentials for Amazon S3 cloud storage.
- **Azure Credential**: Selecting an Azure credential in preferences enables you to discover the Azure BLOB storage. Click **+** to create credentials for Microsoft Azure BLOB storage and Azure Data Lake Storage.
  Click **Check Object Store Access** to check if you have access to the object store. You will view a Policy Check wizard with the following bucket lists:
  - OCI list buckets
  - AWS list buckets
  - Azure list buckets

  The Object Storage Listing Policy checker evaluates if you have access to the buckets. If yes, the checker box is highlighted in green with a check sign, else the checker box displays the policy highlighted in red with a red cross.

Under **Cloud Service Provider for AI**: Select any of the following options from the drop-drown field:

- Oracle Cloud Interface (OCI)
- Amazon Web Services (AWS)

Click **Check Cloud Service Access**.

You will view the following options in the Cloud Services Policy Checker dialog box. The AI cloud policies listed below must be checked against the listed policies and approved by the Cloud Service vendors for support. The Data Studio tool will test these services and inform you if the data you load in the tool is able to access the services.

- **Language**- Checks if the Data Studio tool has accessed the Cloud's language processing abilities. See Use OCI Language Service Capabilities in Data Studio for more information on this service.
- **Document Understanding**- Checks if the data you load into the Data Studio tool has accessed the Cloud's Document Understanding AI ability. See Use OCI Document Understanding to extract Tables from Images for more information on this service.
- **Translation**- Checks if the Data Studio tool accesses the OCI Translation feature. See Perform Text Translation for more information on this service.

On the **AI Profile** tab of the Data Studio Settings wizard, you can configure your AI preferences.

1. Select **AI Profile (Experimental)** from the drop-down use this AI profile for facilitating and configuring the translation of natural language prompts to SQL statements.
   - To utilize the AI services in Data Studio, you must have either an **OCI AI** or **OpenAI** or **Cohere** or **Azure OpenAI service** account.
   - You must also have access to `DBMS_CLOUD_AI` package to set your AI profile.
   - If you have Select AI set up using either **OpenAI** or **Cohere** or **Azure OpenAI** service, you can enhance your data by loading small reference data tables that are generated by large language models. Try out the suggested prompts or come up with your own to generate data to load into the autonomous database.
     You can also use this AI profile to generate Marketplace dataset description and tags. This AI profile is used to generate table description in Catalog App.

See Use DBMS_CLOUD_AI to Configure AI Profiles to create and configure your AI profile.

2. Click **Check AI Access**.

   The AI Policy checker checks for the following service:

   - **Generative AI**- The Data Studio tool checks if it has access to the Generative AI service.

3. Click **View Profile** to view the details related to the AI profile you use. You can view the **AI Provider**, **AI Model** and the **Credential** you use to create the AI profile.

4. Click **Create Profile** to create a new AI profile. Specify the following field values:

   - **Profile Name**: Enter the name of the AI profile. For example, *OCI AI*.
   - **AI Provider**: Select any of the following providers from the drop-down field:
     - *OCI AI*
     - *AI Azure*
     - *Cohere*
   - **Model**: Select any of the following models from the drop-down field:
     - *Llama*
     - *Cohere*
   - **Credential**: Select any of the credentials from the drop-down to create an AI profile.
   - Click **Show Code** to view the equivalent PL/SQL code to create AI profile.

     ```
     BEGIN
         C##CLOUD$SERVICE.DBMS_CLOUD_AI.CREATE_PROFILE(
             PROFILE_NAME => 'new_profile',
             ATTRIBUTES =>
     '{"credential_name":"my_credential_name","model":"","provider":""}');
     END;
     /
     ```

   Click **Create AI Profile**.

You have successfully set and configured your AI profile in the Data Studio Settings wizard.

On the SMTP tab of the Data Studio Preferences wizard, you can configure email service like a local email client on your system (e.g., Thunderbird).

> **Note:**
>
> You must configure SMTP only once and then the system uses the saved configuration from that point forward.

Specify the following field values in the SMTP tab of the wizard:

1. **Server host**: Enter the endpoint used to send the email. For example, *internal-mailrouter.oracle.com*.
2. **Server port**: Enter the SMTP port used to accept an email. Email Delivery supports TLS on port 25. Sender:
3. Enter the email address of the sender. For example, *oaapgens_us@oracle.com*.

4. **Server Encryption**: This field indicates if TLS, the standard means of performing encryption in transit for emails, is being used. Providers must encrypt emails while it's in transit to the Oracle Cloud Infrastructure Email Delivery service. Encrypted emails are protected from being read during transit. If there is no encryption, enter, *None*.
5. Select credential to use for SMTP connection from the drop-down. If there is no available credential in the drop-down, you can create credential by clicking Create Credential icon.

Click **Save** to save your changes and **Test** the test the ongoing process of setting preferences.

On the **Marketplace** tab, you can set the configurations you require for the Data Marketplace tool.

Access to data in Data Marketplace is restricted based on user authorizations. As non admin users you have permission to access the Data Marketplace tool based on the Data Studio Settings configuration. You can view the following options:

- **Restricted, Requires Authorization**:
  Selecting this option enables you to publish restricted cloud links and, consequently, restricted marketplace listings. You will have the ability to grant and revoke access to these listings to specific OCIDs.

- **Publish**:
  This option specifies who and from where you are allowed to access the registered table.

  You can choose from any of the following available options:

  – **TENANCY**: You can grant remote data access to any resource, tenancy, compartment, or database in the tenancy of the Autonomous Database instance that is registering the data set. This scope is more restrictive than the **REGION** scope.

  – **COMPARTMENT**: You can grant remote data access to any resource, compartment, or database in the compartment of the Autonomous Database instance that is registering the data set.

  – **REGION**: You can grant remote data access via cloud links to other tenancies in the Autonomous Database instance that is registering the data set.

  > **Note:**
  >
  > Marketplace listing currently can only be published up to the compartment level. The tool does not allow you to publish a marketplace listing to entire region for security reasons.

  – **None**: You do not have access to any resource.

- **Read**:
  This option allows **READ** access for the specific user to access the registered table.

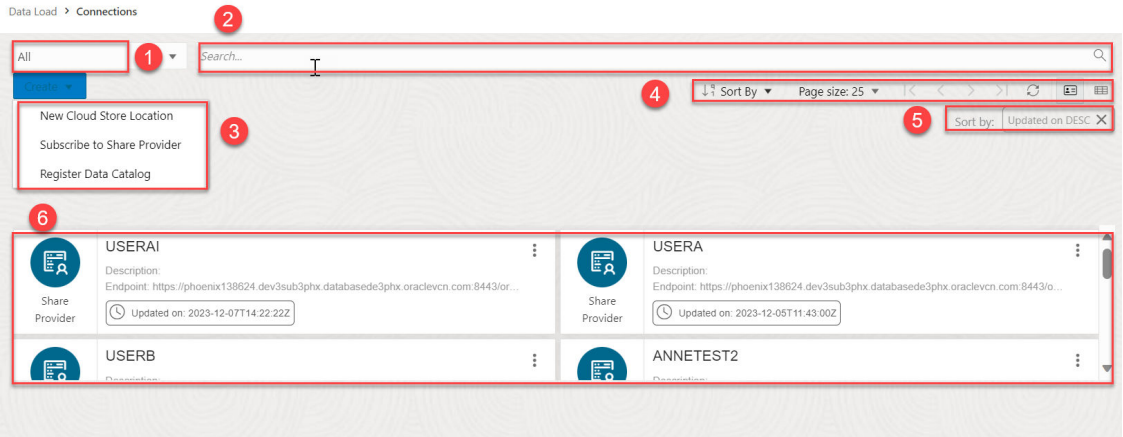
Select **Apply Changes** to save and apply the configuration you made in the Marketplace tab of the Data Settings dialog.

You have successfully set and configured your preferences in the Data Studio Settings wizard.

Managing Connections

Connections that are established from the Data Studio to the cloud, catalogs and shares are listed on this page.

The Connections page contains the Search for entities such as Cloud Storage Links, Data Catalog Links, and Share Providers, and a list of entity cards. You can enter the entity you are looking for in the field or click one of the entity card from the list. You can register the cloud store you want to use from this page. You can also register data catalogs and subscribe to a share provider.



The Connections page consists of the following:

1. **Entity Selector**: You can Use the drop-down lists to select the entity from which the connection is created. You can select from Cloud Storage Links, Data Catalog Links, Share Provider or all.
2. **Search field**: Searches for entities in the field by name. The search functionality is not case-sensitive, retrieves all matching entries, and does not require the use of wild card characters.
3. **Create drop-down**: The options available in the Create drop-down are:
   - **New Cloud Store Location**: Before you can load data from a cloud store, you must establish a connection to the cloud store you want to use. You can select cloud store location from the cloud store locations field. A cloud storage link is a connection to a bucket in a cloud store. You can view the existing cloud storage links and add new ones. Refer to the Create Credentials to add a cloud store location.
   - **Subscribe to a Share Provider**: Upload the JSON profile file and create a share provider description To subscribe, you need to use the information contained in the uploaded JSON profile you received from the share provider. From the Consume Share feature of the Data Share tool, you upload the JSON profile and follow the subscribe wizard. Refer to the Consume Versioned Share for more details.
   - **Register Data Catalogs**: You can create a connection by registering a data source as a data asset in your data catalog. You can view, delete and rename the data catalogs. Refer to the Register Data Catalog chapter.
4. The toolbar consists of the following buttons:
   - **Sort By**

     To select sorting values, click the **Sort By** button to open the list of options. Then click the **Ascending** or **Descending** icon next to one or more of the sorting values. For example, if you select the **Ascending** icon next to **Entity name** and the **Descending** icon next to **Entity type**, the entities will be sorted in alphabetical order by entity name and then in reverse alphabetical order by entity type.

     Click **Reset** in the list to clear the choices in the list.

- **Page size**

  By default, up to 25 entities are displayed on the page. If you want more entities on a page, select a number from this list.

- **Previous** and **Next**

  If the search results are displayed on multiple pages, click these buttons to navigate through the pages.

- **Refresh**

  Click to refresh the data load jobs shown on the page, based on the current search field.

5. The sorting values you choose are listed next to the **Sort by** label beneath the toolbar. Click the **X** icon on a sorting value to remove it.
6. Display area: The area below the Create drop-down field displays the entity carts returned by a search.

View Entity Details

- To view details about the existing entities, click **Actions**.
- For a Share Provider, refer to the View Share Provider Entity detail.
- For a Cloud Storage Link, click **Actions** to perform the following operations:
  - Select **View Details** to view details about the table.
  - Select **Objects** to view objects available in the selected storage link. You can click a file on the navigator pane to view it on the display area.
  - Selecting **Link Tables** opens the Link Data page on the Data Load tool with the selected cloud storage link on the Cloud Location URL field. You can link data present in the cloud storage to the Autonomous Database. See Linking to Objects in Cloud Storage.
  - Selecting **Load Tables** opens the Load Data page on the Data Load tool with the selected cloud storage link on the Cloud Location URL field. You can load data present in the cloud storage to the Autonomous Database. See Loading Data from Cloud Storage.
  - Selecting **Create Live Table Feed** opens the Create Live Table feed wizard with he selected cloud storage link on the Cloud Location URL field. See Feeding Data to view more details.
  - Select **Edit** to update any details on the cloud storage location. See Create Credentials to view details on creation of cloud storage location.
  - Select **Rename** to rename the cloud store location to a different name.
  - Select **Delete** to delete the cloud store location.
- Register Data Catalog
  You can register data catalogs you want to use with registering the data catalog.
- Create Credentials
  The procedure for creating a credential varies depending on the cloud storage provider. If your source files reside in a cloud store provided by one of the following, see the example for that provider.
- Create Oracle Cloud Infrastructure Native Credentials
  To establish cloud storage connection from Data Studio to Oracle Cloud Infrastructure (OCI) Object storage service, you must configure the cloud storage location with your OCI

authentication details. You can create Oracle Cloud Infrastructure (OCI) Native Credentials by using the CREATE_CREDENTIAL procedure of DBMS_CLOUD package.

- Create Amazon Web Services (AWS) Credentials
  To access AWS Glue Catalog you must have AWS credentials that you can create via the Create Credential wizard. This wizard stores cloud service credentials in the Autonomous Database.
- Manage Resource Principal with DBMS_CLOUD
  The Oracle Cloud Infrastructure (OCI) Resource Principal allows applications and services running within the OCI to access an Autonomous Database without the need for traditional database credentials (e.g., username and password).

Create Credentials

The procedure for creating a credential varies depending on the cloud storage provider. If your source files reside in a cloud store provided by one of the following, see the example for that provider.

- **Oracle Cloud Infrastructure (OCI)**, see Create an OCI Cloud Store Location.
- **Amazon S3**, or you are calling an AWS API, see Create an Amazon S3 Cloud Store Location.
- **Microsoft Azure Blob Storage** or you are calling an Azure API, see Create a Microsoft Azure Cloud Store Location.
- **Google Cloud Storage**, see Create a Google Cloud Storage Location.
- **Other (Swift compatible)** cloud storage, see Create an Other (Swift Compatible) Cloud Store Location.
- **OCI cloud storage** by using native OCI credentials, see the subsequent section.

> **Note:**
>
> This method is suggested if you need to use the OCI REST APIs. You need to use the Native OCI credentials if you want to use the OCI REST APIs. To create an OCI cloud storage location using Oracle Cloud Infrastructure Signing Keys, you must first Create Oracle Cloud Infrastructure Native Credentials.

- **OCI Cloud Shell**, see Create Credentials using OCI Cloud Shell.

Create an OCI Cloud Store Location

1. On the **Connections** page, click **Create** and select **Create New Cloud Store Location**. This opens the Add Cloud Store Location wizard.
2. In the Storage Settings tab of the Add Cloud Store Location box, enter a name for the cloud storage link. For example:

   ```
   My_Cloud_Store
   ```

3. (Optional) In the **Description** field, enter a description for the link. For example:

   ```
   My cloud storage link.
   ```

4. Click **Select Credential** and **Create Credential** to create a credential. This opens a Create Credential wizard.

In the Credentials section, select **Cloud Username and Password**.

5. From the **Cloud Store** drop-down list, select **Oracle**.
6. Enter a name in the **Credential Name** field. The name must conform to Oracle object naming conventions. For example:

   ```
   my_credential
   ```

7. For an OCI cloud store, in the **Oracle Cloud Infrastructure User Name** field, enter your OCI user name. You must use Oracle Cloud Infrastructure User Name from your profile in the OCI console. For example:

   ```
   oracleidentitycloudservice/foo@example.com
   ```

   or

   ```
   default/foo@example.com
   ```

8. For an OCI cloud store, in the **Auth Token** field, enter your auth token. For example:

   ```
   LPB>Ktk(1M1SD+a]+r
   ```

9. In the **Bucket URI** field, enter the URI and bucket for your OCI instance bucket.
   a. To get the URI and bucket, go to the bucket in the Object Storage compartment in your Oracle Cloud Instance.
   b. In the Objects group, click the Actions (three vertical dots) icon for a file in the bucket, then click **View Object Details**.
   c. Copy all of the URL Path (URI) except for the file name. Be sure to include the trailing slash. For example, for the file `https://objectstorage.us-phoenix-1.oraclecloud.com/n/myoci/b/my_bucket/o/MyFile.csv`, select the following:

      ```
      https://objectstorage.us-phoenix-1.oraclecloud.com/n/myoci/b/
      my_bucket/o/
      ```

   d. Paste the string into the **URI + Bucket** field.
10. Click **Next**.
    The dialog box progresses to the Cloud Data tab. This tab lists the objects available on this cloud storage location in the display area.

    > **Note:**
    >
    > The display area is blank when we create a new cloud storage location.

11. Click **Create** create the cloud storage location.

Create an Amazon S3 Cloud Store Location

1. On the **Connections** page, click **Create** and select **Create New Cloud Store Location**. This opens the Add Cloud Store Location wizard.

2. In the Storage Settings tab of the Add Cloud Store Location box, enter a name for the cloud storage link. For example:

   ```
   My_Cloud_Store
   ```

3. (Optional) In the **Description** field, enter a description for the link. For example:

   ```
   My cloud storage link.
   ```

4. Click **Select Credential** and **Create Credential** to create a credential. This opens a Create Credential wizard.
   In the Credentials section, select **Cloud Username and Password**.
5. From the **Cloud Store** drop-down list, select **Oracle**.
6. Enter a name in the **Credential Name** field. The name must conform to Oracle object naming conventions. For example:

   ```
   my_credential
   ```

7. From the **Cloud Store** drop-down list, select **Amazon S3**.
8. Enter a name in the **Credential Name** field. The name must conform to Oracle object naming conventions. For example:

   ```
   my_credential
   ```

9. In the **AWS Access Key ID** field, enter your AWS access key ID. For example:myAccessKeyID
10. In the **AWS Secret Access Key** field, enter your AWS secret access key. For information on AWS access keys, see Managing access keys for IAM users.
11. In the **Bucket URI** field, enter the URI and bucket for your Amazon S3 bucket. For example:

    ```
    https://s3.us-west-2.amazonaws.com/my_bucket
    ```

12. Click **Next**.
    The dialog box progresses to the Cloud Data tab. This tab lists the objects available on this cloud storage location in the display area.

    > **Note:**
    >
    > The display area is blank when we create a new cloud storage location.

13. Click **Create** create the cloud storage location.

Create an Microsoft Azure Cloud Store Location

1. On the **Connections** page, click **Create** and select **Create New Cloud Store Location**. This opens the Add Cloud Store Location wizard.
2. In the Storage Settings tab of the Add Cloud Store Location box, enter a name for the cloud storage link. For example:

   ```
   My_Cloud_Store
   ```

3. (Optional) In the **Description** field, enter a description for the link. For example:

   ```
   My cloud storage link.
   ```

4. Click **Select Credential** and **Create Credential** to create a credential. This opens a Create Credential wizard.
   In the Credentials section, select **Cloud Username and Password**.
5. From the Cloud Store drop-down list, select **Microsoft Azure**.
6. Enter a name in the **Credential Name** field. The name must conform to Oracle object naming conventions. For example:

   ```
   my_credential
   ```

7. In the **Azure Storage Account Name** field, enter the name of your Azure storage account. For example:

   ```
   myaccount
   ```

8. In the **Azure Storage Account Access Key** field, enter your Azure access key.
   For information on Azure storage accounts, see Create a storage account.
9. In the **Bucket URI** field, enter the URI and bucket for your Microsoft Azure bucket. For example:

   ```
   https://myaccount.blob.core.windows.net/mycontainer
   ```

10. Click **Next**.
    The dialog box progresses to the Cloud Data tab. This tab lists the objects available on this cloud storage location in the display area.

    > **Note:**
    >
    > The display area is blank when we create a new cloud storage location.

11. Click **Create** to create the cloud storage location.

Create an Microsoft Azure Cloud Store Location

1. On the **Connections** page, click **Create** and select **Create New Cloud Store Location**. This opens the Add Cloud Store Location wizard.
2. In the Storage Settings tab of the Add Cloud Store Location box, enter a name for the cloud storage link. For example:

   ```
   My_Cloud_Store
   ```

3. (Optional) In the **Description** field, enter a description for the link. For example:

   ```
   My cloud storage link.
   ```

4. Click **Select Credential** and **Create Credential** to create a credential. This opens a Create Credential wizard.
   In the Credentials section, select **Cloud Username and Password**.
5. From the Cloud Store drop-down list, select **Microsoft Azure**.

6. Enter a name in the **Credential Name** field. The name must conform to Oracle object naming conventions. For example:

   ```
   my_credential
   ```

7. In the **Azure Storage Account Name** field, enter the name of your Azure storage account. For example:

   ```
   myaccount
   ```

8. In the **Azure Storage Account Access Key** field, enter your Azure access key.
   For information on Azure storage accounts, see Create a storage account.
9. In the **Bucket URI** field, enter the URI and bucket for your Microsoft Azure bucket. For example:

   ```
   https://myaccount.blob.core.windows.net/mycontainer
   ```

10. Click **Next**.
    The dialog box progresses to the Cloud Data tab. This tab lists the objects available on this cloud storage location in the display area.

    > **Note:**
    >
    > The display area is blank when we create a new cloud storage location.

11. Click **Create** to create the cloud storage location.

Create a Google Cloud Store Location

1. On the **Connections** page, click **Create** and select **Create New Cloud Store Location**. This opens the Add Cloud Store Location wizard.
2. In the Storage Settings tab of the Add Cloud Store Location box, enter a name for the cloud storage link. For example:

   ```
   My_Cloud_Store
   ```

3. (Optional) In the **Description** field, enter a description for the link. For example:

   ```
   My cloud storage link.
   ```

4. Click **Select Credential** and **Create Credential** to create a credential. This opens a Create Credential wizard.
   In the Credentials section, select **Cloud Username and Password**.
5. Enter a name in the **Credential Name** field. The name must conform to Oracle object naming conventions. For example:

   ```
   my_credential
   ```

6. From the **Cloud Store** drop-down list, select **Google**.

7. In the **HMAC Access Key** field, enter your HMAC access ID. For example:

   ```
   GOOGTS1C3LPB3KTKSDMB2BFD
   ```

8. In the **HMAC Access Secret** field, enter your HMAC secret. For information on HMAC keys, see HMAC Keys.
9. In the Storage Settings tab of the Add Cloud Store dialog box, enter a name for the cloud storage link. For example:

   ```
   My_Cloud_Store
   ```

10. (Optional) In the **Description** field, enter a description for the link. For example:

    ```
    My cloud storage link.
    ```

11. In the **Bucket URI** field, enter the bucket and URI for your Google bucket. For example:

    ```
    https://my_bucket.storage.googleapis.com
    ```

12. Click **Next**.
    The dialog box progresses to the Cloud Data tab. This tab lists the objects available on this cloud storage location in the display area.

    > **Note:**
    >
    > The display area is blank when we create a new cloud storage location.

13. Click **Create** to create the cloud storage location.

Create an Other (Swift Compatible) Cloud Store Location

1. On the **Connections** page, click **Create** and select **Create New Cloud Store Location**. This opens the Add Cloud Store Location wizard.
2. In the Storage Settings tab of the Add Cloud Store Location box, enter a name for the cloud storage link. For example:

   ```
   My_Cloud_Store
   ```

3. (Optional) In the **Description** field, enter a description for the link. For example:

   ```
   My cloud storage link.
   ```

4. Click **Select Credential** and **Create Credential** to create a credential. This opens a Create Credential wizard.
   In the Credentials section, select **Cloud Username and Password**.

> **Note:**
>
> If you have the user OCID, tenancy OCID, private key and fingerprint, select **Oracle Cloud Infrastructure Signing Keys** and refer to the *Create an OCI Cloud storage location using Oracle Cloud Infrastructure Signing Keys* section of this topic.

5. From the Cloud Store drop-down list, select **Other (Swift Compatible)**.
6. Enter a name in the **Credential Name** field. The name must conform to Oracle object naming conventions. For example:

   ```
   my_credential
   ```

7. In the **Access User Name** field, enter your access user name. For example:

   ```
   OTHER_KEY123...
   ```

8. In the **Access Key** field, enter your access key.
9. In the **Bucket URI** field, enter the URI and bucket for your cloud store bucket. For example:

   ```
   https://someswiftcompatibleprovider.com/my_bucket
   ```

10. Click **Next**.The dialog box progresses to the Cloud Data tab. This tab lists the objects available on this cloud storage location in the display area. The display area is blank when we create a new cloud storage location.
11. Click **Create** to create the cloud storage location.

Create an OCI Cloud storage location using Oracle Cloud Infrastructure Signing Keys

To create an OCI Cloud storage location using Oracle Cloud Infrastructure Signing Keys:

1. On the **Connections** page, click **Create** and select **Create New Cloud Store Location**. This opens the Add Cloud Store Location wizard.
2. In the Storage Settings tab of the Add Cloud Store Location box, enter a name for the cloud storage link. For example:

   ```
   My_Cloud_Store
   ```

3. (Optional) In the **Description** field, enter a description for the link. For example:

   ```
   My cloud storage link.
   ```

4. Click **Select Credential** and **Create Credential** to create a credential. This opens a Create Credential wizard.
   If you have the user OCID, tenancy OCID, private key and fingerprint, select **Oracle Cloud Infrastructure Signing Keys**.

> **Note:**
>
> If you only have a username and password select **Cloud Username and Password** in this step and refer to the *Create an Other (Swift Compatible) Cloud Store Location* section of this topic.

5. Specify the following information about your OCI account:

   Credential Name: Specify a name to identify the credentials. See Create an OCI Credential Object to enter the Credential name to specify the credential name.

6. **Fingerprint:** The fingerprint of the RSA key pair. See Create an OCI Credential Object to enter the Credential name to enter the fingerprint.

7. **Private Key:** The unencrypted private key in the RSA key pair. This should not be encrypted by using any passphrase. See Create an OCI Credential Object to enter the Credential name to enter the private key.

8. **Oracle Cloud Infrastructure Tenancy:** The OCID of the tenant. See Where to Get the Tenancy's OCID and User's OCID for details on obtaining the Tenancy's OCID.

9. **Oracle Cloud Infrastructure User Name:** The OCID of the user. See Where to Get the Tenancy's OCID and User's OCID for details on obtaining the User's OCID.

10. Select **Create Credential**.

11. In the **Bucket URI** field, enter the URI and bucket for your cloud store bucket. For example:

    ```
    https://objectstorage.<region>.oraclecloud.com/n/<namespace>/b/<bucket>/
    ```

12. Click **Next**.

13. The dialog box progresses to the Cloud Data tab. This tab lists the objects available on this cloud storage location in the display area.

> **Note:**
>
> The display area is blank when we create a new cloud storage location.

14. Click **Create**.

You will receive a notification that the cloud storage location is created successfully.

Create Credentials using OCI Cloud Shell

You can create OCI native credentials using a Cloud Shell script.

To create credentials using OCI Cloud shell, you must run `adb-create-cred.sh` to generate credential scripts. The `adb-create-cred.sh` script searches for existing credentials, if found you are asked if you want to reuse them or if you want new credentials created. Depending on your decision, the generated credential scripts include new or existing credentials. Download the OCI Native Credential script or copy it to run it directly in your database using any SQL or JSON tool or utility.

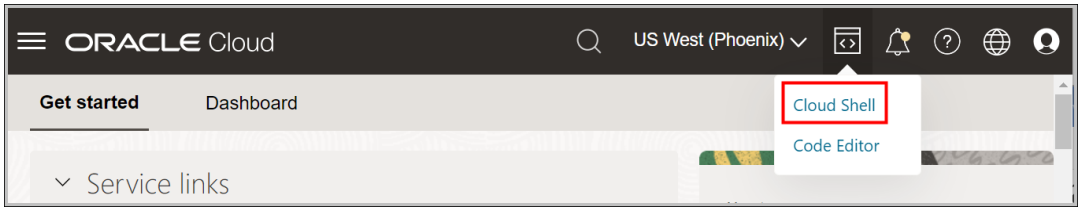
> **Note:**
>
> For a list of arguments supported by the Create Credential script, enter `adb-create-cred.sh --help`.

In this example, OCI Native Credential scripts are generated in the script area for your tenancy without running them in your database.

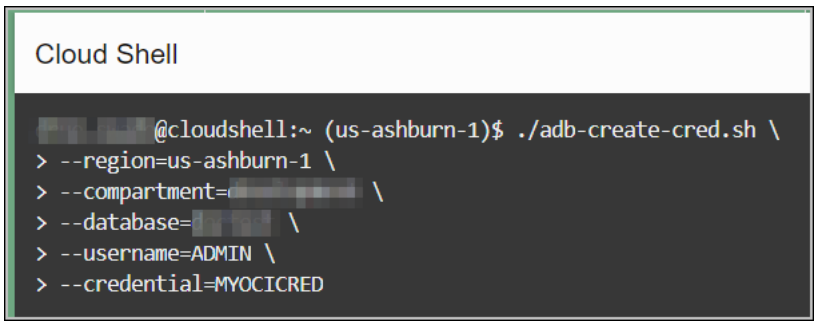
1. On the **Connections** page, click **Create** and select **New Cloud Store Location**. You need to know the Name and Cloud Store Location of the Credential to create a new one.
2. Select **Create Credential**.
3. On the Create Credential wizard, select **Create using OCI Cloud Shell**. You will view a script in the script editor field.

> **Note:**
>
> - If you are in your home region and the Autonomous Database is in your home region, run the script without the region argument. The script assumes the Autonomous Database is in your home region and will not prompt for an alternate region.
> - If you are in your home region and your Autonomous Database is in an alternate region, then you must pass in the --region argument when running the script.

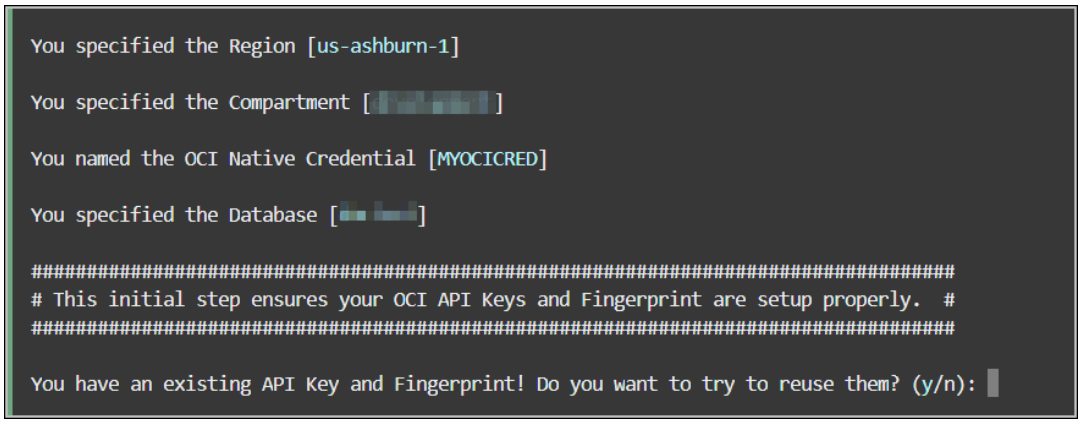
4. Click **Copy Script**.
5. Sign into your OCI Tenancy, select the **Developer Tools** icon and click **Cloud Shell** from the drop down list.



6. Paste the **`adb-create-cred.sh`** script you copied in the previous step that includes arguments the tool uses to locate the database, connect to the database and the credential name you provide in the **Credential Name** field.

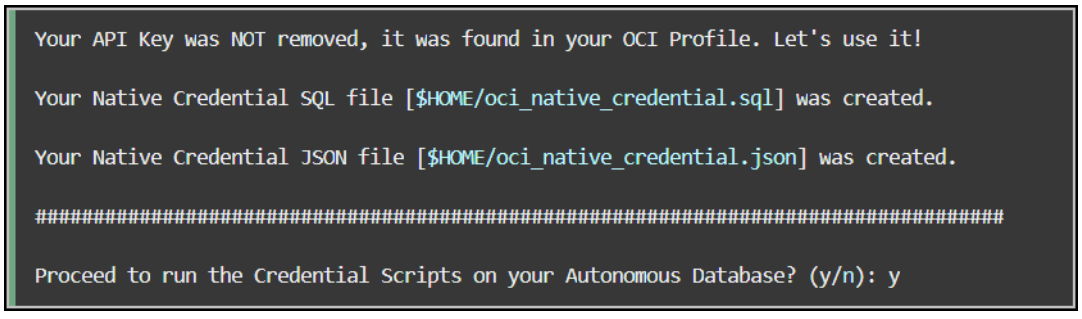


7. The Cloud Shell script informs you of existing OCI API Keys and Fingerprint and if you want to reuse them. Enter `y` to reuse the existing keys and fingerprint.



If you choose to create new credentials, and decide after overwriting the credentials that you want to use the overwritten credentials, you can go back to the directory where the backup is and reuse them. The backup file name or folder has a suffix with this syntax: `_bkp_YYYYMMDD_abc`. For example, this backup file was created on June 06, 2024: `_bkp_20240603_woT`.

8. The Cloud Shell script window displays if you want to run the credential scripts on your Autonomous Database. Enter `y` to run the credential script in the database with the region, compartment and database arguments.



9. You can view a notification that informs you that you do not have a wallet file. The script creates a wallet file. You must specify the ADMIN password you use to connect to the database. In case the script detects the wallet file, the script prompts you to either reuse it or set up a new one.

```
You do not have a Wallet File. Let's set it up for you.

Fixing path in [ _wallet/sqlnet.ora] to that of your Wallet File.

###############################################################################
# We setup your credentials, generated SQL scripts to add the credentials, we  #
# also setup your Autonomous Database connection, and are now connecting you to #
# your Autonomous Database so you can execute the credential scripts.          #
###############################################################################


You specified the Database Username [ADMIN]

Enter Autonomous Database password:
```

> **Note:**
>
> If you have an Autonomous Database private endpoint in a Virtual Cloud Network, the `adb-create-cred.sh` script generates the SQL and JSON scripts. However, it does not perform the steps required to access the Autonomous Database private end point. Instead, it prompts you to execute `cat ~/oci_native_credential.sql` to copy the SQL into whatever SQL tool you have access to via a Bastion or Jump Host.

10. The Cloud Shell script displays a notification of successful login and creation of credentials in the specific database.
    The script prompts you if you want to run it on **another** database. Enter `n` to not run credentials in another database.

    The script exits and displays the path to run the `oci_native_credential.sql` script to create a credential.

```
Login successful.

OCI Native Credential [MYOCICRED] creation successful.

###############################################################################
# Please be aware that it may take up 15 minutes to propagate your credentials to  #
# OCI Services. During this period of time your new credentials will not work.     #
###############################################################################

Proceed to run the Credential Scripts on another Autonomous Database? (y/n): n

Exiting...

To view PL/SQL script for creating credentials,

        execute: cat $HOME/oci_native_credential.sql

 @cloudshell:~ (us-ashburn-1)$
```

> **Note:**
>
> If the connection to the Autonomous Database is unsuccessful, you can run `cat ~/oci_native_credential.sql`, and copy the SQL to run it directly in your database using any SQL tool.

Create Oracle Cloud Infrastructure Native Credentials

To establish cloud storage connection from Data Studio to Oracle Cloud Infrastructure (OCI) Object storage service, you must configure the cloud storage location with your OCI authentication details. You can create Oracle Cloud Infrastructure (OCI) Native Credentials by using the CREATE_CREDENTIAL procedure of DBMS_CLOUD package.

Create an Oracle Cloud Infrastructure (OCI) Credential Object

To access Object Storage, you must have credentials that you can create via the CREATE_CREDENTIAL procedure of DBMS_CLOUD package. DBMS_CLOUD supports creation of credential objects that contains OCI native authentication. The DBMS_CLOUD procedure stores cloud service credentials in Autonomous Database.

The DBMS_CLOUD.CREATE_CREDENTIAL procedure is overloaded with the Oracle Cloud Infrastructure-related parameters, including: user_ocid, tenancy_ocid, private_key, and fingerprint. This is for using Oracle Cloud Infrastructure Signing Keys authentication.

Let us create an OCI native authentication credential when creating an object store credential object. In the OCI native authentication, the DBMS_CLOUD.CREATE_CREDENTIAL procedure includes the following parameters:

Table 3-2 DBMS_CLOUD.CREATE_CREDENTIAL parameter descriptions

Parameter	Description
`credential_name`	The `credential_name` parameter must conform to Oracle object naming conventions.
`user_ocid`	Specifies the user's OCID. See Where to Get the Tenancy's OCID and User's OCID for details on obtaining the User's OCID.
`tenancy_ocid`	Specifies the tenancy's OCID. See Where to Get the Tenancy's OCID and User's OCID for details on obtaining the Tenancy's OCID.
`private_key`	Specifies the generated private key. Private keys generated with a passphrase are not supported. You need to generate the private key without a passphrase. See How to Generate an API Signing Key for details on generating a key pair in PEM format.
`fingerprint`	Specifies a fingerprint. After a generated public key is uploaded to the user's account the fingerprint is displayed in the console. Use the displayed fingerprint for this argument. See How to Get the Key's Fingerprint and How to Generate an API Signing Key for more details.

Here is the syntax of the DBMS_CLOUD.CREATE_CREDENTIAL procedure:

```
DBMS_CLOUD.CREATE_CREDENTIAL (
    credential_name IN VARCHAR2,
    user_ocid       IN VARCHAR2,
    tenancy_ocid    IN VARCHAR2,
    private_key     IN VARCHAR2,
    fingerprint     IN VARCHAR2);
```

Once you obtain all the necessary inputs and generate your private key, here is a sample of your CREATE_CREDENTIAL procedure:

```
BEGIN  DBMS_CLOUD.CREATE_CREDENTIAL (    credential_name =>
      'OCI_NATIVE_CRED',    user_ocid              =>

'ocid1.user.oc1..aaaaaaaatfn77fe3fxux3o5lego7glqjejrzjsqsrs64f4jsjrhbsk5qzndq'
,    tenancy_ocid           =>

'ocid1.tenancy.oc1..aaaaaaaapwkfqz3upqklvmelbm3j77nn3y7uqmlsod75rea5zmtmbl574v
e6a',    private_key            =>

'MIIEogIBAAKCAQEAsbNPOYEkxM5h0DF+qXmie6ddo95BhlSMSIxRRSO1JEMPeSta0C7WEg7g8SOSz
hIroCkgOqDzkcyXnk4BlOdn5Wm/BYpdAtTXk0sln2DH/GCH7l9P8xC9cvFtacXkQPMAXIBDv/
zwG1kZQ7Hvl7Vet2UwwuhCsesFgZZrAHkv4cqqE3uF5p/
qHfzZHoevdq4EAV6dZK4Iv9upACgQH5zf9IvGt2PgQnuEFrOm0ctzW0v9JVRjKnaAYgAbqa23j8tKa
pgPuREkfSZv2UMgF7Z7ojYMJEuzGseNULsXn6N8qcvr4fhuKtOD4t6vbIonMPIm7Z/
a6tPaISUFv5ASYzYEUwIDAQABAoIBACaHnIv5ZoGNxkOgF7ijeQmatoELdeWse2ZXll+JaINeTwKU1
fIB1cTAmSFv9yrbYb4ubKCJuYZJeC6I92rT6gEiNpr670Pn5n43cwblszcTryWOYQVxAcLkejbPA7j
Zd6CW5xm/vEgRv5qgADVCzDCzrij0t1Fghicc+EJ4BFvOetnzEuSidnFoO7K3tHGbPgA+DPN5qrO/
8NmrBebqezGkOuOVkOA64mp467DQUhpAvsy23RjBQ9iTuRktDB4g9cOdOVFouTZTnevN6JmDxufu9L
ov2yvVMkUC2YKd+RrTAE8cvRrn1A7XKkH+323hNC59726jT57JvZ+ricRixSECgYEA508e/
alxHUIAU9J/uq98nJY/
6+GpI9OCZDkEdBexNpKeDq2dfAo9pEjFKYjH8ERj9quA7vhHEwFL33wk2D24XdZl6vq0tZADNSzOtT
rtSqHykvzcnc7nXv2fBWAPIN59s9/oEKIOdkMis9fps1mFPFiN8ro4ydUWuR7B2nM2FWkCgYEAxKs/
zOIbrzVLhEVgSH2NJVjQs24S8W+99uLQK2Y06R59L0Sa90QHNCDjB1MaKLanAahP30l0am0SB450kE
iUD6BtuNHH8EIxGL4vX/SYeE/
AF6tw3DqcOYbLPpN4CxIITF0PLCRoHKxARMZLCJBTMGpxdmTNGyQAPWXNSrYEKFsCgYBp0sHr7TxJ1
WtO7gvvvd91yCugYBJAyMBr18YY0soJnJRhRL67A/
hlk8FYGjLW0oMlVBtduQrTQBGVQjedEsepbrAcC+zm7+b3yfMb6MStE2BmLPdF32XtCH1bOTJSqFe8
FmEWUv3ozxguTUam/
fq9vAndFaNre2i08sRfi7wfmQKBgBrzcNHN5odTIV8l9rTYZ8BHdIoyOmxVqM2tdWONJREROYyBtU7
PRsFxBEubqskLhsVmYFO0CD0RZ1gbwIOJPqkJjh+2t9SH7Zx7a5iV7QZJS5WeFLMUEv+YbYAjnXK+d
OnPQtkhOblQwCEY3Hsblj7Xz7o=',    fingerprint            =>
      '4f:0c:d6:b7:f2:43:3c:08:df:62:e3:b2:27:2e:3c:7a');END;/
    PL/SQL procedure successfully completed.
```

You can now retrieve the new credentials with the following query:

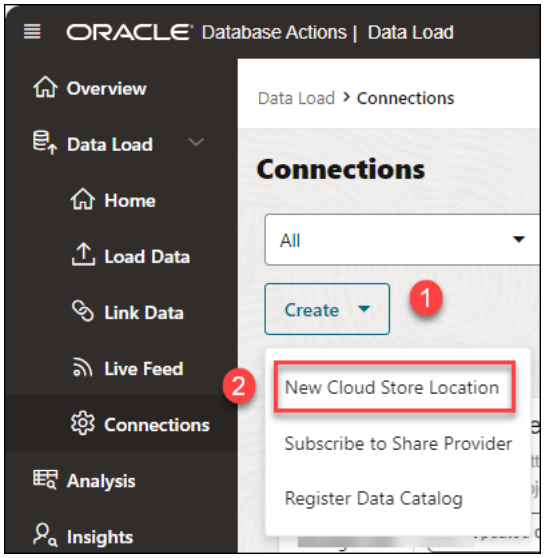
```
SELECT owner, credential_name FROM dba_credentials WHERE credential_name LIKE
'%NATIVE%';  OWNER CREDENTIAL_NAME----- ---------------
    ADMIN OCI_NATIVE_CRED
```

Create Amazon Web Services (AWS) Credentials

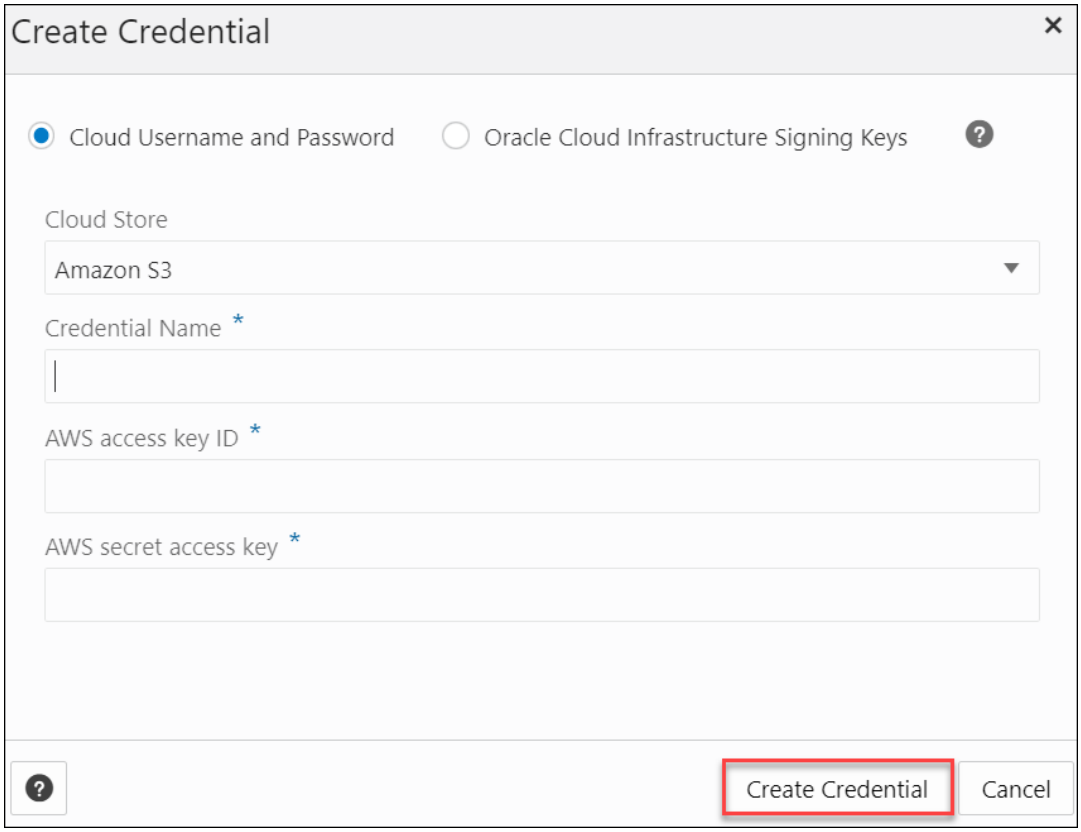
To access AWS Glue Catalog you must have AWS credentials that you can create via the Create Credential wizard. This wizard stores cloud service credentials in the Autonomous Database.

Use the `CREATE CREDENTIAL` wizard depending on where your source files reside. To create AWS credentials:

1. On the **Connections** page, click **Create** and select **New Cloud Store Location**.



2. Select **Create Credential** from **Add Cloud Store Location** wizard.



3. Select **Cloud Username and Password**.
4. Select **Amazon S3** from the Cloud Store drop-down.
5. Enter a name in the Credential Name field. The name must conform to Oracle object naming conventions. For example, `MY_AWS_CRED`.
6. In the **AWS Access Key ID** field, enter your AWS access key ID. For example: `myAccessKeyID`, see Managing access keys for IAM users.
7. In the **AWS Secret Access Key** field, enter your AWS secret access key. For information on AWS access keys, see Managing access keys for IAM users.
8. Click **Create Credential** to create AWS credentials.

Manage Resource Principal with DBMS_CLOUD

The Oracle Cloud Infrastructure (OCI) Resource Principal allows applications and services running within the OCI to access an Autonomous Database without the need for traditional database credentials (e.g., username and password).

You do not need to create a credential object when you use a resource principal. The Autonomous Database creates and secures the resource principal credentials you use to access the specified Oracle Cloud Infrastructure resources. A resource principal consists of a temporary session token and secure credentials that enable the database to authenticate itself to other Oracle Cloud Infrastructure services.

Prerequisites to Use Resource Principal with Autonomous Database:

- There are several steps required to set up a resource principal on an Autonomous Database: You must create and define Oracle Cloud Infrastructure Identity and Access

Management (IAM) policies. See Perform Prerequisites to Use Resource Principal with Autonomous Database for more information.

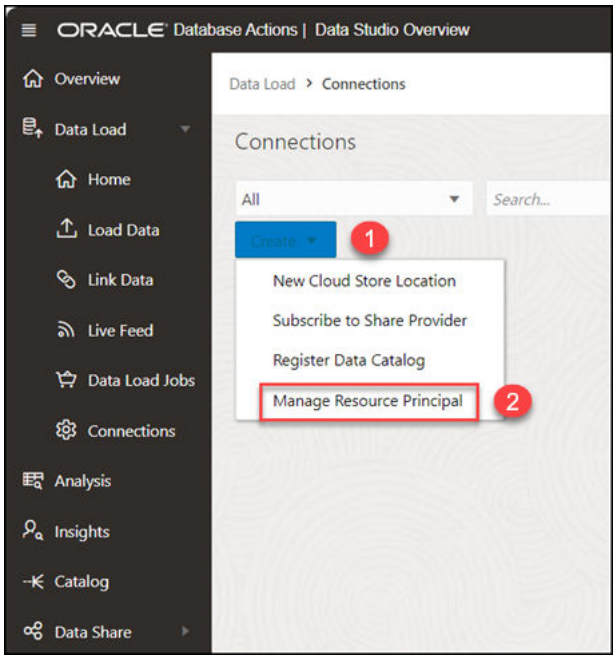
- Enable the resource principal for the ADMIN user, and optionally enable the resource principal for a database user. See Enable Resource Principal to access OCI Resources.

When you authenticate using a resource principal, you do not need to create and manage credentials to access OCI resources. The Autonomous Database makes the resource principal available to you and secures the resource principal for you.

Manage Resource Principal using Data Studio

You can also enable Resource Principal using Data Studio.

1. From the Database Actions Launchpad, select the **Data Load** tile.
2. Click **Connections** This opens the Connections page.
3. On the **Connections** page, click **Create** and select **Manage Resource Principal.**



This opens the Manage Resource Principal dialog.



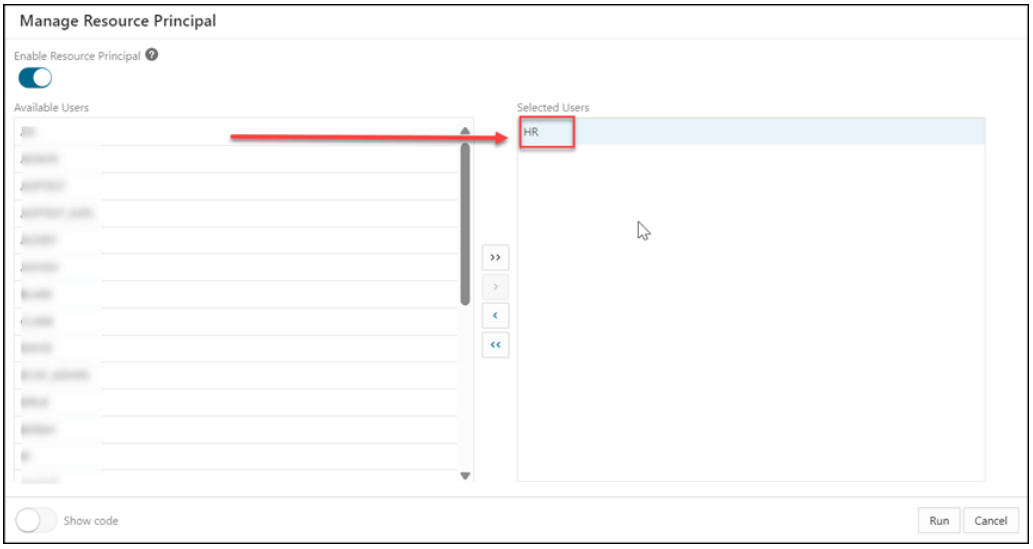
4. Select the user or users you want to add from the **Available Users** list to the **Selected Users** list to enable their resource principal.
Choose any of the available options:

   - **>**: This option enables you to move the user to **Selected Users**.

     <: To remove the selected user from **Selected Users**, select this option.

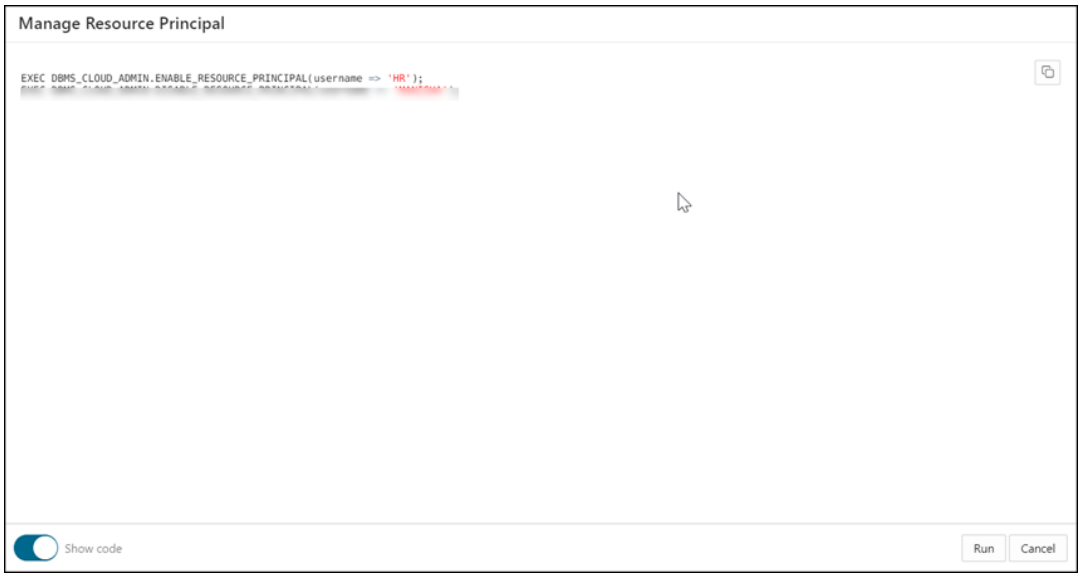
     >>: This option allows you to move all the tables to the **Selected Users** list.

     <<: To remove all the selected users from **Selected Users**, select this option.



> **Note:**
>
> You can grant access to the resource principal credential to a database user only if the ADMIN user has enabled the resource principal credential.

5. Select the **Show Code** option to view the PL/SQL code equivalent of the Manage Resource Principal dialog box. You can copy and execute this PL/SQL code in the worksheet to perform the same action that occurs when you click **Enable Resource Principal** in the Manage Resource Principal dialog box.



6. Click **Run** to complete the process of enabling the resource principal of the selected users. A confirmation notification is displayed that confirms that the resource principal of selected users is enabled.
7. Select a user from Selected Users to Available Users to **disable** the resource principal for the selected user. This removes the credential `OCI$RESOURCE_PRINCIPAL`. Refer to Disable Resource Principal on Autonomous Database for more information on this topic.

Register Data Catalog

You can register data catalogs you want to use with registering the data catalog.

Register OCI Data Catalog

To register a data catalog you need to specify the details of the credentials you want to register your data source. A credential object manages a data catalog instance. In OCI native authentication, the DBMS_CLOUD.CREATE_CREDENTIAL procedure includes these parameters: `credential_name`, `user_ocid`, `tenancy_ocid`, `private_key`, and `fingerprint`.

See REF DBMS_DCAT Package to refer to the procedures to add custom parameters such as region and data catalog ID to the Data Catalog. The Data Catalog ID is a unique Oracle cloud Identifier for data catalog instance and region is the data catalog region.

To register Data catalog:

1. From Connections page, click **Create** and select **Register Data Catalog**. This opens Register Data Catalog wizard.
2. In the Catalog Settings tab, specify the following details:
   - **Catalog Name:** `Test`. Enter a name of your choice.
   - **Description:** Specify a description. This is an optional field.
   - Under Data Catalog details, fill in the following field values:

  - **Credential for Data Catalog Connection:** Select a value from the drop-down. The drop-down lists the values of the credentials you create. If you do not have credentials, you can create one. Refer to Create Credentials for more information.
  - **Region:** `us-ashburn-1`. Enter the region name you use while you create the credentials using the DBMS_DCAT package.
  - **Data Catalog ID:** `ocid1.datacatalog.oc1.iad.amaaaaaa7ratcziayxh7uzll24cp3uwzsugfj7qlubak77toiehidpsqsygq`. Enter the data catalog ID from the DBMS_DCAT package.
  - Select the **Register Data Catalog Connection** button to register the data catalog within the autonomous database. You can view this option when you select **OCI Data Catalog** from Catalog Type.
  - Select the **Use separate credential** for object Storage to select the database credentials from the drop-down. You can view this option when you select **OCI Data Catalog** from Catalog Type.

3. Click **Next** to progress to the Register Assets tab. This tab creates a connection with the source Data Catalog objects you select from the list of objects.
4. After successful registering of the data catalog objects, click **Create**.
5. You will receive a notification that the data catalog is created successfully.
   After successful creation of the data catalog, you can view the data catalog entity in the list of entities in the **Connections** page.

Register AWS Glue Catalog

You can integrate Oracle Data Catalog with Amazon's Glue data catalog.

To access the AWS Glue Catalog, register the catalog on the **Connections** page.

This enables you to synchronize Data Catalog metadata with AWS Glue and query data stored in S3 from an Autonomous Database without manually deriving the schema for the external data sources and creating external tables.

Refer to Query External Data with Glue Data CatalogQuery External Data with Glue Data Catalog for more details.

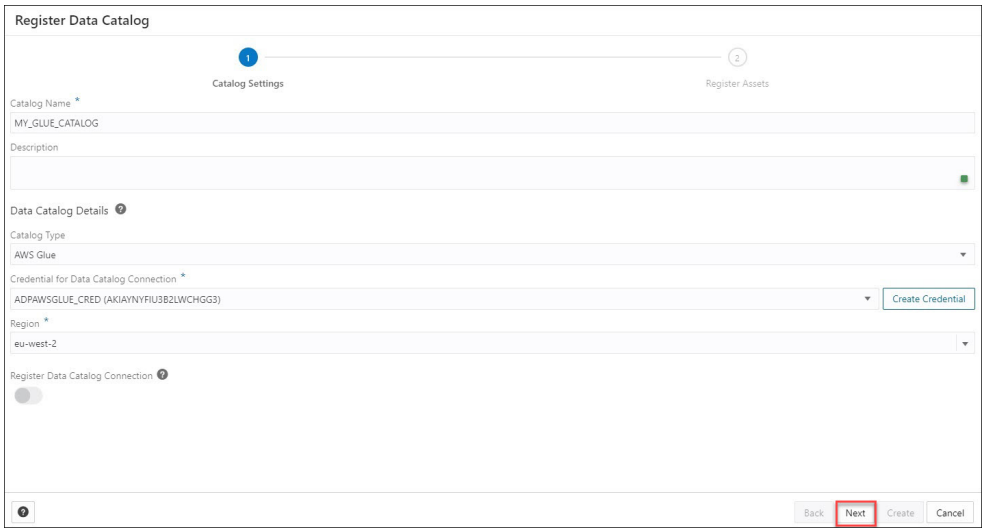
You also need an AWS Credential to be associated with the AWS Glue Catalog. The tool initiates the connection by specifying the following field values:

- Cloud Store
- Credential Name
- AWS access key ID
- AWS Secret Access Key

To register the Data Catalog:

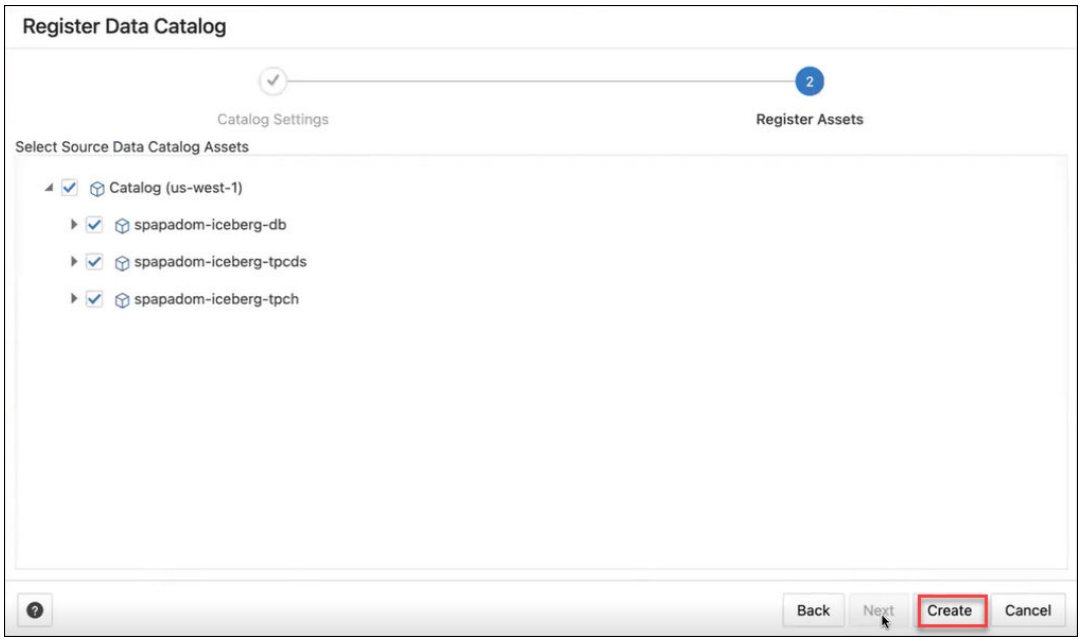
1. From the Connections page, click **Create** and select **Register Data Catalog**. This opens the Register Data Catalog wizard.
2. In the Catalog Settings tab, specify the following details:
   - **Catalog Name**: `MY_GLUE_CATALOG`
   - **Description**: Specify a description. This is an optional field.
   - Under Data Catalog Details, fill in the following field values:
     - **Catalog Type**: Select AWS Catalog from the drop-down.

- **Credential for Data Catalog Connection**: Select a credential value from the drop down. The drop-down lists the values of the credentials you create. If you do not have credentials, you can create one. Refer to Create Amazon Web Services (AWS) Credentials for more information.
- **Region**: `eu-west-2`. Enter the region name you use while you create the AWS credentials.



The **Register Data Catalog Connection** field will be greyed out since you selected AWS Catalog. You can select this field only if you want to register data catalog connection within the Autonomous Database (i.e., when you select OCI Data Catalog).

3. Click **Next** to progress to the Register Assets tab. The Register Assets tab creates a connection with the source data catalog objects you select from the list of objects.
4. After successfully registering the data catalog objects, click **Create**.



5. You will receive a notification that the data catalog has been created successfully. After the successful creation of the data catalog, you can view the data catalog entity in the list of entities on the **Connections** page.

Loading Data

You can load data from files on your local device, from remote databases, or from cloud storage buckets from directories and share providers. The file formats that you can load are CSV, XLS, XLSX, TSV, TXT, XML, JSON, GEOJSON, AVRO, PARQUET, GZ, GZIP, ZIP, PDF, PNG, JPG, JPEG and TIFF.

The Data Load tool supports the loading of source files with boolean format. The data type boolean has the truth values `TRUE` and `FALSE`. If there is no `NOT NULL` constraint, the boolean data type also supports the truth value `UNKNOWN` as the null value. You can use the boolean data type wherever the datatype appears in Oracle SQL syntax.

It also supports the `VECTOR` data type, which allows you to store vector embeddings directly within Oracle Database tables.

- Use OCI Language Service Capabilities in Data Studio
  You can utilize OCI Language Service Capabilities such as Sentiment Analysis, Key Phrase Extraction and Language Detection to analyze data without machine learning (ML) or artificial intelligence (AI) expertise.
- Use OCI Document Understanding to extract Tables from Images
  You can utilize Oracle Cloud Infrastructure(OCI) Document Understanding Capabilities such as Table Extraction to detect tables from the data you load.
- Use GeoJSON in Data Load
  A GeoJSON object accommodates information about the specific geometry (e.g. Point, Line, Polygon, etc.) along with optional metadata (e.g. ID, etc.).
- Loading Data From Local Files
  To load data from local files into your Oracle Autonomous Database, on the Data Load page, select **LOAD DATA** and **LOCAL FILE**.
- Loading Data from Other Databases
  To load data from tables in another database into your Oracle Autonomous Database, on the Data Load page, select **LOAD DATA** and **DATABASE**. Select a database link from the drop-down list. Drag one or more tables from the list of database tables and drop them in the Data Load Cart.
- Loading Data from Cloud Storage
  You can load data from a cloud store to a table in your Autonomous Database.
- Load Apache Iceberg Tables
  Data Studio supports loading and linking Iceberg tables from the object store.
- Loading Data from File System
  You can load files from file system directories to your Autonomous Database.
- Loading Data from AI Source
  You can use Data Studio tools to load data from AI source.
- Loading Data from Share
  You can select tables from a Share. You need to subscribe and access provided data share.
- Create Live Feed from Data Load
  The Data load tool loads the data from folders in cloud object stores and enables it to schedule repeated data loads in real time. This is the creation of Live Feed from a data load job.

Use OCI Language Service Capabilities in Data Studio

You can utilize OCI Language Service Capabilities such as Sentiment Analysis, Key Phrase Extraction and Language Detection to analyze data without machine learning (ML) or artificial intelligence (AI) expertise.

For example, you can use it for feedback on a product. A phone manufacturer has launched a new phone model and they want to know what the customer sentiment is on their product. If a large percentage of sentiment is negative it could signal a potential fault with the camera which wasn't detected in Quality Control (QC).

- Overview of Sentiment Analysis, Key Phrase Extraction and Language Detection
  Sentiment Analysis, Key Phrase Extraction, Language Detection and Text Translation are currently supported in loading data from local files and loading data from cloud storage.
- Parameters to Analyze Data
  When you invoke an **Add Expression** from the Settings tab, you must configure the model using parameters.
- Perform Sentiment Analysis
  To determine the Sentiments of input data:
- Perform Key Phrase Extraction
  To extract Key Phrase information from input data:
- Perform Language Detection
  The Data Studio supports the detection of multiple languages.
- Perform Text Translation
  The Data Studio supports the translation of multiple languages supported by OCI Language Service.

Overview of Sentiment Analysis, Key Phrase Extraction and Language Detection

Sentiment Analysis, Key Phrase Extraction, Language Detection and Text Translation are currently supported in loading data from local files and loading data from cloud storage.

Sentiment Analysis

Sentiment Analysis analyses the text to define your sentiment on a topic or product. The Language service sentiment analysis uses natural language processing (NLP). The Data Studio tool uses the Oracle Cloud Infrastructure (OCI) Language service to analyze and understand the input data. The Data Studio tool dynamically adds new columns to the data load that contains the output of the OCI Language service. You can detect the sentiments of any column of the source data. For example, when searching through a column containing reviews for an application, assume that you want a general opinion about the application. The Data Studio tool performs sentiment analysis on the input data and creates a new expression column defined in the target table that consists of the sentiment.

For more details, refer to Sentiment Analysis in OCI.

Key Phrase Extraction

Key phrase extraction identifies the main concepts in a text. Keyword extraction is the automated process of extracting the words with the most relevance, and expressions from the input text. It helps summarize the content and recognizes the main topics. A word cloud can be generated with key phrases to help visualize key concepts in text comments or feedback. For example, a Movie review could generate a word cloud based on key phrases identified in their comments and might see that people are commenting most frequently about the direction, acting, and cinematography staff.

For more details, refer to Key Phrase Extraction in OCI.

Language Detection

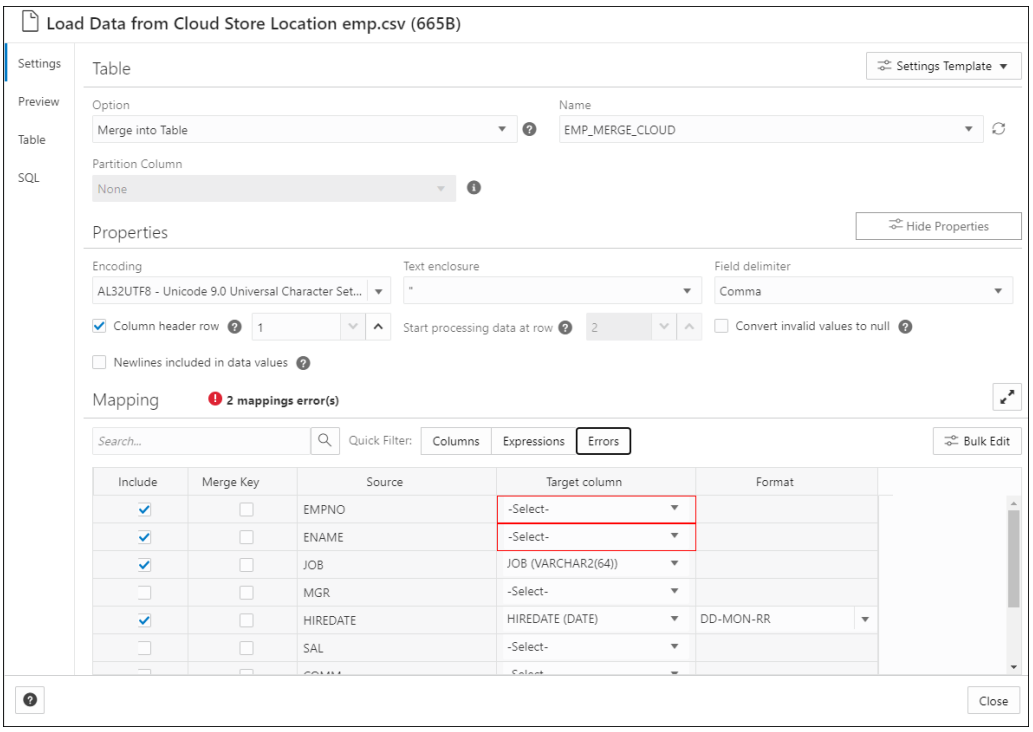
You can utilize OCI Language Service Capabilities such as Language Detection to detect the language of the input text. It returns which natural language the text is in. You can use it to overcome language barriers thus improving communication with people from other countries. You can determine the language of the input text and translate content to different languages. For more details refer to Language Extraction in OCI.

Text Translation

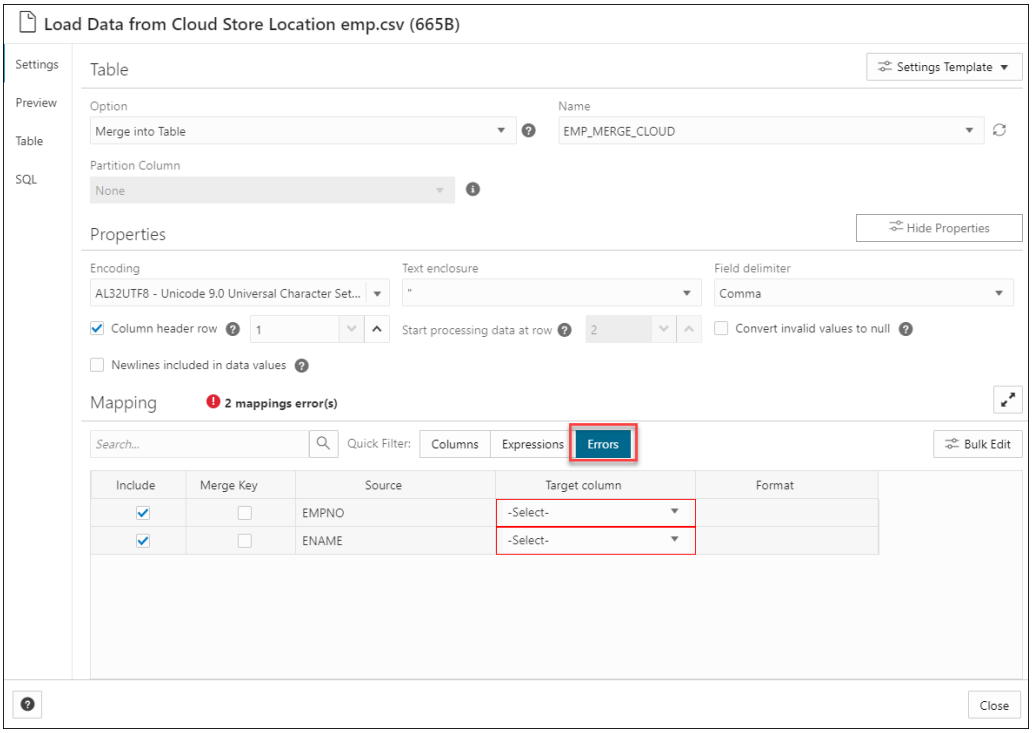
The Text Translation translates input text into any one of the OCI Language service-supported target languages. The Data Studio tool uses the Oracle Cloud Infrastructure (OCI) Language service to analyze and understand the input data. The Data Studio tool dynamically adds new column to the data load that contains the translated input source column. For example, when you need to know what customers are saying about your product in the local market language – French. The Data Studio tool performs language translation on the input data and creates a new expression column defined in the target table that consists of the translated source language in the column.

Before you start:

- **Load Data from Local Files or Cloud storage:** Load the data you wish to analyze into your Oracle Autonomous Database from Local files or Cloud storage. Make sure the data is loaded to the data load cart without any errors. You can view and fix mapping errors from the **Errors** quick filter in the **Load Data from Cloud Store Location**. After you load data into the Data load tool,
  - Click the **Settings** icon on the Data Load job cart to review the settings.
  - On the Settings pane of the **Load Data from Cloud Store Location** wizard, if there is a mapping error, the mapping grid cell will be highlighted with red to indicate an invalid value that must be fixed.



– Click the **Errors** quick filter, it shows only the columns with errors.



- The tools perform sentiment analysis, Key Phrase extraction and Language Detection while you specify settings for the data load job.

Parameters to Analyze Data

When you invoke an **Add Expression** from the Settings tab, you must configure the model using parameters.

Table 3-3 Parameters for Sentiment Analysis, Key Phase Extraction, and Language Detection

Parameter	Description
Expression Type	Select any one of the following operations you wish to perform on the input text: Sentiment Analysis, Key Phrase Extraction, and Language Detection.
Input Column	Select the column which you wish to analyze. The input column drop-down only contains columns that Sentiment Analysis, Key Phrase Extraction, and Language Detection support. For Sentiment Analysis, only VARCHAR2, NVARCHAR2, CLOB, or NCLOB target columns will be displayed in the input drop-down.

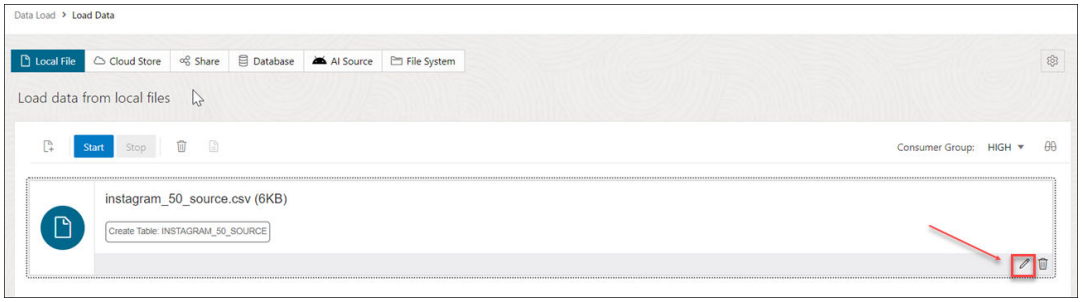
Table 3-3 (Cont.) Parameters for Sentiment Analysis, Key Phase Extraction, and Language Detection

Parameter	Description
Target Column	• Enter the name of the newly created expression column defined in the target table. • For Sentiment Analysis, this column displays the sentiment of the input column. The different types of sentiments the tool identifies are: – Positive – Neutral – Mixed – Negative If the tool cannot determine the sentiment of the input column, it returns `NULL` in the expression column. • For Key Phrase Extraction, this column displays the key phrases of the input column you select. • For Language Detection, this column displays the language of the input column you select.

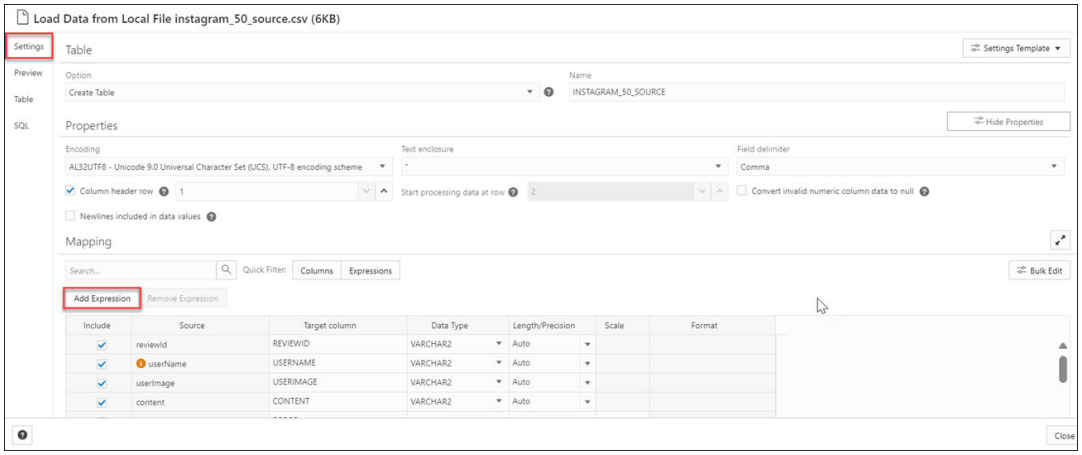
Perform Sentiment Analysis

To determine the Sentiments of input data:

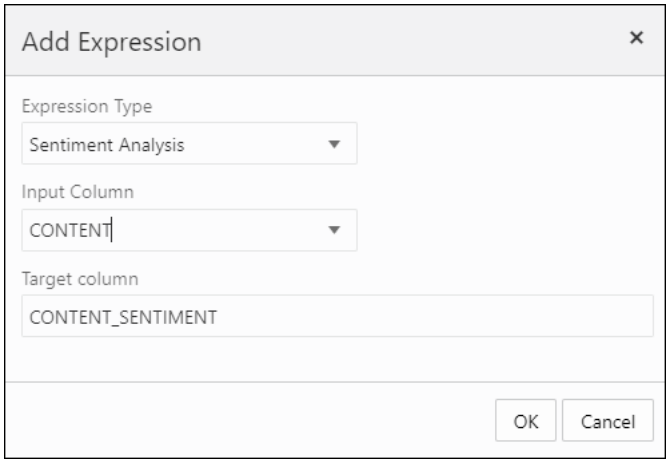
1. After you load data into the Data Load cart from local files or cloud storage, you can view the file in the cart. Click the **settings** icon.



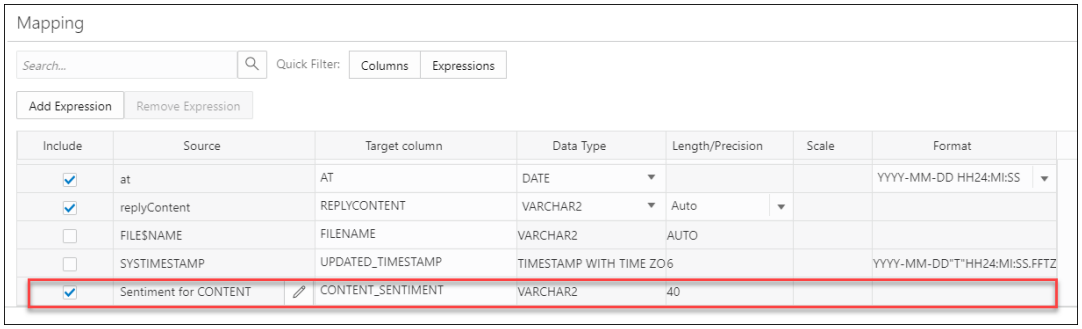
2. Clicking the Setting icon opens a **Load Data from Local File** wizard. In this example, we have loaded data from a local file.
3. On the Settings tab of the wizard, click **Add Expression** under the Mapping section.



This opens the **Add Expression** dialog box.



4. On the **Add Expression** dialog, specify the following fields:
   - **Expression Type:** From the Expression Type drop-down, select **Sentiment Analysis**.
   - **Input Column:** Select the column from the drop-down that you wish to analyze. For example, *CONTENT*.
   - **Target column:** Enter the name of the newly created expression column. For example, *CONTENT_SENTIMENT.*
     Refer to the Parameters to Analyze Data for more details.
5. Click **OK**. You will see a new row added to the mapping grid. This row determines the output expression column generated by the OCI Language service.



6. Click **Close**.
7. Click **Start** in the Data Load menu cart to load data from local files. You will view a confirmation message asking if you wish to start the load from local files.
8. Click **Run** to confirm.

When the data load job completes, the **Table and View Loads** of the Data Load page display the details of the source table that is loaded into the tool. It displays the number of rows and columns and when the upload was complete.

- Output Data Generated from OCI Sentiment Analysis
  When you analyze columns using OCI Language Service model, the Data Studio generates a new expression column and saves the result in the updated table.

Output Data Generated from OCI Sentiment Analysis

When you analyze columns using OCI Language Service model, the Data Studio generates a new expression column and saves the result in the updated table.

To locate the generated expression column, from the Database Actions Launchpad, navigate to **Data Load**. Select the table you load under the **Table and View Loads** section.

Click the three vertical dots beside the load name, and click **Table** then select **View Details**.

This opens the Preview tab of the data load which displays the updated source file. For example, here is an output dataset from sentiment analysis of the Instagram application. Here, `CONTENT` is the target column, and `CONTENT_SENTIMENT` is the sentiment analysis of the input column. This column displays one of the following values such as positive, neutral, mixed, or negative It displays **Null** when the tool is unable to determine the sentiment.

INSTAGRAM_50_SOURCE

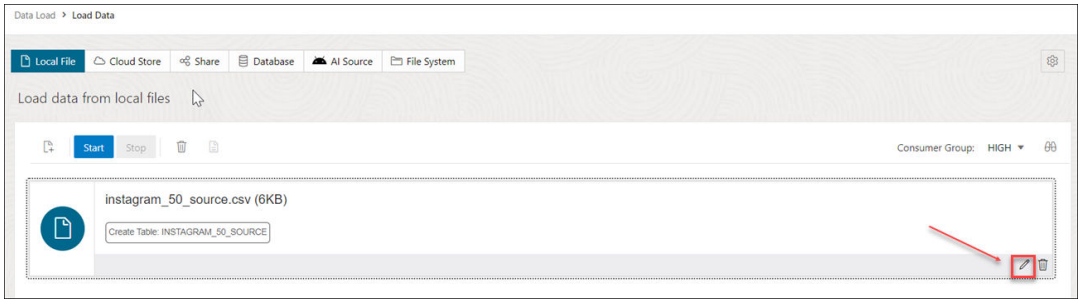
Preview | Lineage | Impact | Statistics | Job Report | Data Definition

		CONTENT	CONTENT_SENTIMENT
1	[redacted]	Nice app pppp	Positive
2	[redacted]	This is very good app	Positive
3	[redacted]	Osm	(null)
4	[redacted]	Mast app aaye 👍	(null)
5	[redacted]	Good	Positive
6	[redacted]	Osm	(null)
7	[redacted]	Good as a social media	Positive
8	[redacted]	Good	Positive
9	[redacted]	Best	Positive
10	[redacted]	💯 💯 💯 👍 👌 👌 👌	(null)
11	[redacted]	Amazing	Positive
12	[redacted]	I really enjoy using this app	Positive

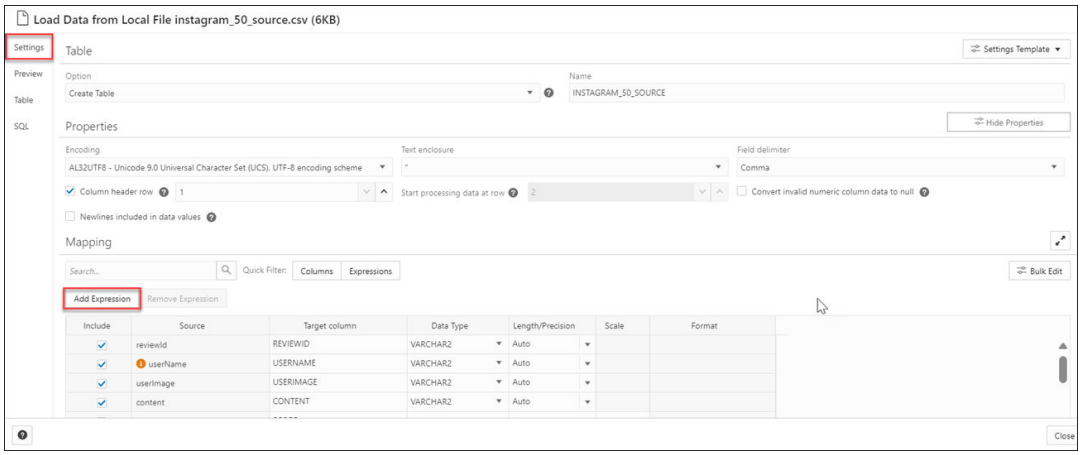
Perform Key Phrase Extraction

To extract Key Phrase information from input data:

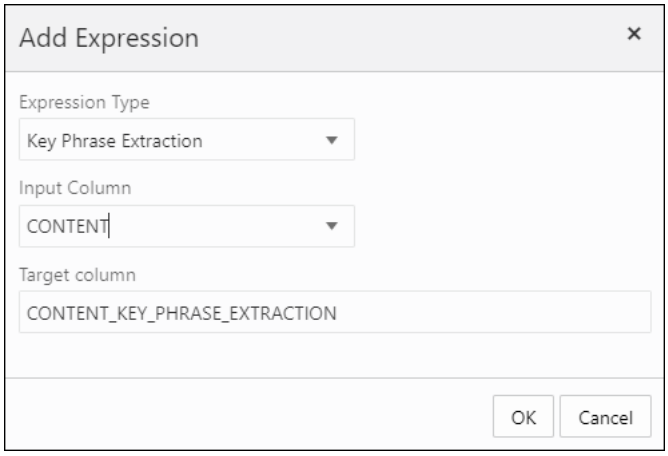
1. After you load data into the Data Load cart from local files or cloud storage, you can view the file in the cart. Click the **settings** icon.



2. Clicking the Setting icon opens a **Load Data from Local File** wizard. In this example, we have loaded data from a local file.
3. On the Settings tab of the wizard, click **Add Expression** under the Mapping section.



This opens the Add Expression dialog box.

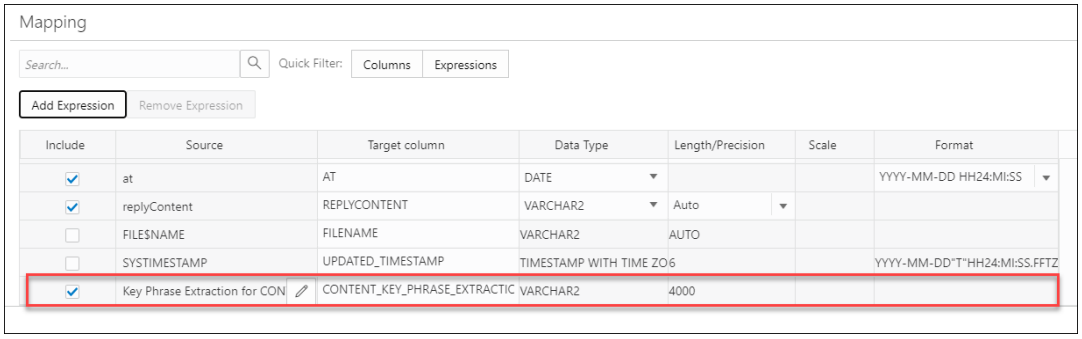


4. On the Add Expression dialog, specify the following fields:
   - **Expression Type:** From the Expression Type drop-down, select Key Phrase Extraction.
   - **Input Column:** Select the column from the drop-down that you wish to analyze. For example, *CONTENT*.

- **Target column:** Enter the name of the newly created expression column. For example, *CONTENT_KEY_PHRASE_EXTRACTION.*

Refer to the Parameters to Analyze Data for more details.

5. Click **OK**. You will see a new row added to the mapping grid .This row determines the output expression column generated by the OCI Language service.



6. Click **Close**.
7. Click **Start** in the Data Load menu cart to load data from local files. You will view a confirmation message asking if you wish to start the load from local files.
8. Click **Run** to confirm.

When the data load job completes, the **Table and View Loads** of the Data Load page display the details of the source table that is loaded into the tool. It displays the number of rows and columns and when the upload was complete.

- Output Data Generated from OCI Key Phrase Extraction
  When you analyze columns using OCI Language Service model, the Data Studio generates a new expression column and saves the result in the updated table.

Output Data Generated from OCI Key Phrase Extraction

When you analyze columns using OCI Language Service model, the Data Studio generates a new expression column and saves the result in the updated table.

To locate the generated expression column, from the Database Actions Launchpad, navigate to **Data Load**. Select the table you load under the **Table and View Loads** section.

Click the three vertical dots beside the load name, and click **Table** then select **View Details**.

For example, here's an output dataset from sentiment analysis of the Instagram application. Here, `CONTENT` is the target column, and `CONTENT_KEY_PHRASE_EXTRACTION` column displays the key phrases extracted from the input column.

INSTAGRAM_50_SOURCE

Preview
Lineage
Impact
Statistics
Job Report
Data Definition

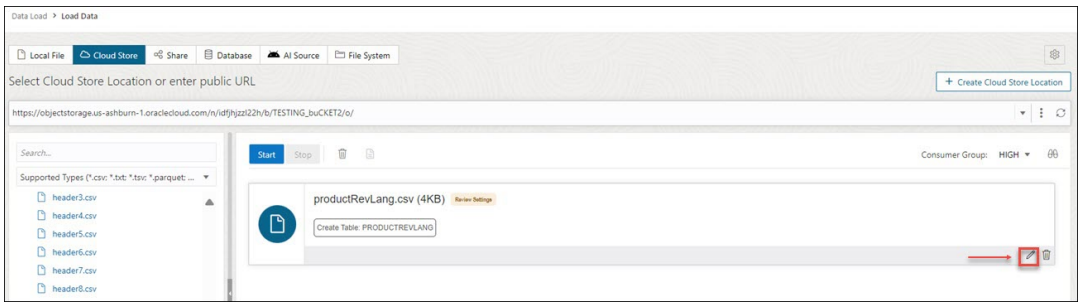
	CONTENT	TENT_KEY_PHRASE_EXTRACTION
1	Nice app pppp	"nice app pppp"
2	This is very good app	"good app"
3	Osm	"osm"
4	Mast app aaye 👍	"mast"
5	Good	(null)
6	Osm	"osm"
7	Good as a social media	"social media"
8	Good	(null)
9	Best	(null)
10	💯 💯 💯 👍 👌 👌 👌	(null)
11	Amazing	(null)
12	I really enjoy using this app	"app"
13	Good	(null)
14	Always my first choice. 👍 👍	"first choice. 👍 "
15	Good	(null)
16	I love talking to my friends	"friends"
17	Best social media app in the	"best social media app"

Perform Language Detection

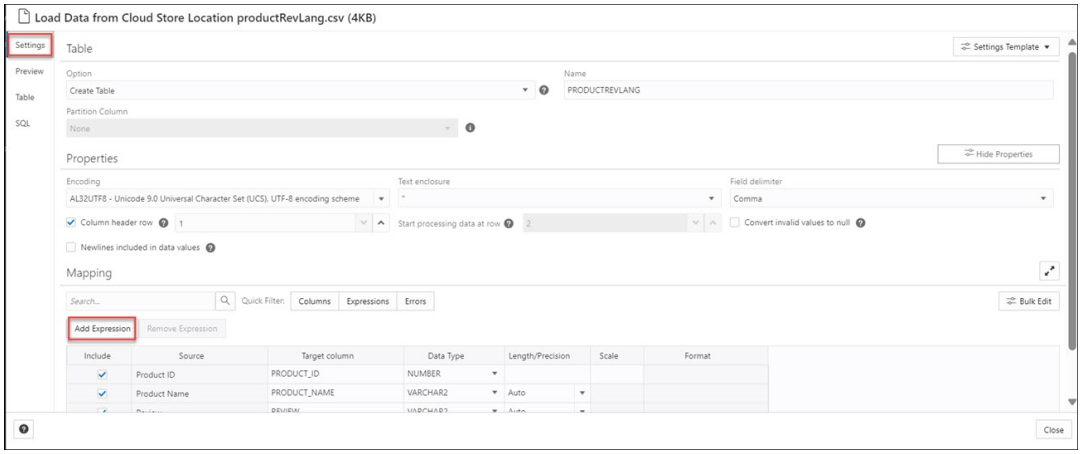
The Data Studio supports the detection of multiple languages.

To determine the language of the input text:

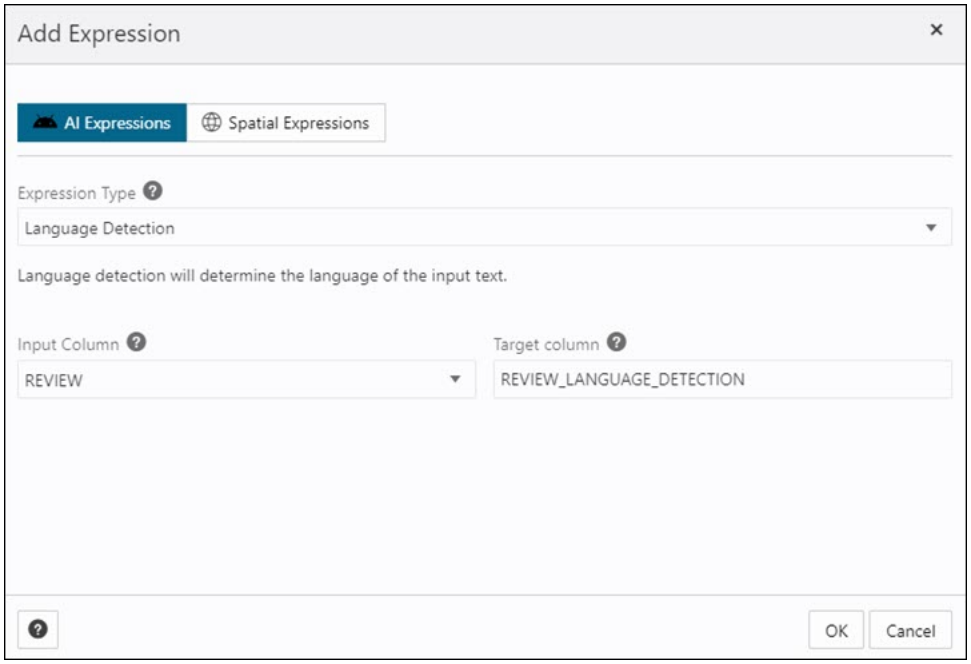
1. After you load data into the Data Load cart from cloud storage, you can view the file in the cart. Click the **Settings** icon.



2. Clicking the Setting icon opens a Load Data from Cloud Store Location wizard.
3. On the Settings tab of the wizard, click **Add Expression** under the Mapping section.

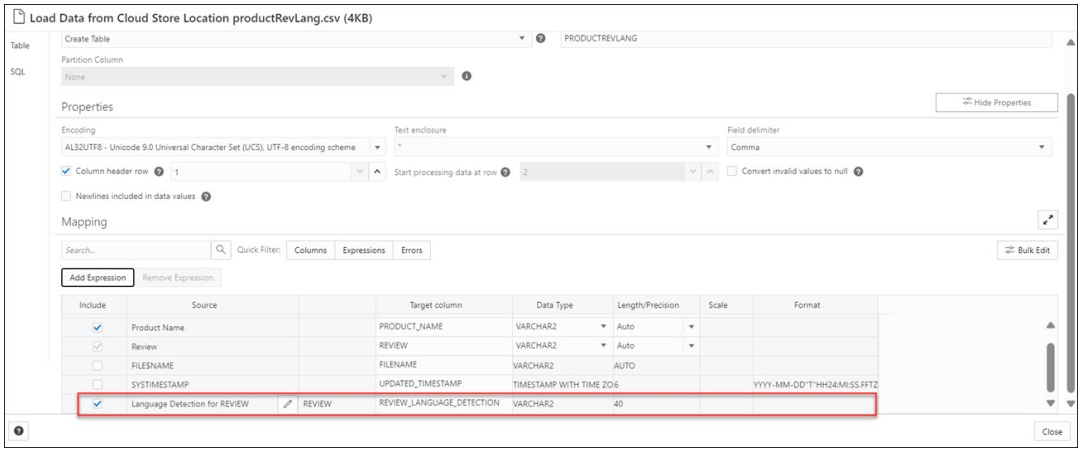


This opens the **Add Expression** dialog box. Select **AI Expressions**.



4. On the Add Expression dialog, specify the following fields:
   - **Expression Type:** From the Expression Type drop-down, select Language Detection.
   - **Input Column:** Select the column from the drop-down that you wish to analyze. In this case, we need to detect the language of the Review column. For example, *REVIEW*.
   - **Target column:** Enter the name of the newly created expression column. For example, *REVIEW_LANGUAGE_DETECTION.*

   Refer to the Parameters to Analyze Data for more details.
5. Click **OK**. You will see a new row added to the mapping grid .This row determines the output expression column generated by the OCI Language service.



6. Click **Close**.
7. Click **Start** in the Data Load menu cart to load data from local files. You will view a confirmation message asking if you wish to start the load from local files.
8. Click **Run** to confirm.

When the data load job completes, the **Table and View Loads** of the Data Load page display the details of the source table that is loaded into the tool. It displays the number of rows and columns and when the upload was complete.

- Output Data Generated from OCI Language Detection
  When you analyze columns using OCI Language Service model, the Data Studio generates a new expression column and saves the result in the updated table.

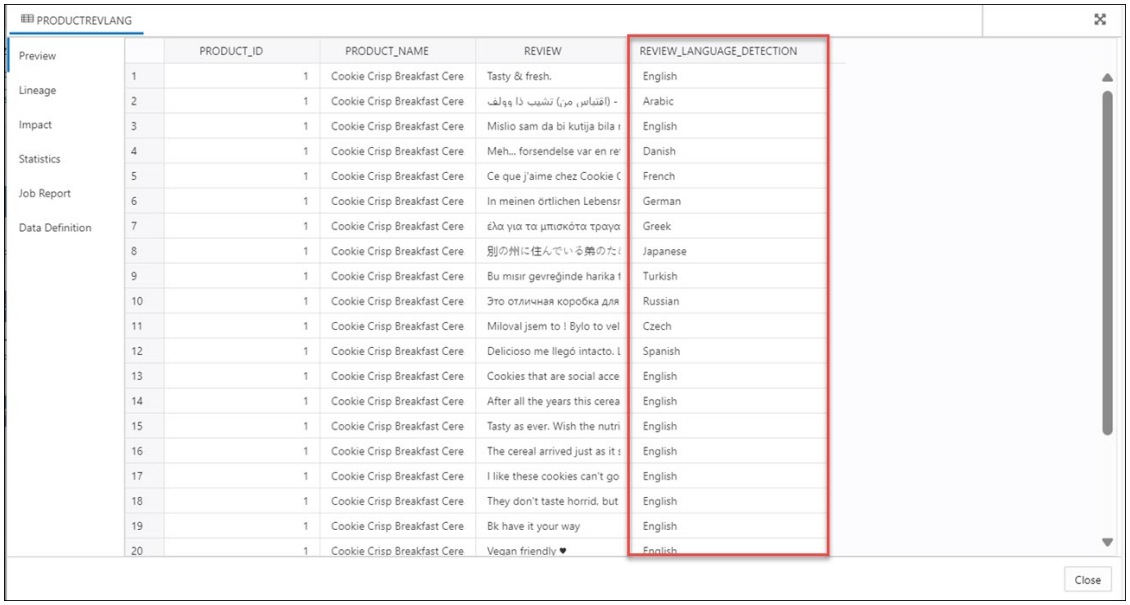
Output Data Generated from OCI Language Detection

When you analyze columns using OCI Language Service model, the Data Studio generates a new expression column and saves the result in the updated table.

To locate the generated expression column, from the Database Actions Launchpad, navigate to **Data Load**. Select the table you load under the **Table and View Loads** section.

Click the three vertical dots beside the load name, and click **Table** then select **View Details**.

For example, here's an output dataset from language detection of the Review column. Here, `REVIEW` is the target column, and `REVIEW_LANGUAGE_DETECTION` column displays the language detected from the input column.

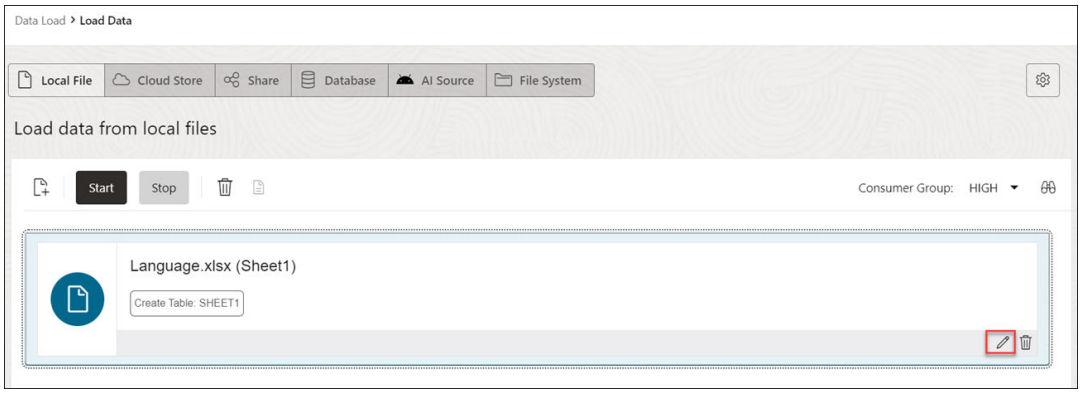


Perform Text Translation

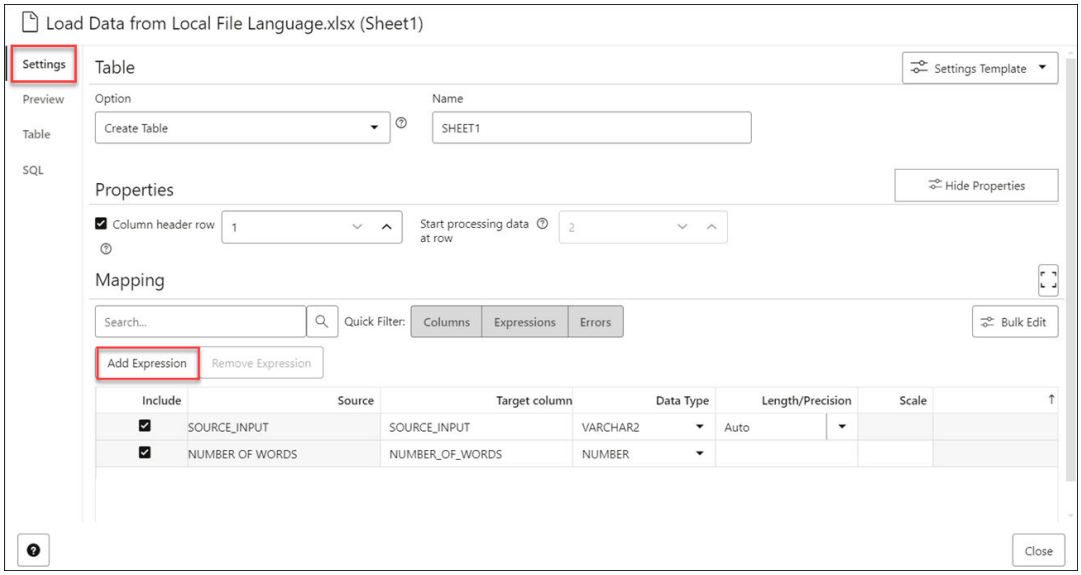
The Data Studio supports the translation of multiple languages supported by OCI Language Service.

To determine the translated input text:

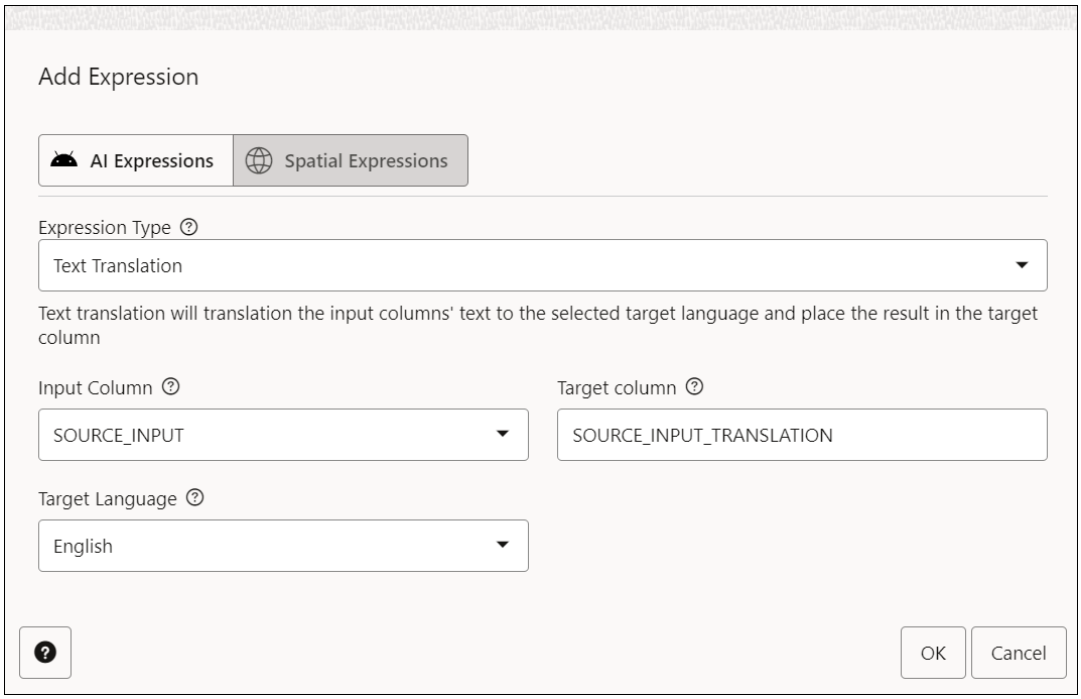
1. After you load data into the Data Load cart from local files, you can view the file in the cart. Click the **settings** icon.



2. Clicking the **Setting** icon opens a Load Data from Local File wizard.
3. On the Settings tab of the wizard, click **Add Expression** under the Mapping section.



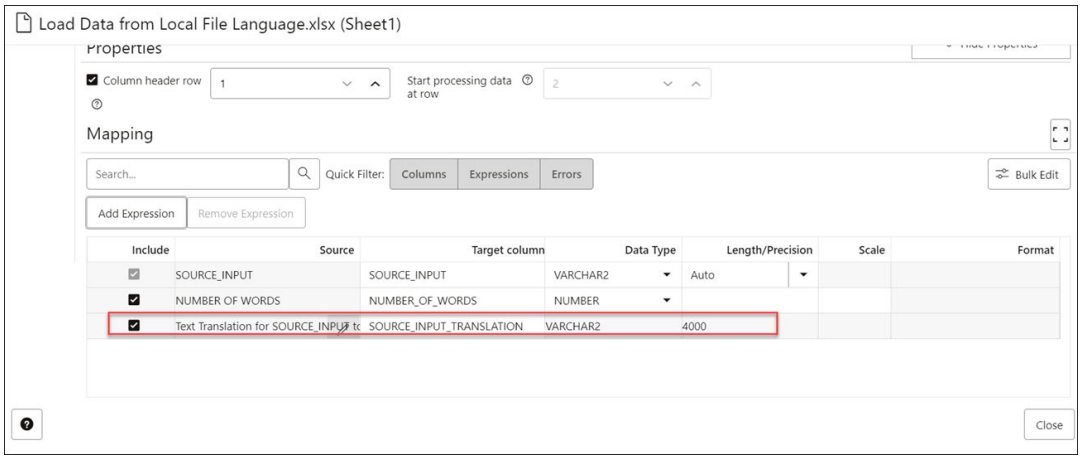
This opens the **Add Expression** dialog box. Select **AI Expressions**.



4. On the **Add Expression** dialog, specify the following fields:
   - **Expression Type:** From the Expression Type drop-down, select **Text Translation**.
   - **Input Column:** Select the column from the drop-down that you wish to analyze. In this case, we need to translate the text of the *SOURCE_INPUT* column.
   - **Target column:** Enter the name of the newly created expression column. For example, *SOURCE_INPUT_TRANSLATION.*
   - **Target Language:** Select the language of the text the tool translates from the drop-down. For example, *English*.

Refer to the Parameters to Analyze Data for more details.

5. Click **OK**. You will see a new row added to the mapping grid .This row determines the output expression column generated by the OCI Language service.



6. Click **Close**.
7. Click **Start** in the Data Load menu cart to load data from local files. You will view a confirmation message asking if you wish to start the load from local files.
8. Click **Run** to confirm.

When the data load job completes, the **Table and View Loads** of the Data Load page display the details of the source table that is loaded into the tool. It displays the number of rows and columns and when the upload was complete.

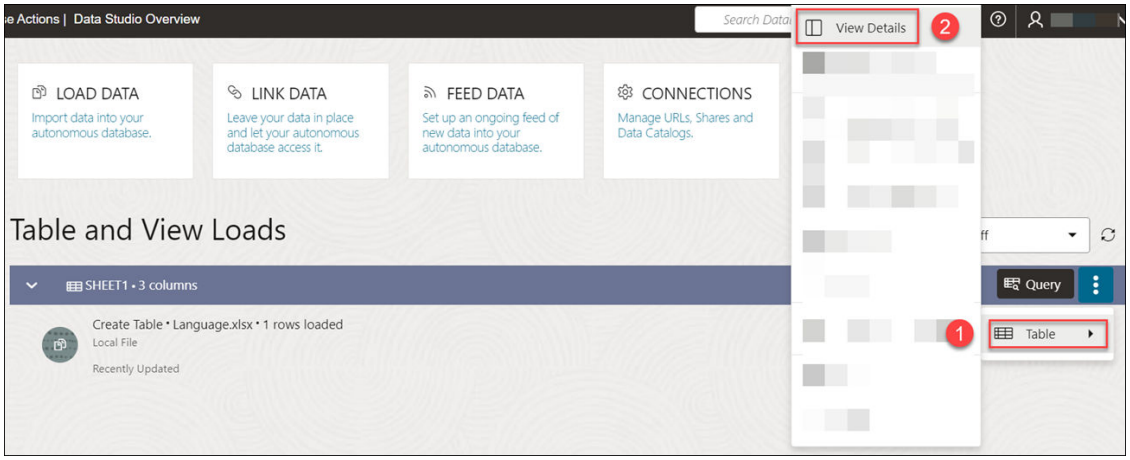
- Output Data Generated from OCI Text Translation
  When you analyze columns using OCI Language Service model, the Data Studio generates a new expression column and saves the result in the updated table.

Output Data Generated from OCI Text Translation

When you analyze columns using OCI Language Service model, the Data Studio generates a new expression column and saves the result in the updated table.

To locate the generated expression column, from the Database Actions Launchpad, navigate to **Data Load**. Select the table you load under the **Table and View Loads** section.

Click the three vertical dots beside the load name, and click **Table** then select **View Details**.



Use OCI Document Understanding to extract Tables from Images

You can utilize Oracle Cloud Infrastructure(OCI) Document Understanding Capabilities such as Table Extraction to detect tables from the data you load.

The supported file formats to use this feature are PDF, JPEG, JPG, PNG and Tag Image File Format (TIFF) formats.

For example, if a PDF document consists of all the employee details including their hiring date and salary, document understanding will identify the table and extract the table structure. It will extract the content of the table while maintaining the row and column relationship of cells.

- Overview of Table Extraction
  The Load Data tool supports Table Extraction in loading data from local files and loading data from cloud storage.
- Perform Table Extraction from Image
  Let us perform table extraction with a sample PDF.

Overview of Table Extraction

The Load Data tool supports Table Extraction in loading data from local files and loading data from cloud storage.

Table Extraction

OCI Document Understanding service allows customers to uncover insights in unstructured documents powered by deep learning models. It enables you to extract text, tables and identify document types among other great capabilities. Table Extraction identifies all the tables in a document and extracts the content in tabular format maintaining the row/column relationship.

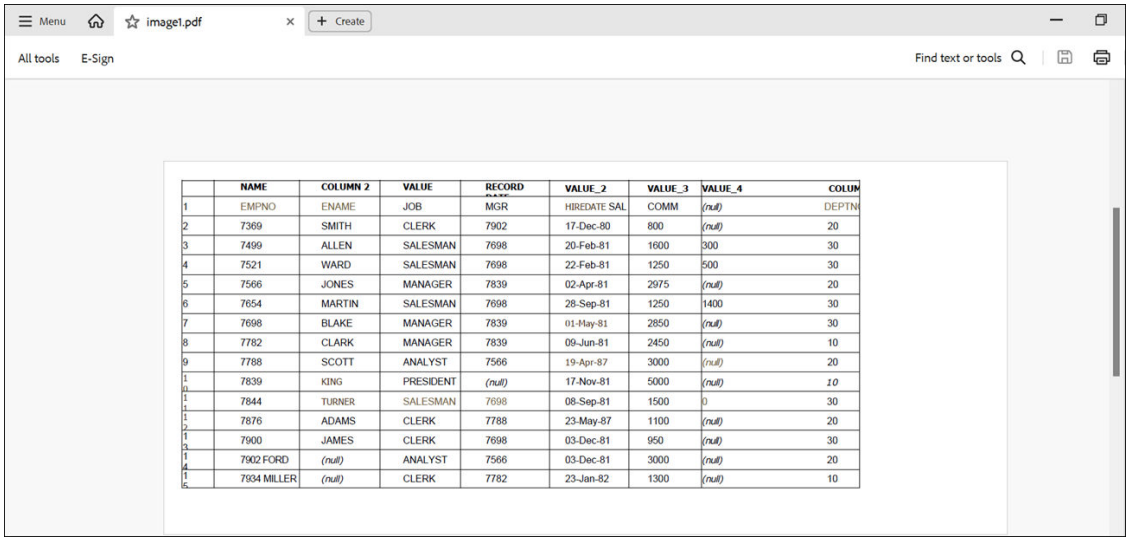
> **Note:**
>
> This service is available for source data in English Language only.

For more details, see Table Extraction.

Perform Table Extraction from Image

Let us perform table extraction with a sample PDF.

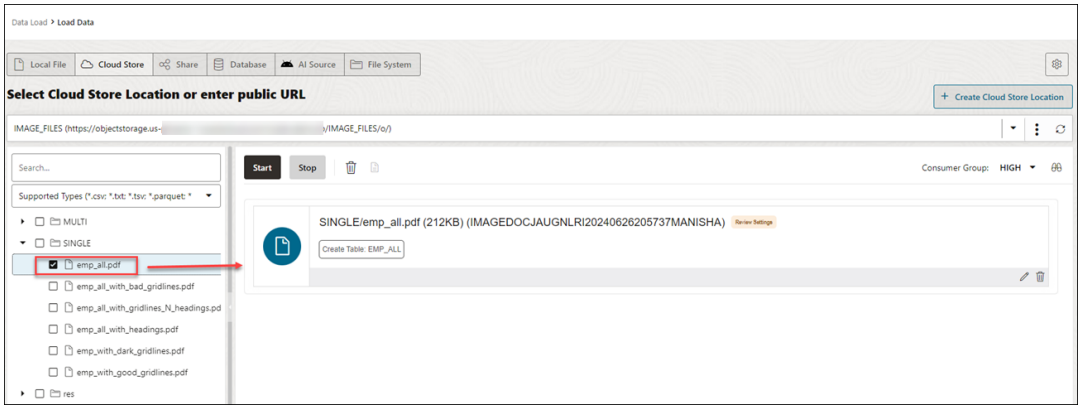
You can see a table in the PDF below.



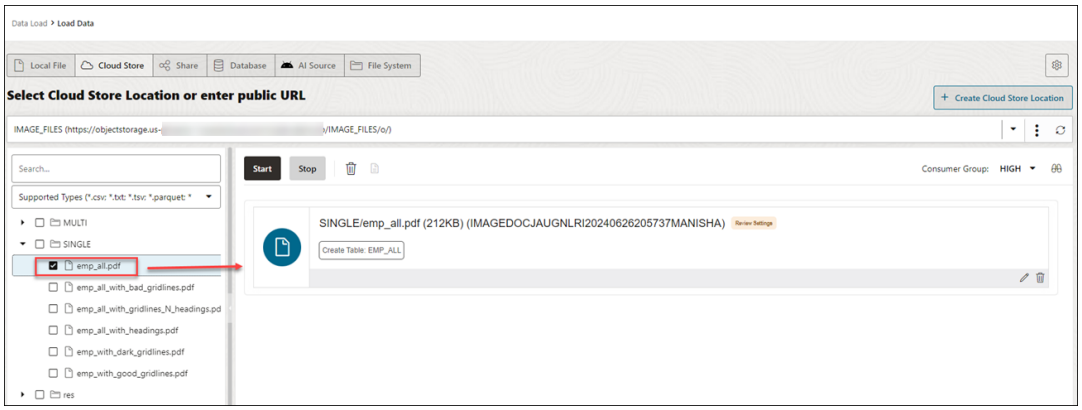
	NAME	COLUMN 2	VALUE	RECORD	VALUE_2	VALUE_3	VALUE_4	COLUM
1	EMPNO	ENAME	JOB	MGR	HIREDATE SAL	COMM	(null)	DEPTN
2	7369	SMITH	CLERK	7902	17-Dec-80	800	(null)	20
3	7499	ALLEN	SALESMAN	7698	20-Feb-81	1600	300	30
4	7521	WARD	SALESMAN	7698	22-Feb-81	1250	500	30
5	7566	JONES	MANAGER	7839	02-Apr-81	2975	(null)	20
6	7654	MARTIN	SALESMAN	7698	28-Sep-81	1250	1400	30
7	7698	BLAKE	MANAGER	7839	01-May-81	2850	(null)	30
8	7782	CLARK	MANAGER	7839	09-Jun-81	2450	(null)	10
9	7788	SCOTT	ANALYST	7566	19-Apr-87	3000	(null)	20
10	7839	KING	PRESIDENT	(null)	17-Nov-81	5000	(null)	10
11	7844	TURNER	SALESMAN	7698	08-Sep-81	1500	0	30
12	7876	ADAMS	CLERK	7788	23-May-87	1100	(null)	20
13	7900	JAMES	CLERK	7698	03-Dec-81	950	(null)	30
14	7902 FORD	(null)	ANALYST	7566	03-Dec-81	3000	(null)	20
15	7934 MILLER	(null)	CLERK	7782	23-Jan-82	1300	(null)	10

To extract a table from the PDF located in the cloud storage location:

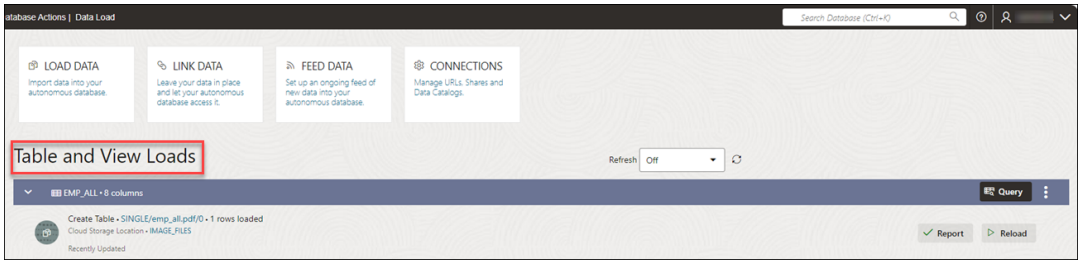
1. Load the PDF file into the Data Load cart from the cloud storage. You can drag and drop the file from the navigator to the cart.



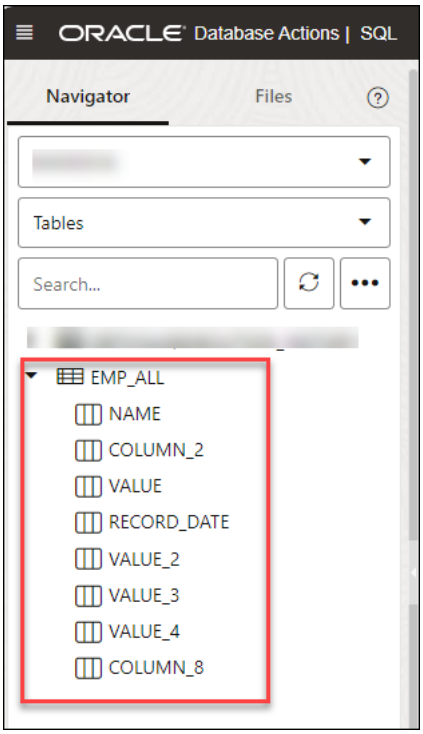
2. Click **Start** to load the data. You will receive a notification that confirms if you wish to start loading the data. Click **Run** to continue.



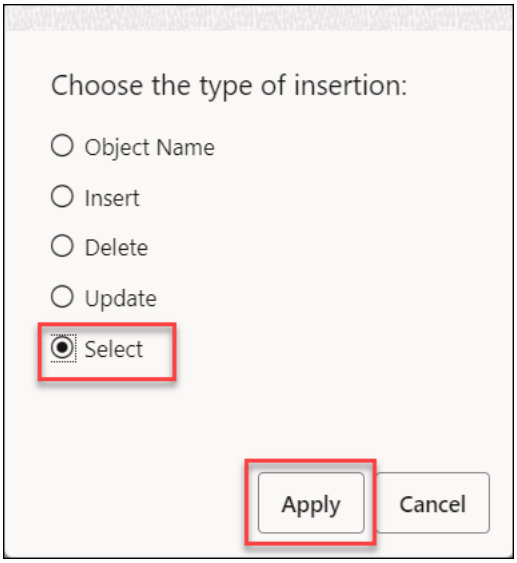
3. After the loading is complete, you can view the loaded file under Table and View Loads in the Data Load dashboard.



4. Extract the table from the PDF you loaded in the autonomous database:
   - Click **Database Actions** in the breadcrumb.
   - Select **Development**, then click **SQL**.
   - Select **Tables**, on the left Navigator, from the object drop-down. From the list of available tables, drag and drop the `EMP_ALL` table to the worksheet.

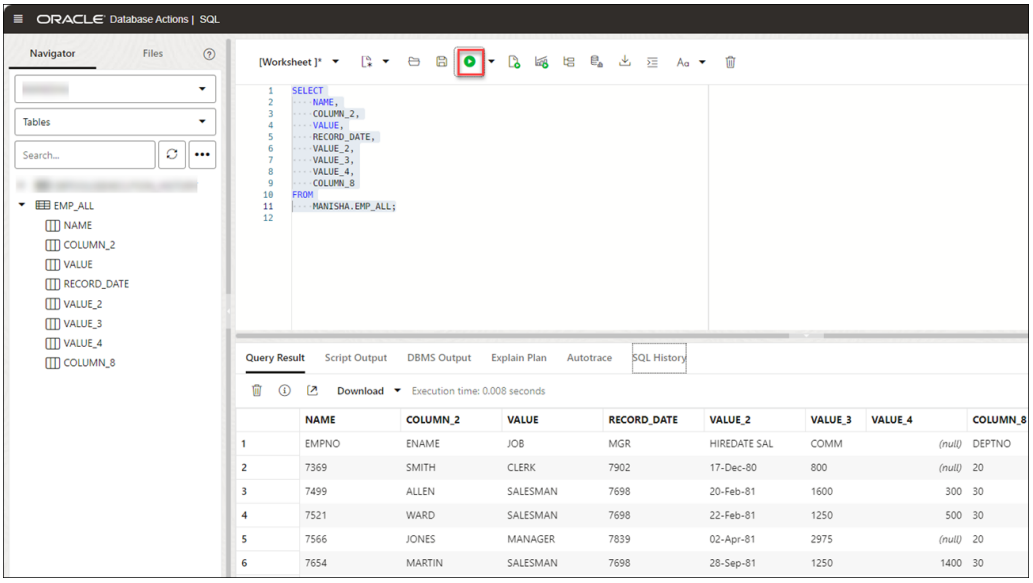


A dialog displays offering the types of available insertions.



Click **Select**, then **Apply**.

- Click **Run** (the arrow icon at the top of the worksheet). The tool displays the extracted table under the Query Results tab.



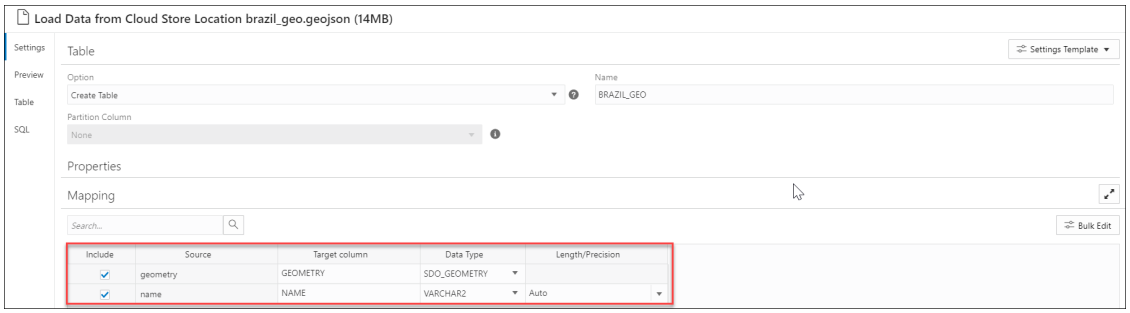
Use GeoJSON in Data Load

A GeoJSON object accommodates information about the specific geometry (e.g. Point, Line, Polygon, etc.) along with optional metadata (e.g. ID, etc.).

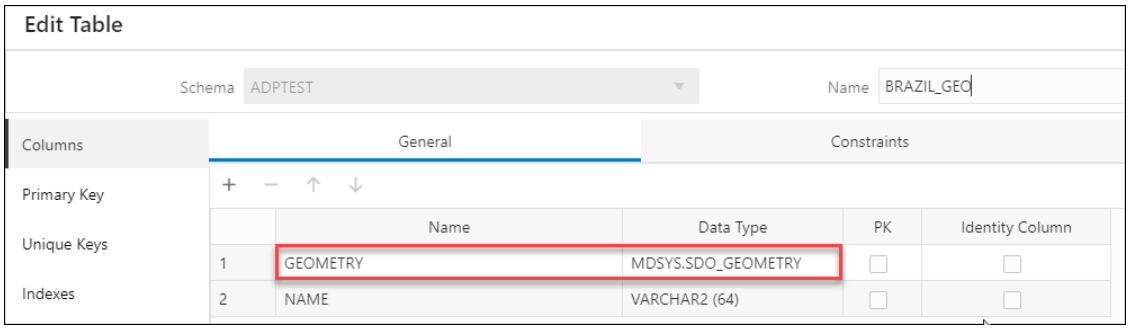
The extension for a GeoJSON file is *.`geojson`. You can load GeoJSON data into the Autonomous Database using the Data Load in Data Studio. If the table contains GeoJSON data, then the data is loaded in a column that projects GeoJSON data from the document set of SQL data type `SDO_GEOMETRY`.

Load a table with GeoJSON Data

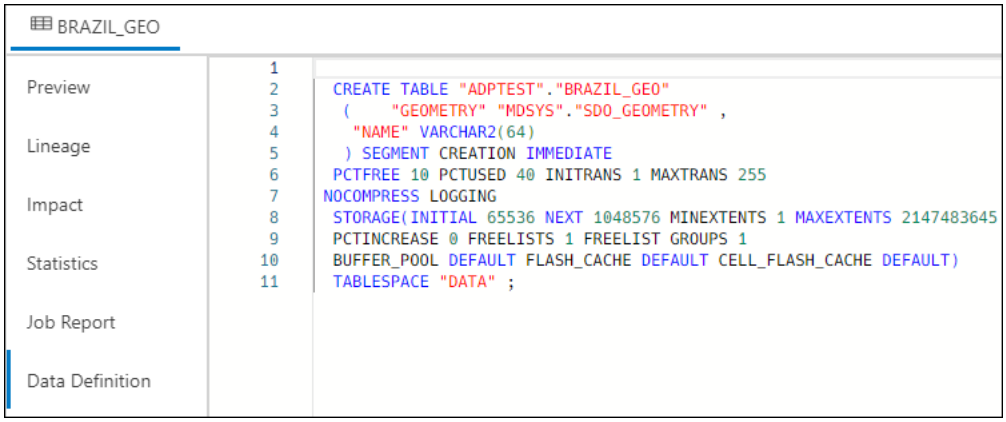
When you load a table in Data Studio with GeoJSON data and review its settings, you will see that it creates table `Brazil_Geo`, which has a column `geometry` of GeoJSON data.



After you load `BRAZIL_GEO` you will view that the tool loads GeoJSON object into a new column `geometry` with the data type `SDO_GEOMETRY`.



You can also view the same in the **Data Definition** tab when you **View Details** of the Table load after it is run.



- Add Spatial Expression for GeoJSON columns
  Oracle Spatial is designed to make spatial data management easier and more natural to users of location-enabled applications and geographic information system (GIS) applications.

- Parameters for Spatial Data Analysis
  When you invoke an **Add Expression** and select **Spatial Expressions** from the **Add Expression** wizard, you must configure the model using parameters.
- Perform Spatial Analysis
  You can load Spatial data using the Data Load tool from local files or cloud storage.
- Output Data Generated from Spatial Analysis
  When you analyze columns using Spatial Expressions, the Data Studio generates a new expression column and saves the result in the updated table.

Add Spatial Expression for GeoJSON columns

Oracle Spatial is designed to make spatial data management easier and more natural to users of location-enabled applications and geographic information system (GIS) applications.

Once spatial data is stored in an Oracle database, it can be easily manipulated, retrieved, and related to all other data stored in the database.

The spatial data model in Oracle Spatial is a hierarchical structure consisting of elements, geometries, and layers. Layers are composed of geometries, which in turn are made up of elements.

Tolerance is used to associate a level of precision with spatial data. **Tolerance** reflects the distance that two points can be apart and still be considered the same (for example, to accommodate rounding errors). The tolerance value must be a positive number greater than zero. See Tolerance in Spatial Data for more information.

Parameters for Spatial Data Analysis

When you invoke an **Add Expression** and select **Spatial Expressions** from the **Add Expression** wizard, you must configure the model using parameters.

Table 3-4 Parameters for Spatial Data Analysis

Parameter	Description
Expression Type	Select any one of the following spatial attributes such as **Length** or **Area**: • **Length**: Selecting this attribute returns and stores the perimeter of a two-dimensional polygon in the mapping table with a target column of type `GEOMETRY_LENGTH`. • **Area**: Selecting this attribute returns and stores the area of a two-dimensional polygon in the mapping table with a target column of type `GEOMETRY_AREA`.
Input Column	Select the column which you wish to analyze for Spatial Data: You will by default view a `GEOMETRY` column in Source mapped to the `SDO_GEOMETRY` column value. This also means that the table is spatially enabled.

Table 3-4 (Cont.) Parameters for Spatial Data Analysis

Parameter	Description
Target Column	This is the newly created expression column defined in the target table which stores the result of the Spatial Expression and depends on the expression type value you select: • `GEOMETRY_LENGTH`: If you select Length as your Expression type, you can view the target column of type `GEOMETRY_LENGTH` which displays the geometry's length value. • `GEOMETRY_AREA`:If you select Area as your Expression type, you can view the target column of type `GEOMETRY_AREA` which displays the geometry's area value.

Perform Spatial Analysis

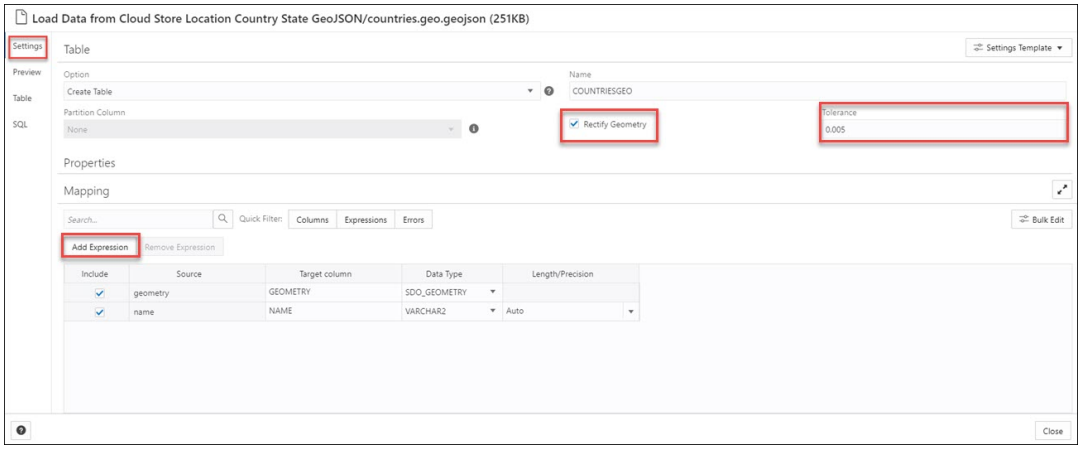
You can load Spatial data using the Data Load tool from local files or cloud storage.

To perform Spatial Analysis on GeoJSON data:

1. After you load data into the Data Load cart from local files or cloud storage, you can view the file in the cart. Click the **Settings** icon.

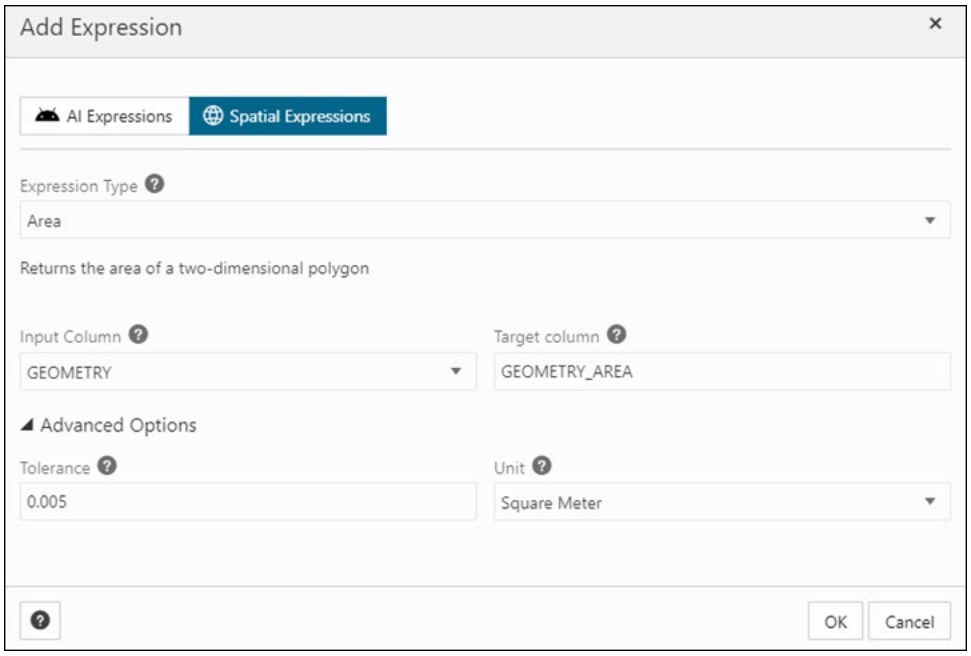


2. On the Settings tab of the wizard, select **Rectify Geometry** to fix invalid geometries. This icon corrects the most common errors such as duplicate points, polygon orientation errors, polygon construction errors etc. If it detects any uncorrectable error, it fails with an `ORA-13199: the given geometry cannot be rectified` exception.



> **Note:**
>
> You cannot fix geometries having wrong orientation with the. **Rectify Geometry** icon.

3. Specify the **Tolerance** value. In the above example, we specify 0.005 as tolerance, i.e., 5mm.
4. Click **Add Expression** under the Mapping section and select **Spatial Expressions**.



On the **Add Expression** dialog, specify the following fields:

- **Expression Type**: From the Expression Type drop-down, select **Area** to display the area of a two-dimensional polygon. You can also select **Length** to display the perimeter of a two-dimensional polygon.
- **Input Column**: You will view a `GEOMETRY` column by default. This source column stores the area of the geometry. In case you select **Length** as an expression type, this source column stores the length of the geometry.
- **Target column**: Specifies the name of the newly created special expression column. For example, `GEOMETRY_AREA` for **Area** type expression parameter and `GEOMETRY_LENGTH` for **Length** type expression parameter.

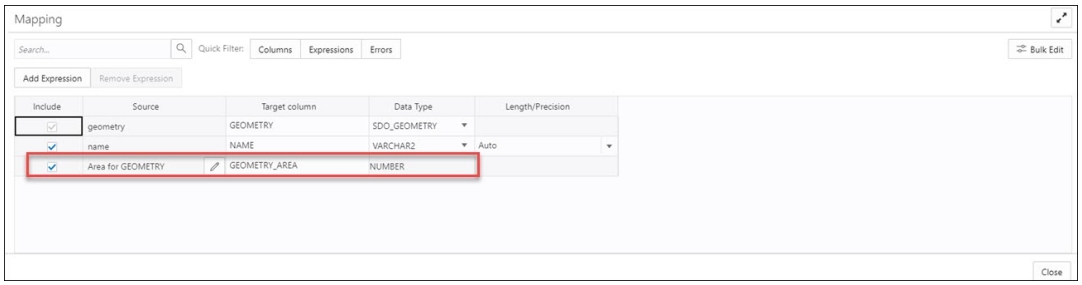
Under the Advanced Options section, specify the following field values:

- **Tolerance**: Enter the tolerance value. For example, `0.005`.
- **Unit**: Select a unit of measure of the tolerance from the drop-down. For example, `Square meter`.

> **Note:**
>
> If the file you load does not contain spatial data, you will view the following error on the **Add Expression** wizard "`No input columns have the SDO_GEOMETRY data type for the currently selected expression type of Area`".

Refer to the Parameters for Spatial Data Analysis for more details.

5. Click **OK**. You will see a new row added to the mapping grid. This row determines the output spatial expression column generated by the OCI Language service. In this example, you can view a new `GEOMETRY_AREA` will be created in the target table after we load the table.



For more information on Oracle Spatial, refer to Oracle Spatial Developer's Guide.

When the data load job completes, the **Table and View Loads** of the Data Load page display the details of the source table that is loaded into the tool. It displays the number of rows and columns and when the upload was complete.

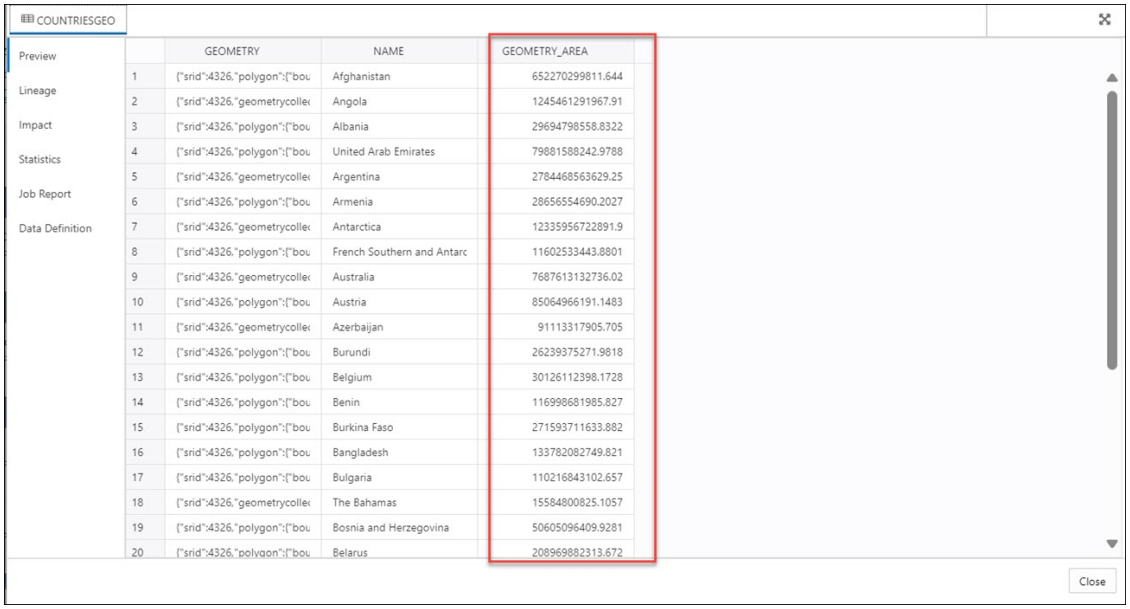
Output Data Generated from Spatial Analysis

When you analyze columns using Spatial Expressions, the Data Studio generates a new expression column and saves the result in the updated table.

To locate the generated expression column, from the Database Actions Launchpad, navigate to **Data Load**. Select the table you load under the **Table and View Loads** section.

Click the three vertical dots beside the load name, and click **Table** then select **View Details**.

This opens the Preview tab of the data load which displays the updated source file. For example, here is an output dataset from spatial analysis of the `COUNTRIESGEO.GeoJSON` file. Here, `GEOMETRY_AREA` is the spatial analysis of the input column. This column displays the area of the two-dimensional polygons.

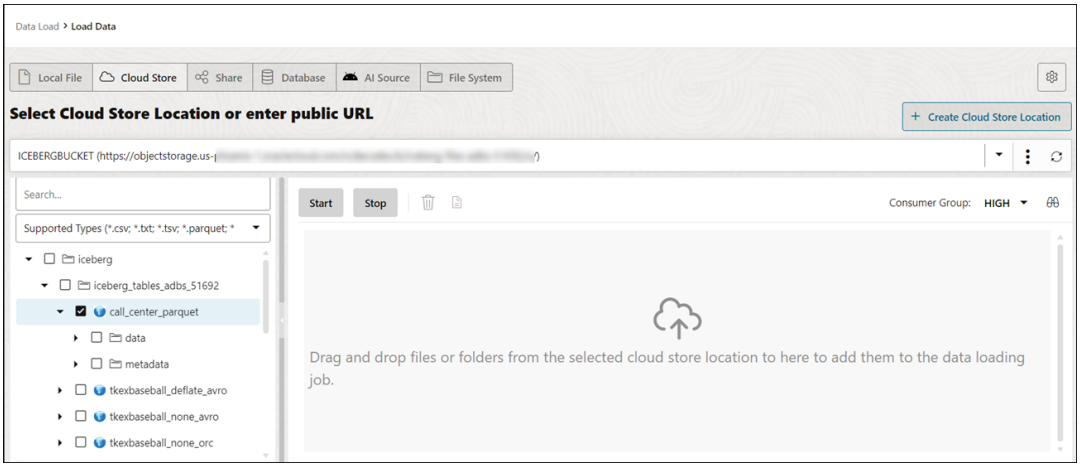


Load Apache Iceberg Tables

Data Studio supports loading and linking Iceberg tables from the object store.

The tool can load multiple versions of the same file.

1. Select the Cloud Store Location that has iceberg tables defined. In this example, we will select the `ICEBERGBUCKET` stored in OCI.

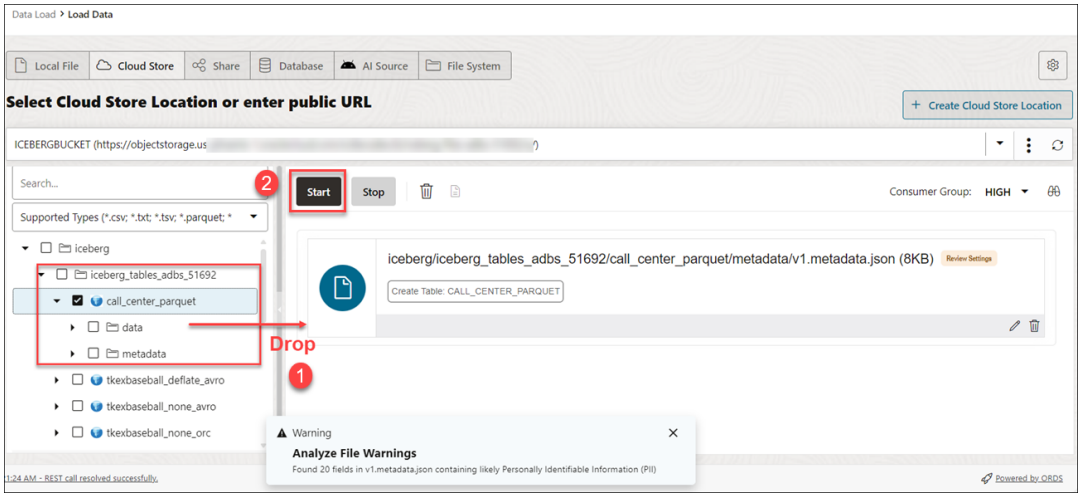


You can detect the Iceberg files with the Iceberg symbol



beside its name.

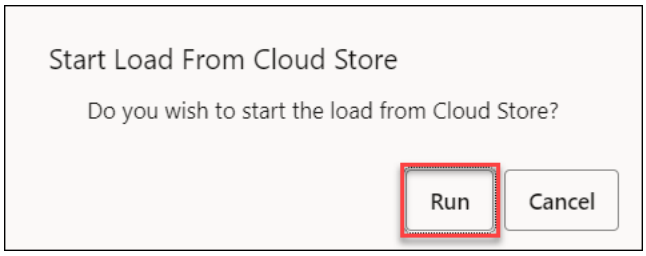
2. Drag the folder to the file navigator on the left and drop it into the cart on the right. The folder also consists of the metadata JSON file that keeps track of the table snapshots, partitioning scheme, and schema information.
When you add the folder to the cart, a warning about the number of files consisting of a Personal Identifiable Information (PII) prompt is displayed.



Note:

: If you load multiple files instead of a folder, the tool asks if you want to load all the objects from the multiple source files into a single target table. Click **OK** to continue or Escape to cancel.

3. Click **Start** in the data link cart menu bar. You will view a notification that says if you wish to start the load from the cloud store. Select **Run** to continue.

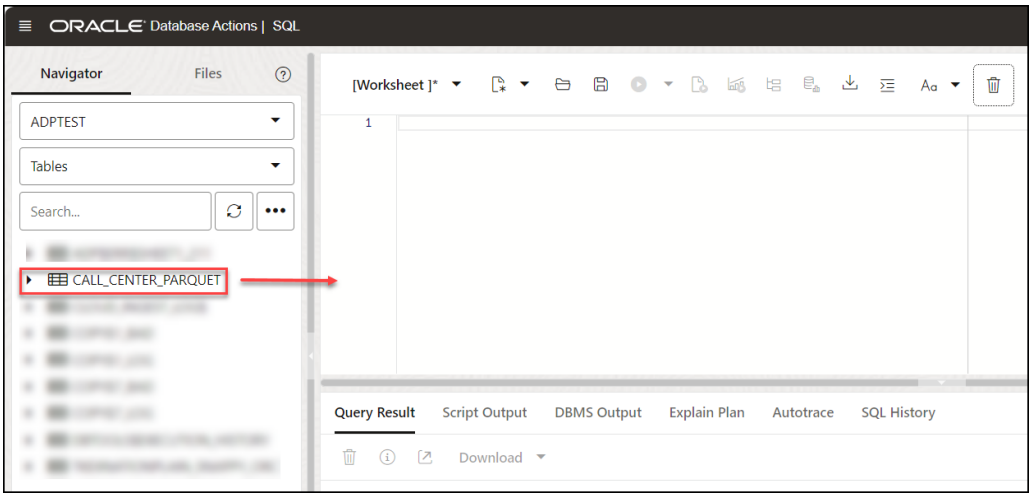


4. You will view the loading progress under Table and View Loads. After loading, you will view the Report and Reload icons on the table load.

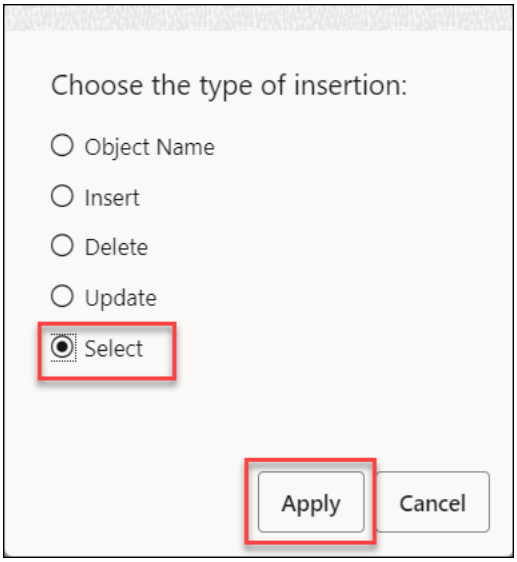


5. Query Iceberg tables you loaded in the Autonomous Database:
   - Click **Database Actions** in the breadcrumb.

- Select **Development**, then click **SQL**.
- Select **Tables**, on the left Navigator, from the object drop-down. From the list of available views, drag and drop the **CALL_CENTER_PARQUET** table to the worksheet.

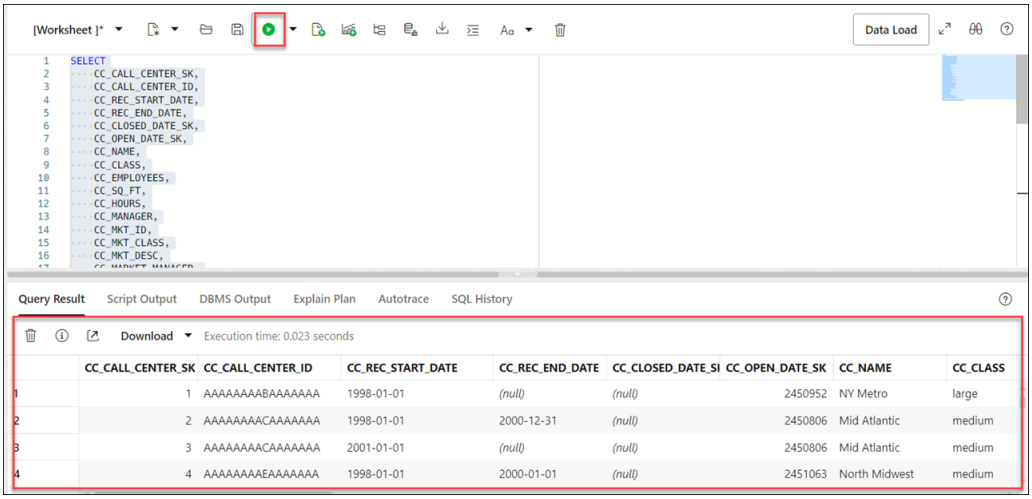


- A dialog displays offering the types of available insertions.



Click **Select**, then **Apply**.

- Click **Run** (the arrow icon at the top of the worksheet). The tool displays the iceberg data under Query Results.



Loading Data from File System

You can load files from file system directories to your Autonomous Database.

You can set filters on the data for a table to load only the specified data. For example, to limit the files to only those that are CSV files, enter *.CSV in the file extension filter.

Configure and run a data load job from the Load Cloud Object page. To open that page:

1. On the Data Studio tab and select **Data Load**. You will view the Data Load dashboard.
2. Click **LOAD DATA** and select the **File System** option.

On the top of the page you need to select the directory from where you need to load the files. On the left side of the page is a navigator pane, where you choose the files in the directory containing the data. On the right of the page is the data load cart, where you stage the files and folders for the data load job. You can set options for the data load job before running it. The Autonomous Database comes with predefined CPU/IO shares assigned to different consumer groups. You can set the consumer group to either low, medium or high while executing a data load job depending on your workload. To load files from a directory into your database, do the following:

1. Prepare the Data Load Job: Refer to the Prepare the Data Load Job section for more details.
2. Add Files or Folders for the Data Load Job: Refer to the Add Files or Folders for the Data Load Job section for more details.
3. Enter Details for the Data Load Job: Refer to the Enter Details for the Data Load Job for more details.
4. Run the Data Load Job: Refer to the Run the Data Load Job section for more details.
5. View Details About the Data Load Job After It Is Run: Refer to the View Details About the Data Load Job After It Is Run section for more details.
6. View the Table Resulting from the Data Load Job: Refer to the View the Table Resulting from the Data Load Job section for more details.

- Create Directories in Database Actions
  In the Autonomous Database, there is a preconfigured `data_pump_dir` where you can store files. You can also create directories, drop directories, and attach network file systems.

Create Directories in Database Actions

In the Autonomous Database, there is a preconfigured `data_pump_dir` where you can store files. You can also create directories, drop directories, and attach network file systems.

For example, you can use the `CREATE DIRECTORY` command to create additional directories. Use the database `DROP DIRECTORY` command to drop directories and use `DBMS_CLOUD.LIST_FILES` to list the contents of a directory.

Create a Directory

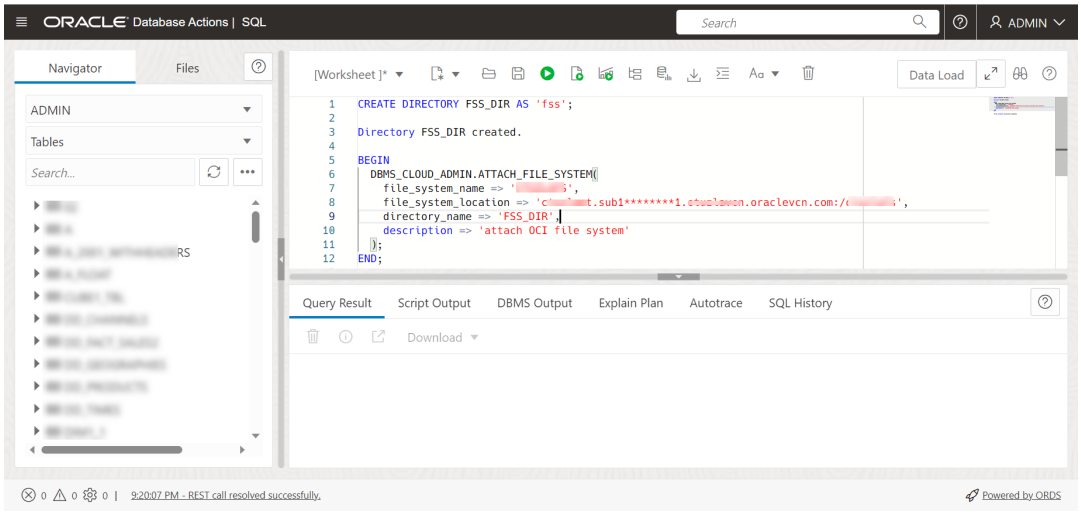
To add a directory, you must have the `CREATE ANY DIRECTORY` system privilege. The ADMIN user is granted the `CREATE ANY DIRECTORY` system privilege. The ADMIN user can grant `CREATE ANY DIRECTORY` system privileges to other users.

See CREATE DIRECTORY for more information.

> **Note:**
>
> - `CREATE DIRECTORY` creates the database directory object in the database and also creates the file system directory. For example the directory path could be:
>
>   ```
>   /u03/dbfs/7C149E35BB1000A45FD/data/stage
>   ```
>
> - You can create a directory in the root file system to see all the files with the following commands:
>
>   ```
>   CREATE OR REPLACE DIRECTORY ROOT_DIR AS '';
>   ```
>
>   After you create the `ROOT_DIR` directory, use the following command to list all files:
>
>   ```
>   SELECT * FROM DBMS_CLOUD.list_files('ROOT_DIR');
>   ```
>
> - To run `DBMS_CLOUD.LIST_FILES` with a user other than ADMIN you need to grant read privileges on the directory to that user.

Let us demonstrate how to create a directory and access it from Data Studio:

- **Create a directory in Database Actions:**
  Login to your Database Actions instance and select the SQL card under **Development**. You can view the SQL worksheet. Now create a directory and attach a file system name of your choice to the directory you create. In the below given example, `FSS_DIR` is the name of the directory.



Run the above command. The above command gives the following output:

```
PL/SQL procedure successfully completed.
```

- **Attach the file system**
  Attach your file system with the name of your choice to the `FSS_DIR` directory via the DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM procedure.

```
BEGIN
  DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM(
    file_system_name => '********',
    file_system_location => '*******.sub1********1.********.oraclevcn.com:/
********',
    directory_name => 'FSS_DIR',
    description => 'attach OCI file system'
  );
END;
/
```

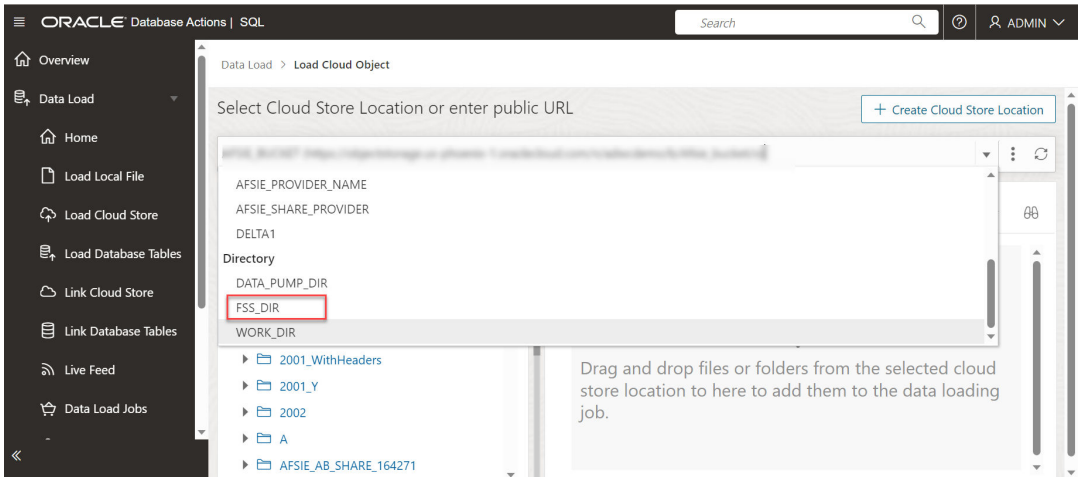
Run the above command to view the following output:

```
PL/SQL procedure successfully completed.
```

- To view the attached file system, run the following command:
  `SELECT file_system_name, file_system_location, directory_path FROM dba_cloud_file_systems;`

  You will view the file system name, file system location and directory path.
- You can view the new directory along with the files attached to it by navigating to Load Cloud Object under Data Load menu of the Data Studio tool. Click the **Select Cloud Store**

location drop-down.

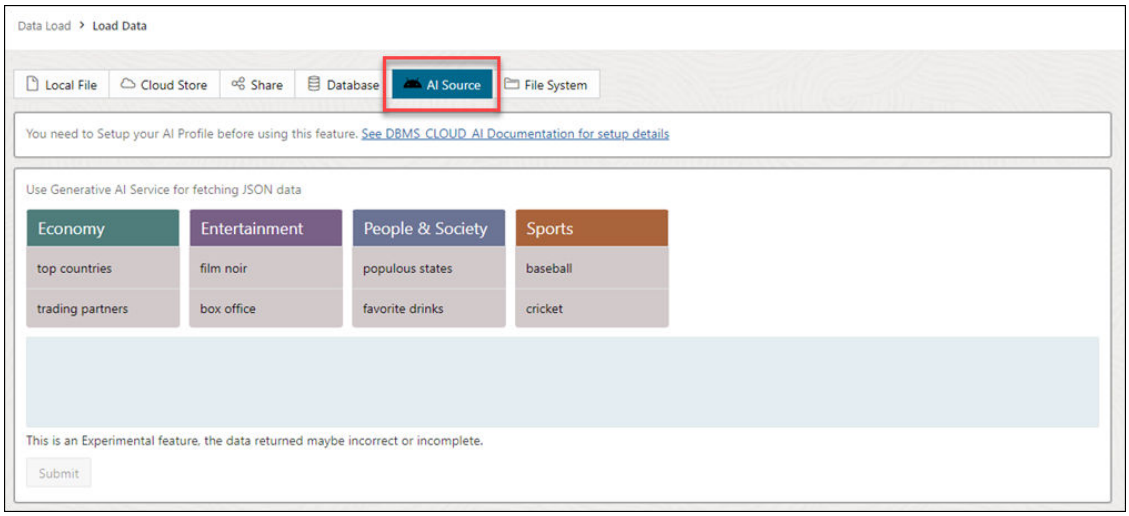


You can view the newly created directory `FSS_DIR`. You can load data from the file system directories to the autonomous database using the Data Load tool. See Loading Data from File System.

Loading Data from AI Source

You can use Data Studio tools to load data from AI source.

On the Data Load dashboard, click **Load Data** card, and select **AI Source**.



You need to perform a pre-requisite of setting up your AI profile before using this feature. See Configure DBMS_CLOUD_AI Package for details on the setup of this feature.

To set an AI profile, you must have:

- An OpenAI account, or a Cohere account, or Azure OpenAI service
- An access to `DBMS_CLOUD_AI` package.

Follow the steps in Use DBMS_CLOUD_AI to Configure AI Profiles section to create and configure your AI profile and utilize this feature.

Loading Data from Share

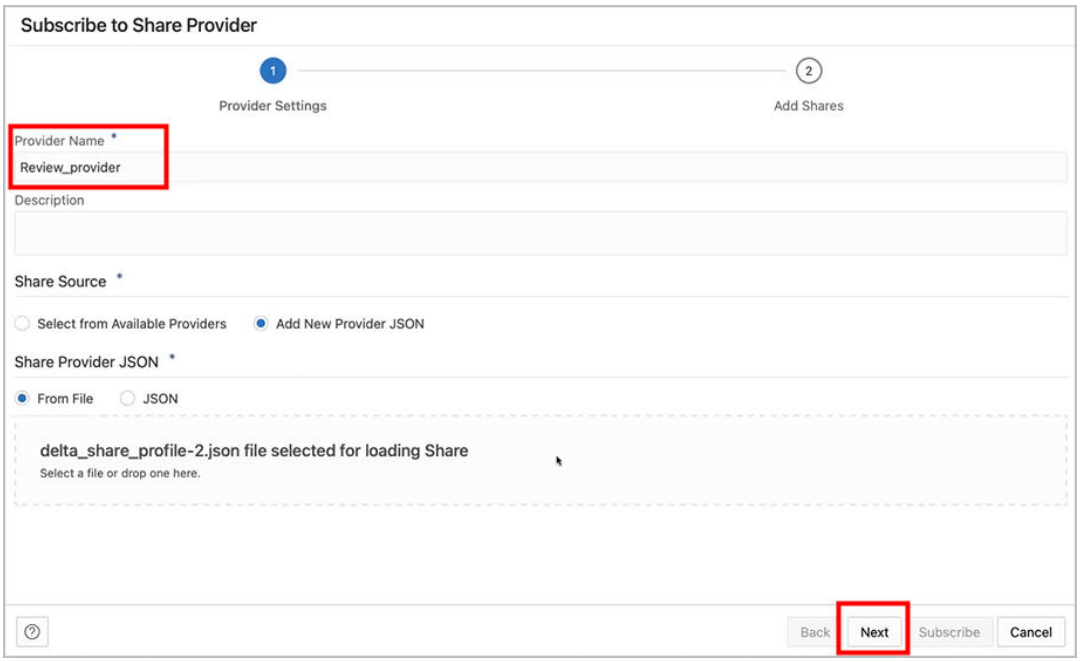
You can select tables from a Share. You need to subscribe and access provided data share.

To load tables from a share, click **Load Data** on the Data Load Dashboard. Click **Share** on the Load Data page. Click **+ Subscribe to Share Provider** to subscribe to a Share Provider.

Subscribe to the Data Share Provider

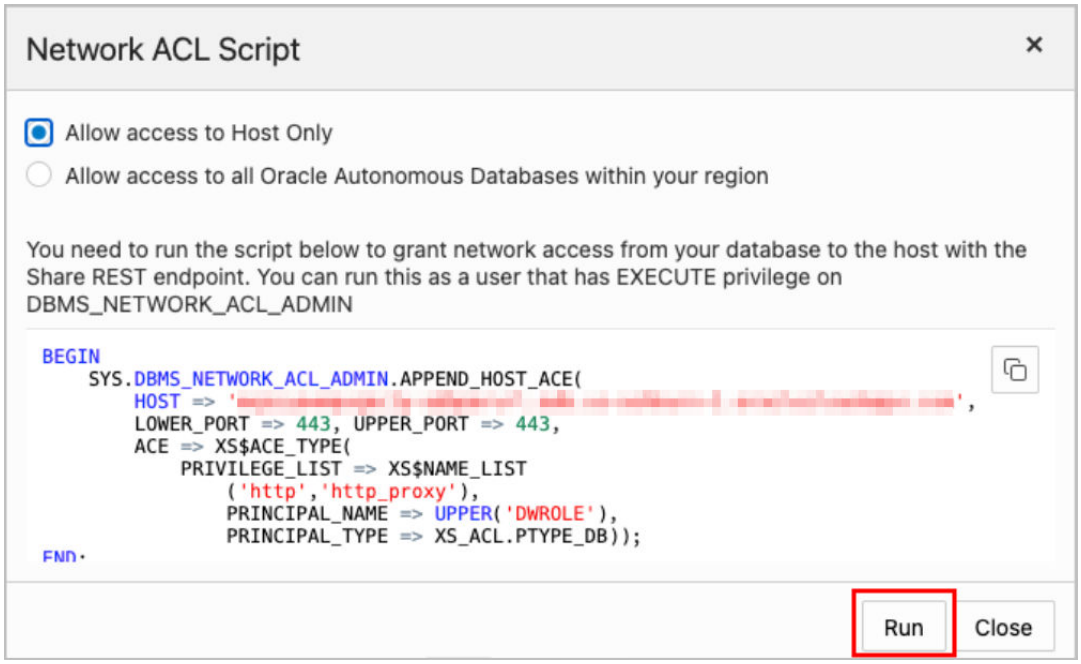
To subscribe, you need to use the information contained in the uploaded JSON profile you received from the share provider.

1. On the Subscribe to Data Share Provider, enter the **Provider Name**.

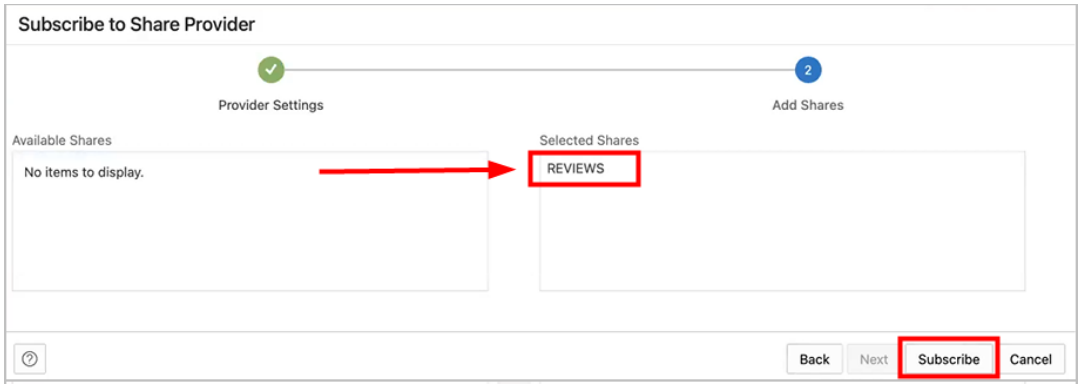


2. Select Add New Provider JSON and click on the Delta Share Profile JSON to drag and drop the JSON profile.
3. Click **Next** to progress to the Add Shares tab.
4. Select the level of network access you want to allow from your database to the host with the Share REST endpoint, and click **Run**. In this example, **Allow access to Host Only** is

selected.

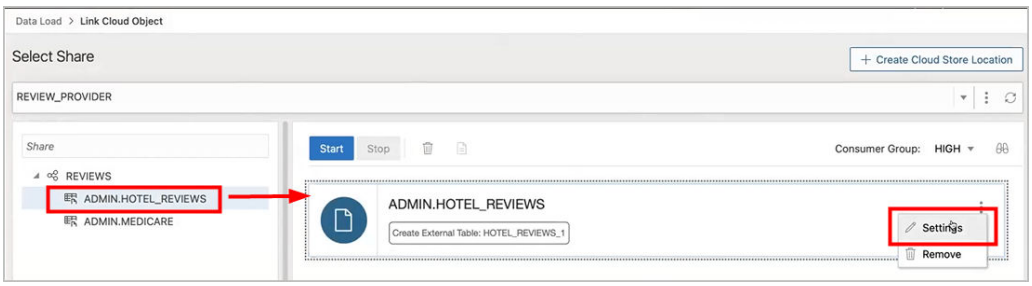


5. To register the shares available to you, move shares from **Available Shares** to **Selected Shares** and click **Subscribe**.
   The screenshot below shows the **REVIEWS** share moved from **Available Share**s to **Selected Shares** before clicking **Subscribe**.



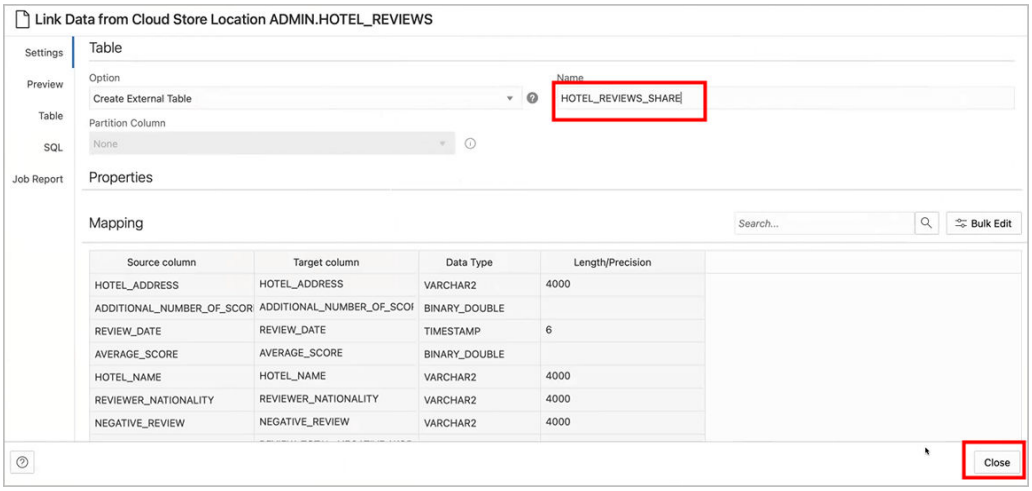
6. Create external tables derived from tables selected from the data share.

   a. Drag and drop tables from the selected share. You can optionally click **settings** to view the table details.
      In this example, the only table selected is **HOTEL_REVIEWS**.

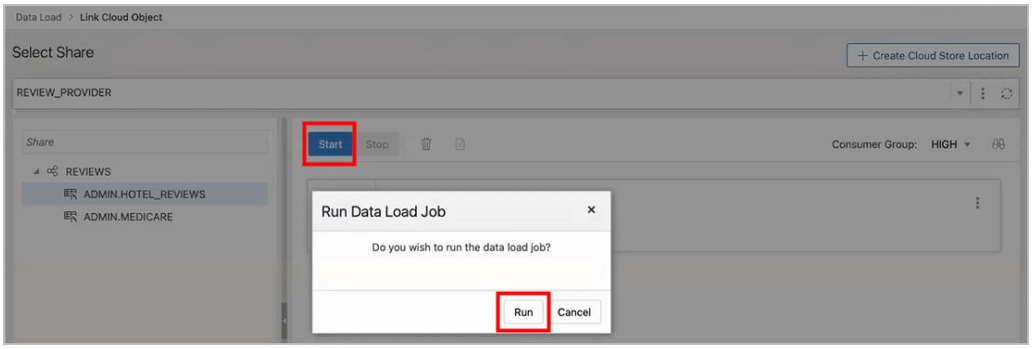


   b. You can optionally change the name of your table and click **Close**.

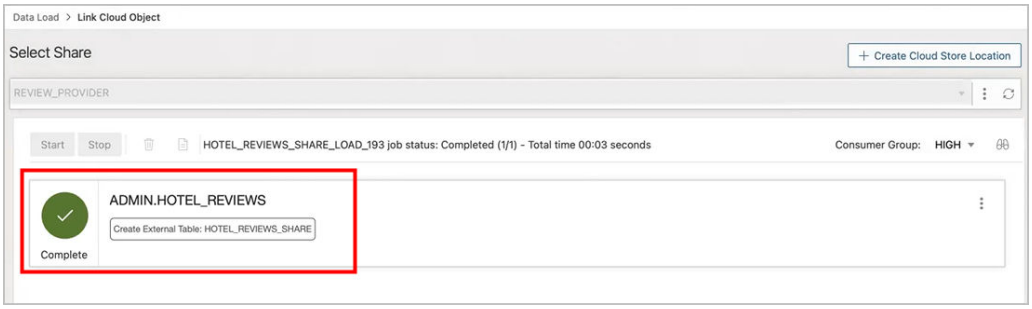
In this example, the name is changed from **HOTEL_REVIEWS** to **HOTEL_REVIEWS_SHARE**.



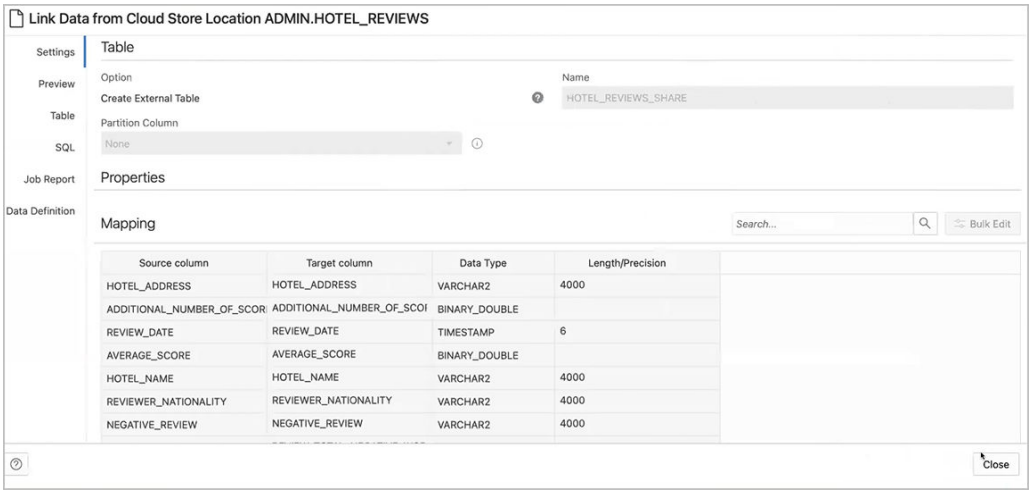
c. Create the external table by clicking **Start**, on the Select Share page, and then clicking **Run** on the Run Data Load Job dialog.



d. When the external tables are created the Complete message is displayed.



e. Click on the external table name to view the table details.

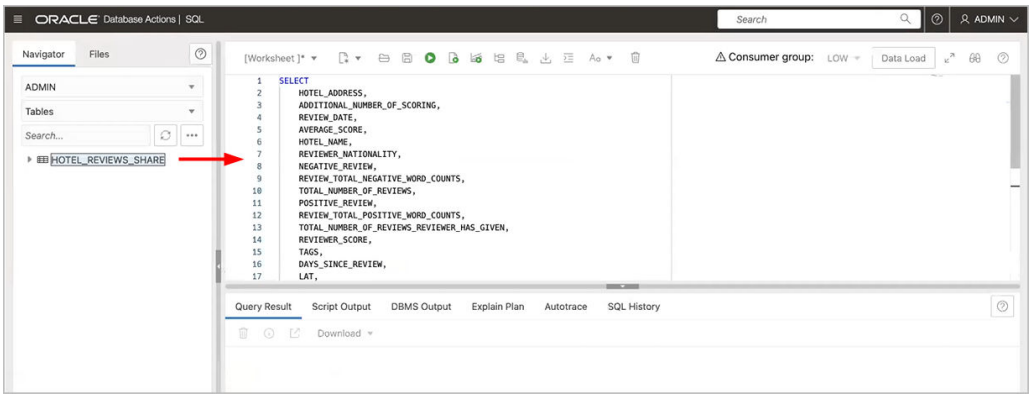


7. View the created tables from Database Actions.

   a. Click on **Database Actions**, in the breadcrumb, to go back to the Database Actions launchpad.

   b. Click on the **SQL** tile.

   c. Select the external table, drag and drop it into the worksheet.
   The SQL Select statement for the table appears. This SQL statement can be run to consume the shared data.



8. Drag and drop tables from selected share

Create Live Feed from Data Load

The Data load tool loads the data from folders in cloud object stores and enables it to schedule repeated data loads in real time. This is the creation of Live Feed from a data load job.

Once the data load is complete, you can create a Live Feed from a cart item that loaded an object store folder using the Create Table or Drop Table and Create New Table options.

To create a live feed from Data Load:

1. Click **Selector** to display the navigation menu. Under Data Studio, select **Data Load**.
2. Select the **Load Data** tile to load data from various sources such as local files, databases, cloud store, directories, and Shares.
3. Click **Cloud Stores**to load objects from URLs or Cloud store links.

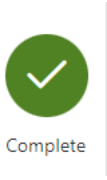
4. Select the cloud store location from the drop-down. If are not able to view the cloud store location, select **Create Cloud Store Location** to create a new cloud store location. Follow the steps described in the Create Oracle Cloud Infrastructure Native Credentials if you do not have a cloud location available.
5. After you select the cloud store location, you can view the list of folders and files present in the cloud storage. Add files from the cloud store to the data load cart, where you can edit the details of the load job.

> **Note:**
>
> The Data load tool does not support the creation of a live feed from a loaded cart item that consists of a single file in CSV, XLS, XLSX, TSV, TXT, XML, JSON, and AVRO format or of a folder that contains a file in XLSX format.

6. To add the folders, drag a folder from the file navigator on the left and drop them into the cart on the right. When you add the folder to the cart, a prompt is displayed that asks if you want to load all the objects from the multiple source files into a single target table. Click **Yes** to continue or **No** to cancel. You can add multiple folders to the cart, the data represented by each card will be loaded into a separate table, but all the items in the cart will be processed as part of the same data load job.
7. Select **Settings** (pencil icon) from the data load cart to enter the details about the data load job.
8. In the **Settings** tab of the Load Data from the Cloud Store Location, you can select **Create Table** or **Drop Table and Create New Table** from the **Option** drop-down.

> **Note:**
>
> The Live feed tool works with the Data load job only if you create a table and insert the data into a new table or drop the existing table and insert the data into a new table.

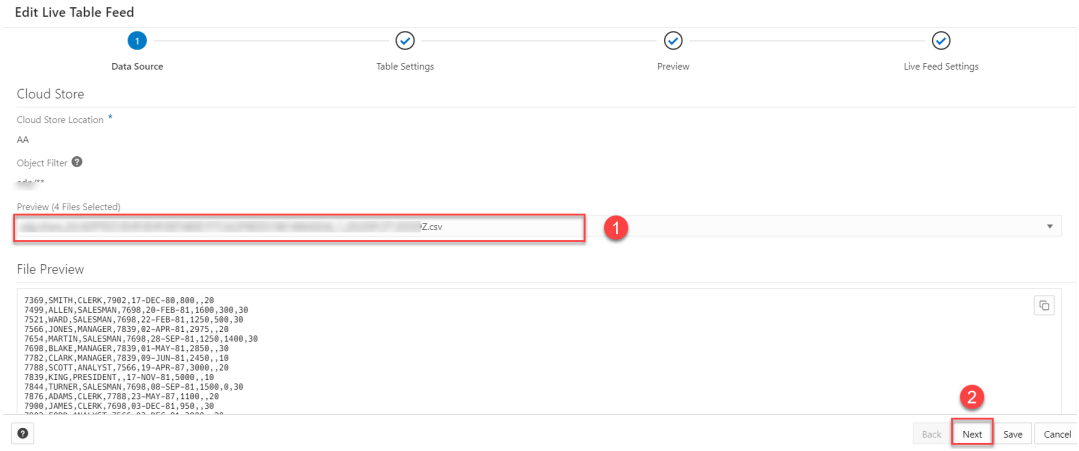
9. Enter the other details for the data load job. For more information on entering the details, refer to the Enter Details for the Data Load Job topic.
10. Once you have added data sources to the data load cart and entered details about the data load job, select **Start** to run the job.
11. After the data load job is run, the data load cart displays a green check mark



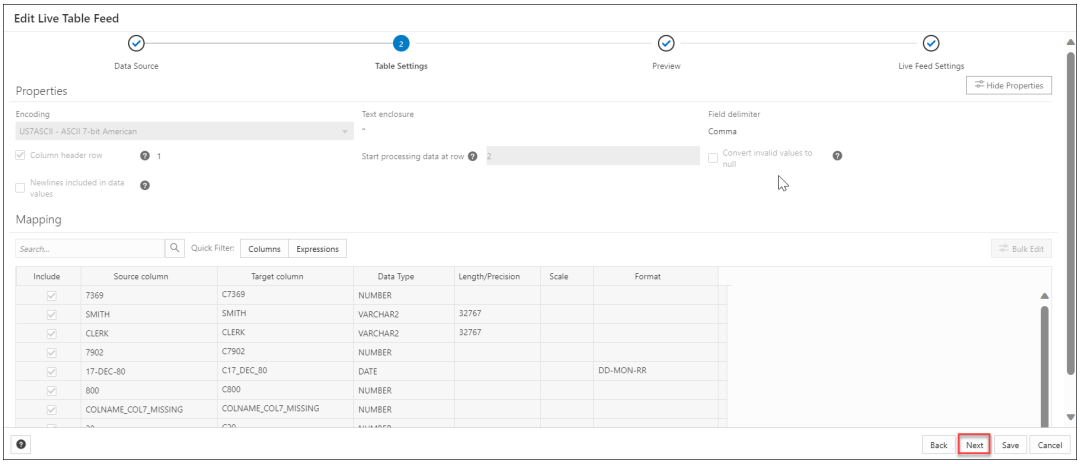
which indicates that the data load job is complete.

12. Click **Create Live Table Feed** on the data load cart to create a Live table feed from the data load job. You will view a successful message that says the creation of Live Table feed is successful and if you wish to edit the live table feed. Click **Yes** to continue and **No** to cancel. Selecting **Yes** opens an Edit Live Table Feed wizard.

13. On the Edit Live Table Feed wizard, you can view the **Cloud Store Location** of the source folder and the **Object Filter** applied to the data.
Select any file whose data you want to preview from the Preview drop-down in the Data Source tab: The field shows the total number of files present in the cloud store folder you loaded. A preview of the data is displayed.



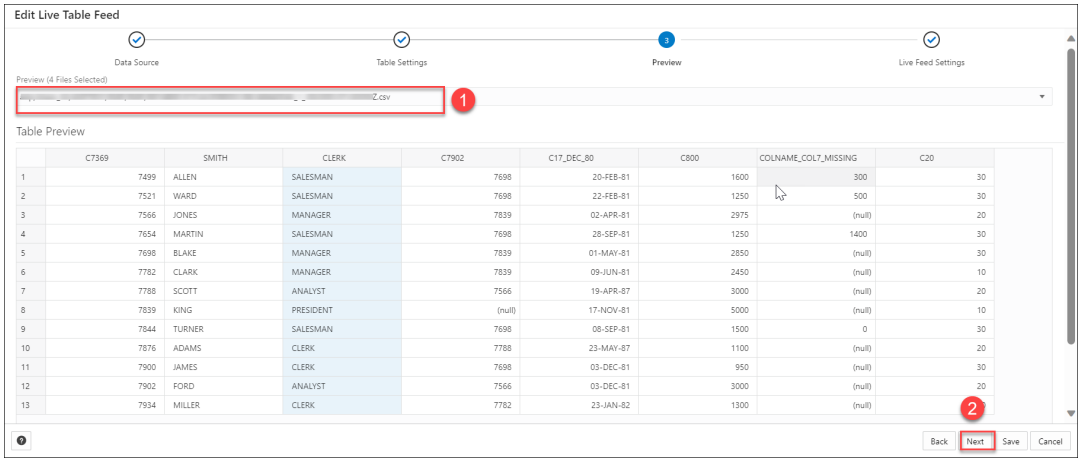
14. Click **Next** to progress to the Table Settings tab.



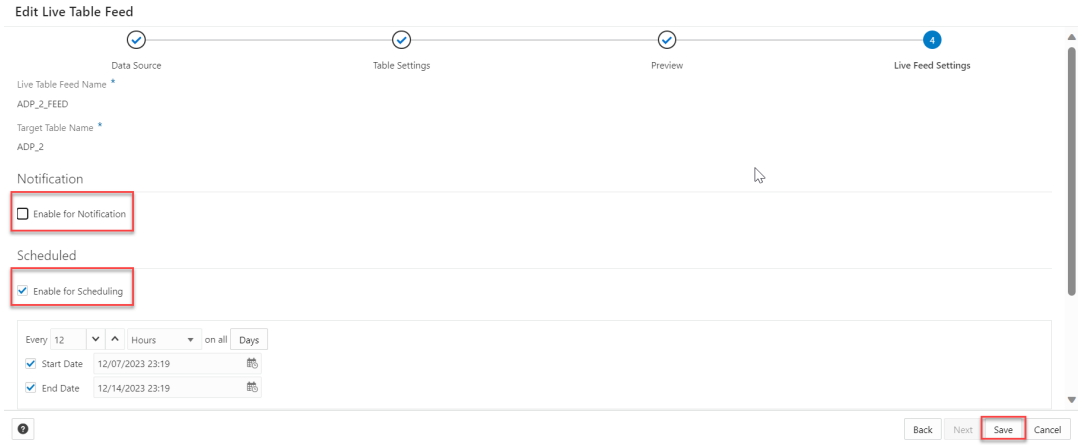
You can view the Properties and Mapping details of the data load job on the Table Settings tab.

> **Note:**
>
> You cannot select or edit any of the details of this tab.

15. Click **Next** to progress to the Preview tab of the wizard.
Select any file from the Preview drop-down to view the file. The Table Preview displays the preview of the file you select from the drop-down.



16. Click **Next** to view the Live Feed Settings tab of the wizard. On the Live Feed Settings tab, specify the following values in the field:



- **Enable for Notification**: Select this option so that new or changed data in the data source will be loaded based on an Oracle Cloud Infrastructure notification. When you select this option, you can avoid delays that might occur when polling is initiated on a schedule (that is, if you selected the live table feed Scheduled option.

  When you select the Enable for Notification option, you must also configure your object store bucket to emit notifications

- **Enable For Scheduling**: Select this option to set up a schedule for data feed. In the time interval fields, enter a number, and select a time type and the days on which to poll the bucket for new or changed files. For example, to poll every two hours on Monday, Wednesday, and Friday, enter **2**, select **Hours**, and then select **Monday**, **Wednesday**, and **Friday** in the appropriate fields. You can select **All Days**, **Monday to Friday**, **Sunday to Thursday**, or **Custom** from the Week Days drop-down. The Custom field enables you to select **Monday**,**Tuesday**, **Wednesday**, **Thursday** and **Friday** .

Select a start and end date. If you don't select a start date, the current time and date will be used as the start date. The end date is optional. However, without an end date, the feed will continue to poll.

The rest of the fields displayed in the wizard such as the **Live Table Feed Name**, **Target Table Name** , and **Consumer Group** are greyed out and disabled for selecting or editing.

Click **Save** to save and create a Live Table Feed from a data load cart.

Loading Data From Local Files

To load data from local files into your Oracle Autonomous Database, on the Data Load page, select **LOAD DATA** and **LOCAL FILE**.

Drag one or more files from your local file system navigator and drop them in the Data Load Cart. You can also click **Select Files** or the Select Files icon, select one or more files from the file system navigator, and then click **Open**.

You can add files in these file formats: AVRO, CSV, JSON, GeoJSON, TSV, delimited TXT, XLS, XLSX, XML, PNG, JPEG and PDF. For information on supported file formats, see Format Specifications for JSON, AVRO, and XML Files.

An item for each file appears in the cart. For an XLS or XLSX spreadsheet, the worksheets of the spreadsheet appear as individual items. The item shows the name of the source file or worksheet and its size, and the name of the table that is the target for the data load. The Data Load tool supports loading tables from only the first worksheet of a multi-worksheet XLSX file when the file is in an object store.

You can add more files to the cart by clicking the Select Files icon. You can add any number of files to the cart and load data from all of them in a single data load job.

To remove a source file from the Data Load Cart, click the Remove (trash can) icon for the source item. To remove all source files from the cart, click the Remove All (trash can) icon in the Data Load Cart menu bar.

To return to the Data Load page, click **Data Load** above the Data Load Cart menu bar.

Specify Processing Options

To specify settings for the data load job or preview the data in the source or the target, select the Settings (pencil) icon for the item in the Data Load Cart.

In the settings pane, on the **Settings** tab, you can view the name and size of the file in the title of the Load Data dialog box.

The **Name** field specifies the name of the target table. The value in the field varies depending on the selection in the **Options** field. If the option is **Create Table**, then the default target value is the name of source file or worksheet.

To specify a different name for the target table, enter it in the **Name** field. For the other target table choices in the **Options** field, the default value is <None>. Expand the drop-down list and select an existing table as the target.

In the **Options** field select **Create Table**, **Insert into Table**, **Replace Data**, **Drop Table and Create New Table**, or **Merge into Table**. Point to the question mark icon to see a brief description of the selected option.

Select a different schema from the **Schema** drop-down to create your target table in another schema.

> **Note:**
>
> The **Schema** drop-down is available only if you have `PDB_DBA` role grant to you. To grant yourself a `PDB_DBA` role, you must log into your Database Actions instance and enter the following command in the **SQL** worksheet area displayed in the **SQL** tab under **Development** tools present in the Launchpad.
>
> ```
> Grant PDB_DBA to Username;
> ```
>
> This drop-down is available for **Create Table** and **Drop Table and Create New Table** options.

The **Source column name** option specifies whether to get the source and target column names from the file or to specify the column names manually. Getting the column names from the header of the source file is the default. If you select the **Column header** option, then the first row in the file is processed as column names. If you deselect the option, then the first row is processed as data. To specify column names manually, enter a name for each target column in the **Mapping** section. You can also select a data type for the column.

The **Start processing data at row** field specifies how many rows to skip when loading the source data into the target. If you have selected the **Column header** option, and if you enter a number greater than 0 in the **Start processing data at row** field, then that number of rows after the first row are not loaded into the target. If you have deselected the **Column header** option, and if you enter a number greater than 0 in the **Start processing data at row** field, that number of rows including the first row are not loaded into the target.

To change the character set encoding for the contents of the file, select a value from the **Encoding** drop-down list.

To specify the characters that enclose text, select the double-quotes or single-quote character or `None` from the **Text enclosure** drop-down list.

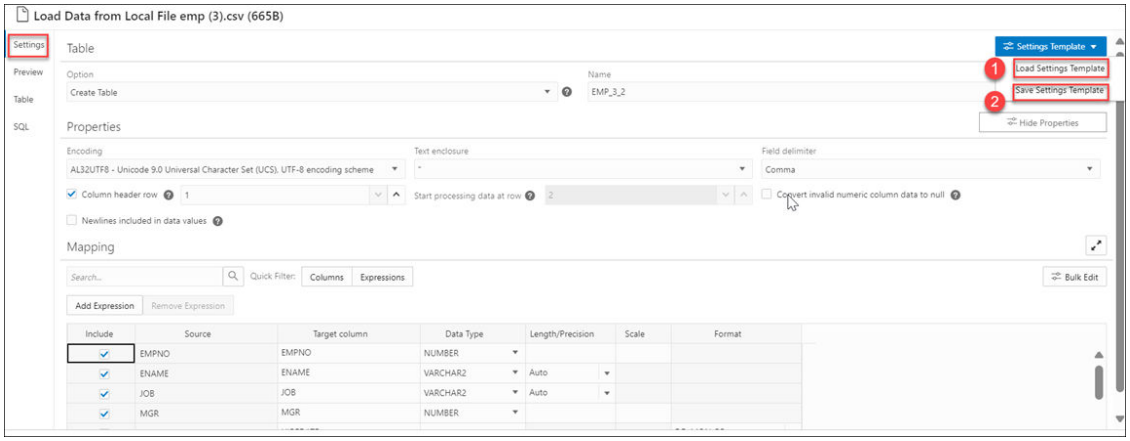
To change the delimiter character that separates columns in the source, expand the **Field delimiter** drop-down list and select a character. For example, if the file has columns delimited by semicolons, change the delimiter from the default comma delimiter to a semicolon.

To convert any invalid value in a numeric source column to a null value in the target column, select the Numeric column **Convert invalid data to null** option.

Settings Template

The save settings feature saves the configuration set in the Cart settings in the form of a JSON file. When opening the Settings template, you have the following options:

1. **Load Settings Template**: Loads a settings template from your local system.
2. **Save Settings Template**: Saves the current existing settings template.



You can use the **Load Settings Template** if you want to use an existing customized template present in your local.

1. From the Settings Template in the Settings tab of the Load Data page, select **Load Settings Template**.
2. You will see a Load Settings Template wizard, click the Settings Template JSON to load a JSON file from your system.
3. Clicking the Settings template JSON will open your local system. Click **OK** to load the JSON file.
4. After you load the JSON file, you can view the updates applied automatically to the settings tab which matches the JSON settings template you load from your local.

You can use the **Save Settings Template** to save the existing current Settings template.

1. From the Settings Template in the Settings tab of the Load Data page, select **Save Settings Template**.
2. The Template file editor appears. Click the Template File name and name the new template.
3. Click **OK** to finish saving the new name of the existing template.
4. You can test the configuration of the new template.

Bulk Edit Settings

You can use the Bulk edit settings to update all the columns at once from the mapping table. Use it to apply changes to the selection currently displayed in the results pane. You can search for the values of the column you want to edit in the search field and click the magnifier icon. The mapping table will display the results of the search. Select the Bulk Edit setting to update the column. The Bulk Edit setting allows you to:

- Update values of all the fields in a group.
- Find and replace, Add Prefix and Add suffix to target column name.
- Include the column(s) for loading data to the target table.
- Exclude the column(s) for loading data to the target table.

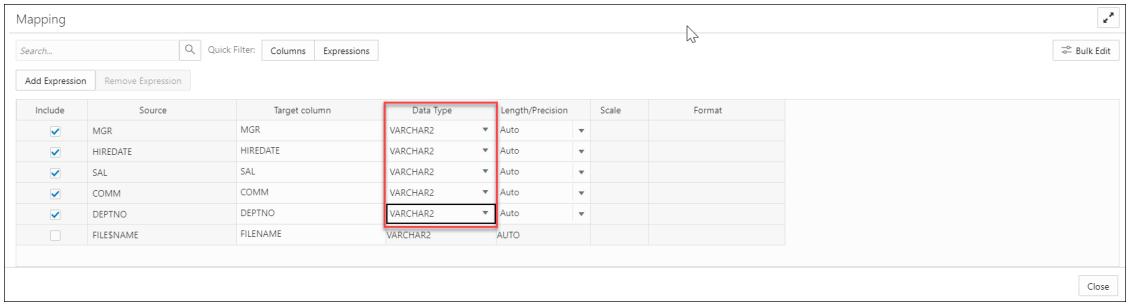
Searching the Column

The Bulk Edit setting updates the columns returned by the search field. The search box besides the Bulk Edit setting icon filters the list of columns you wish to update in a bulk. As soon as you start typing in the search field, the tool returns the field values which contains the letters you type. You can remove the filter by deleting all the content from the search box and clicking the magnifier icon that appears next to the search box.

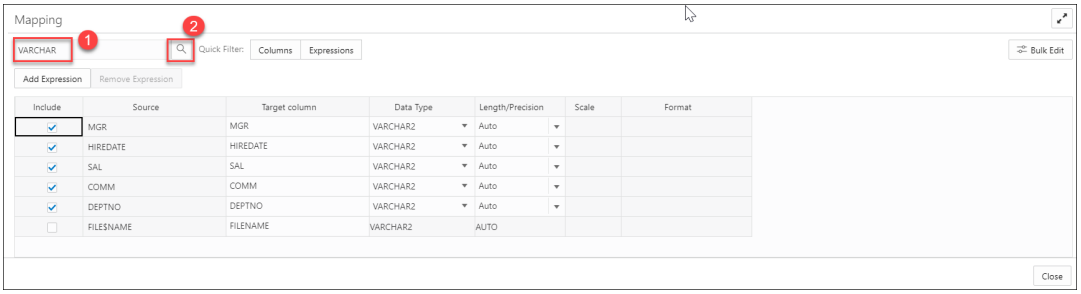
The Bulk Edit setting enables you to update the values of the following columns for all the searches returned by the search field:

- Data Type
- Target Column name
- Include Columns for loading
- Exclude columns for loading

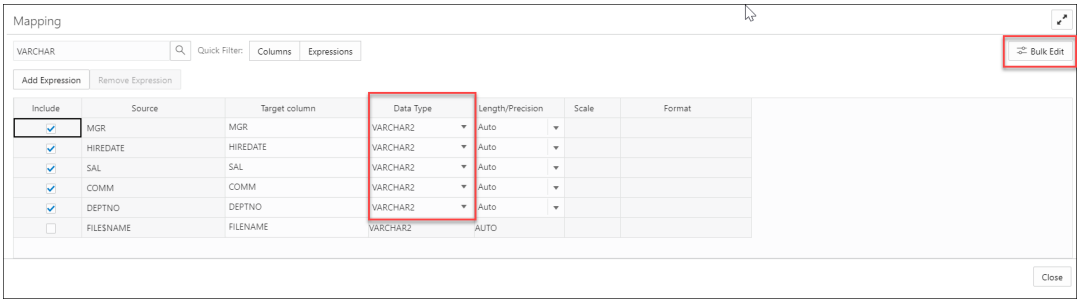
Consider changing the Data Type of first to fifth rows from VARCHAR to NUMBER in the mapping table.



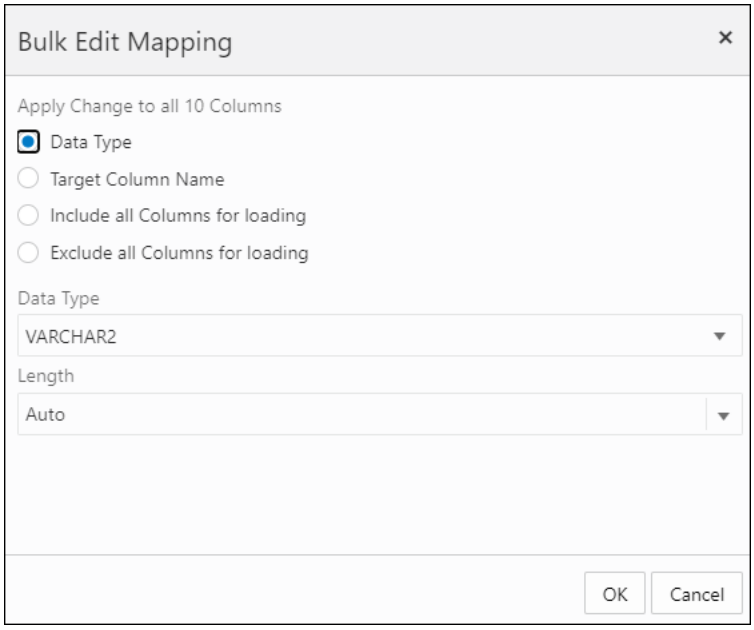
- Type `VARCHAR` in the search field and click on the magnifier icon besides the search field.



- The search will return the rows having `VARCHAR` as the Data Type. Select the **Bulk Edit** icon besides the magnifier icon.



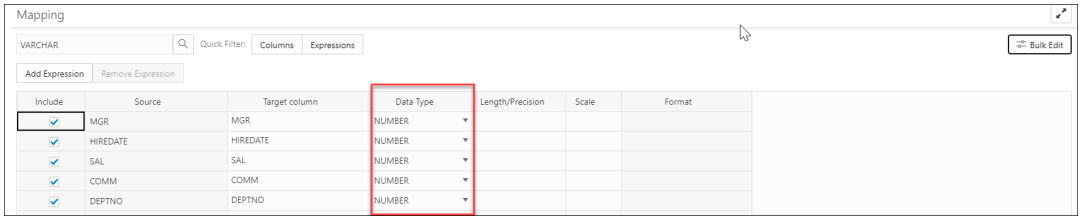
- Selecting the Bulk Edit icon opens the Bulk Edit Mapping dialog box.



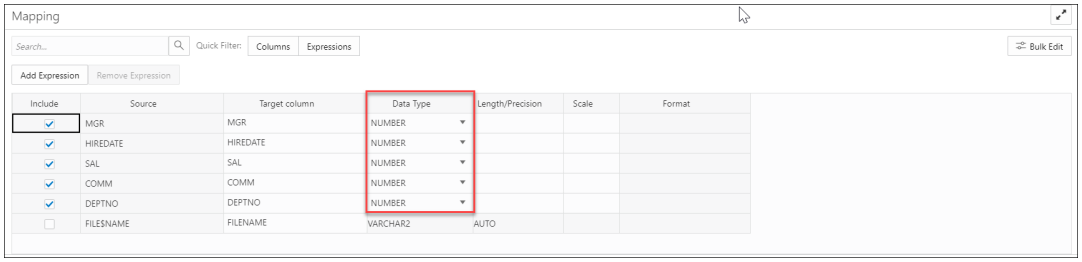
- Click **Data Type** and select `NUMBER` from the Data Type drop-down. Click **OK**.



- The Data Type of the selected rows changes from `VARCHAR` to `NUMBER`.



- Clear the content from the search field and click the magnifier icon to view the bulk edit updates in the Mapping table.

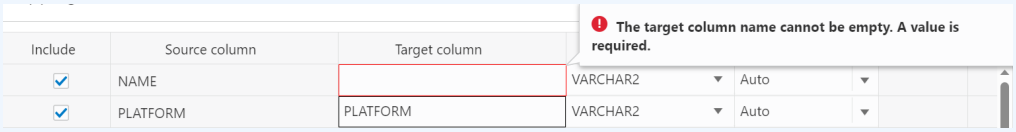


Specify Mappings

You can filter the results in the mapping table with Quick Filter field that enables you to filter out **Columns** or **Expressions**.

Select **Add Expression** to perform **Sentiment Analysis** or **Key Phrase extraction** with the source data. See Use OCI Language Service Capabilities in Data Studio for more details.

If you select the **Create Table** or **Drop Table and Create New Table** option and you are getting the source column names from the file header, then in the **Mapping** section either accept the default values for the target columns and data types or specify different values. To specify different values, in the target column, enter a name for the column. In the Data Type column, select a data type from the drop-down list. If you are not getting the source column names from the file header, then for each source column specify a name for the target column and select a data type from the Data Type drop-down list. For the Date data type, select a date format from the Format drop-down list.

> **Note:**
>
> You will receive a tooltip error message with the exact reason for the error message when you complete editing a cell. The mapping grid cell will be highlighted with red to indicate an invalid value that must be fixed. The highlight is removed after you fix the invalid value. For example, you can view the following tooltip error message when the target column name is not filled in.
>
> 
>

For the **Merge into Table** option, for each source column, select a target column form the drop-down list. You must specify at least one column as a key column. To specify a column as a key column, select the **Merge Key** check box for the column. Merge keys are one or more columns that uniquely identify each row in the table. Merge keys must not contain any null values. For loading tables with primary keys, this option automatically enables the selection of primary key columns as the merge keys.

For the **Insert into Table** or **Replace Data** options, for each source column, select a target column from the drop-down list of existing columns.

Preview Source Data

To view a selection of data in the source file, select the **Preview** tab. The Preview tab displays a File menu which displays the data in tabular format with it's values. You can copy the table. The **Load Preview** tab displays the source data.

Any modifications you make in the source preview do not affect the loading of data from the file.

Preview Target Data

For all options except **Create Table**, to view the existing data in the target table, select the **Load Preview** tab. The load preview displays the data in the target table before you run the data load job.

To close the settings pane, click **Close**.

Run the Data Load Job

When you have added all of the sources for the job and specified the settings for each source, to run the job click the Start icon in the Data Load Cart menu bar. In the Run Data Load Job dialog box, click **Start**. To stop the data load job, click the Stop icon.

Once the data load job starts, you can view the progress of the job in the Data Load dashboard. See The Data Load Page for more details.

When the data load job completes, the Data Load dashboard page displays the results of the job. At the top of the header of the table load, you can view the name of the table along with the total columns present in the table.

Click **Report** to view the total number of rows processed successfully and the count of rejected rows. You can also view the Start time. The SQL pane of the Report displays the equivalent SQL code of the job.

To view information about an item in the job, click the **Actions** icon on the Table Load.

To view a log of the load operation, click the Logging icon. You can save the log, clear it, or refresh it. Click **OK** to dismiss the log.

The list of tables on the Data Load page contains any new tables created. The target tables for the **Insert into Table**, **Replace Data**, **Drop Table and Create New Table**, and **Merge into Table** options contain the loaded data.

Fixing a data load job

After your data load job, you might see errors that you want to correct, or upon inspection, realize that you wanted to name a column differently. In such cases, you will view a warning sign on the Table load. Click the **Reload** icon to reload source with suggested fixes. Click **Actions** icon on the Table load and select **Edit** to make any changes to the data load job (i.e., change a column name).

Click **Apply** to apply any changes. Click **Close** to return to the Database Actions page.

Loading Data from Other Databases

To load data from tables in another database into your Oracle Autonomous Database, on the Data Load page, select **LOAD DATA** and **DATABASE**. Select a database link from the drop-down list. Drag one or more tables from the list of database tables and drop them in the Data Load Cart.

Each table appears as an item in the Data Load Cart. The item shows the name of the table and the number of rows in it, and the name of the table that is the target for the data load.

To remove a table from the Data Load Cart, click the Remove (trash can) icon for the item. To remove all tables from the cart, click the Remove All (trash can) icon in the Data Load Cart menu bar.

To add a remote database to the list of database links, create a database link to the remote database. For information on creating a database link, see Database Links in *Oracle Database Administrator's Guide*.

The databases available to you appear in the drop-down list of the database navigation pane of the Load Tables page.

You can filter the tables displayed in the navigation pane by entering a case-sensitive value in the search field at the top of the navigation tree and pressing Enter. To display all of the tables again, clear the search field and press Enter.

You can add any number of tables from the navigation pane to the Data Load Cart and load data from all of them in a single data loading job. You can set filters on the data for a table to load only the specified data.

Specify Processing Options

To specify settings for the data load job, preview the data in the source or the target, and see statistics about the data, click the Settings (pencil) icon for the item in the Data Load Cart.

In the settings pane, on the Settings tab, you can view the name and size of the file in the title of the Load Data dialog box.

The **Table** field specifies the name of the target table. The value in the field varies depending on the selection in the **Options** field. If the option is **Create Table**, then the default target value is the name of source table. To specify a different name for the target, enter it in the **Name** field. For the other options, the default value is <None>. Expand the drop-down list and select a table as the target.

In the **Options** field for the source, select **Create Table**, **Insert into Table**, **Replace Data**, **Drop Table and Create New Table**, or **Merge into Table**. Point to the question mark icon to see a brief description of the selected option.

Select a different schema from the **Schema** drop-down to create your target table in another schema.

> **Note:**
>
> The **Schema** drop-down is available only if you have `PDB_DBA` role grant to you.
>
> To grant yourself a `PDB_DBA` role, you must log into your Database Actions instance and enter the following command in the **SQL** worksheet area displayed in the **SQL** tab under **Development** tools present in the Launchpad.
>
> ```
> Grant PDB_DBA to Username;
> ```
>
> This drop-down is available for **Create Table** and **Drop Table and Create New Table** options.

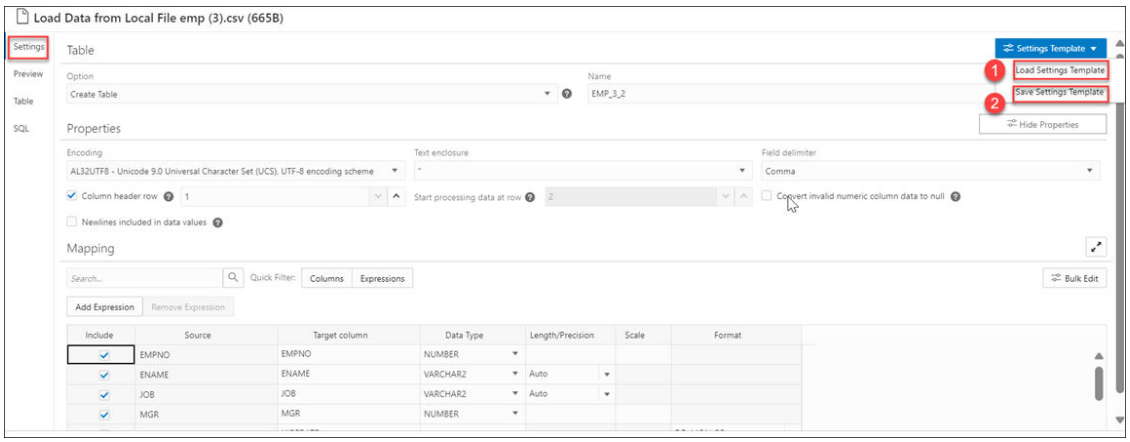
If you select **Create Table**, then in the **Name** field accept the default name, which is the name of the source table, or enter a different name.

If you select one of the other options, then expand the drop-down list of the **Name** field and select a table as the target.

Settings Template

The save settings feature saves the configuration set in the Cart settings in the form of a JSON file. When opening the Settings template, you have the following options:

1. **Load Settings Template**: Loads a settings template from your local system.
2. **Save Settings Template**: Saves the current existing settings template.



You can use the **Load Settings Template** if you want to use an existing customized template present in your local.

1. From the Settings Template in the Settings tab of the Load Data page, select **Load Settings Template**.
2. You will see a Load Settings Template wizard, click the Settings Template JSON to load a JSON file from your system.
3. Clicking the Settings template JSON will open your local system. Click **OK** to load the JSON file.
4. After you load the JSON file, you can view the updates applied automatically to the settings tab which matches the JSON settings template you load from your local.

You can use the **Save Settings Template** to save the existing current Settings template.

1. From the Settings Template in the Settings tab of the Load Data page, select **Save Settings Template**.
2. The Template file editor appears. Click the Template File name and name the new template.
3. Click **OK** to finish saving the new name of the existing template.
4. You can test the configuration of the new template.

Bulk Edit Settings

You can use the Bulk edit settings to update all the columns at once from the mapping table. Use it to apply changes to the selection currently displayed in the results pane. You can search for the values of the column you want to edit in the search field and click the magnifier icon. The mapping table will display the results of the search. Select the Bulk Edit setting to update the column. The Bulk Edit setting allows you to:

- Update values of all the fields in a group.
- Find and replace, Add Prefix and Add suffix to target column name.
- Include the column(s) for loading data to the target table.

- Exclude the column(s) for loading data to the target table.

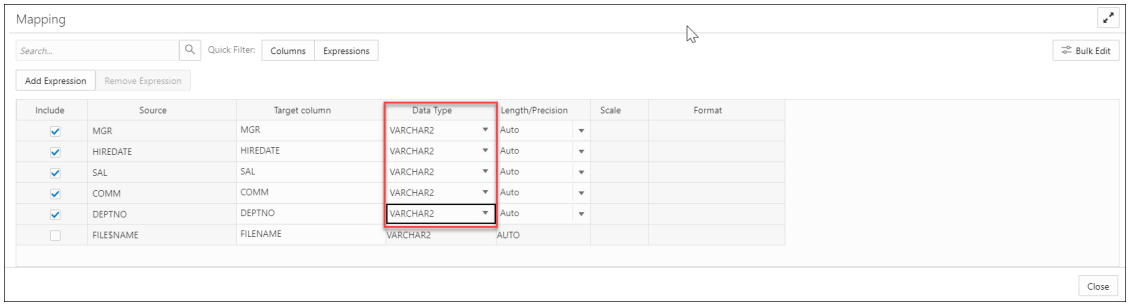
Searching the Column

The Bulk Edit setting updates the columns returned by the search field. The search box besides the Bulk Edit setting icon filters the list of columns you wish to update in a bulk. As soon as you start typing in the search field, the tool returns the field values which contains the letters you type. You can remove the filter by deleting all the content from the search box and clicking the magnifier icon that appears next to the search box.

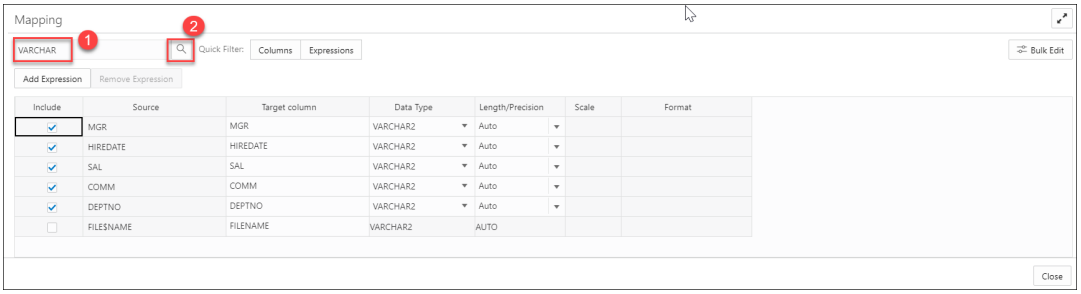
The Bulk Edit setting enables you to update the values of the following columns for all the searches returned by the search field:

- Data Type
- Target Column name
- Include Columns for loading
- Exclude columns for loading

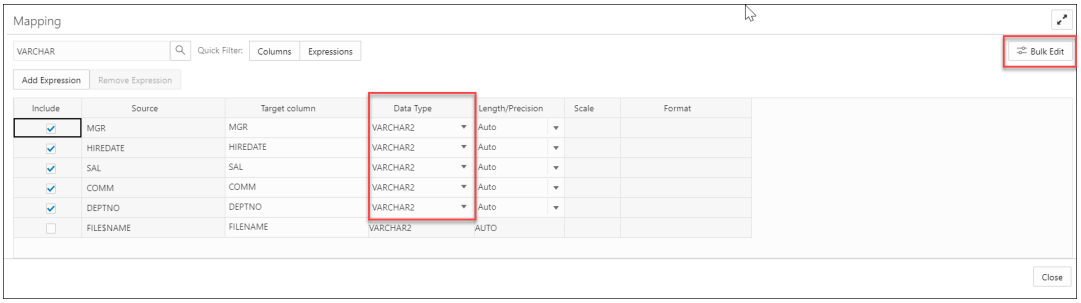
Consider changing the Data Type of first five rows from VARCHAR to NUMBER in the mapping table.



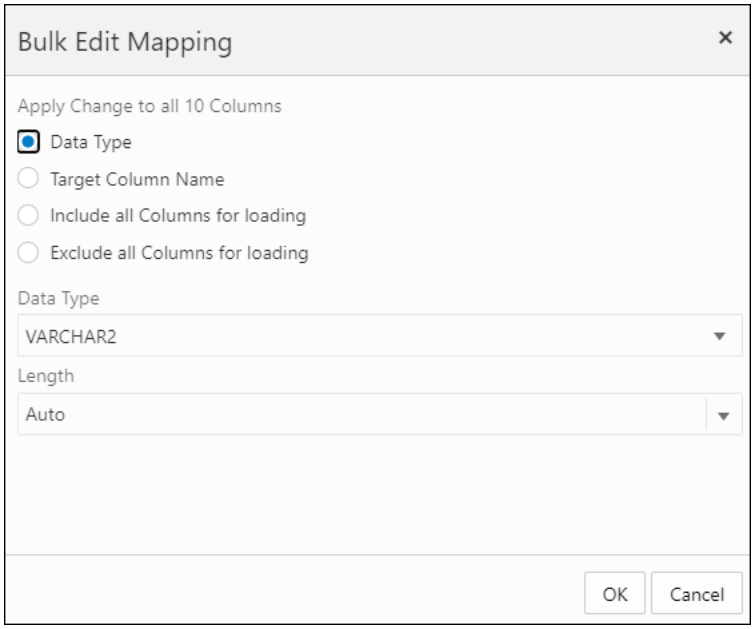
- Type `VARCHAR` in the search field and click on the magnifier icon besides the search field.



- The search will return the rows having `VARCHAR` as the Data Type. Select the **Bulk Edit** icon besides the magnifier icon.



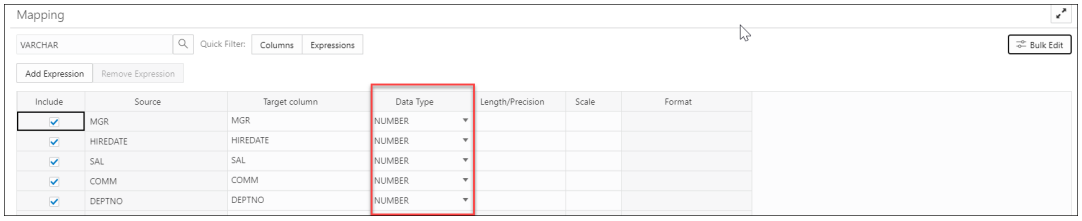
- Selecting the Bulk Edit icon opens the Bulk Edit Mapping dialog box.



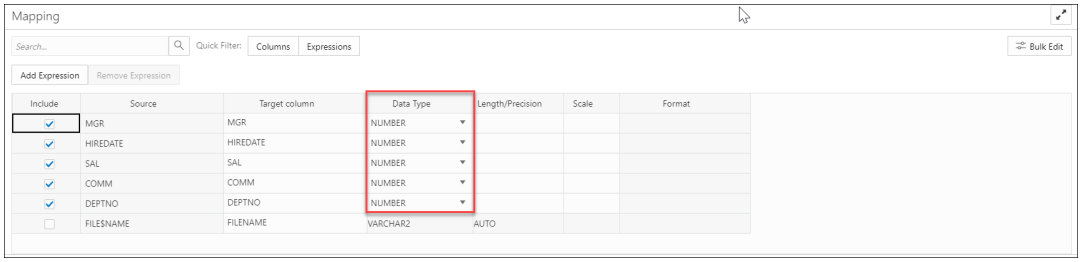
- Click **Data Type** and select `NUMBER` from the Data Type drop-down. Click **OK**.



- The Data Type of the selected rows changes from `VARCHAR` to `NUMBER`.



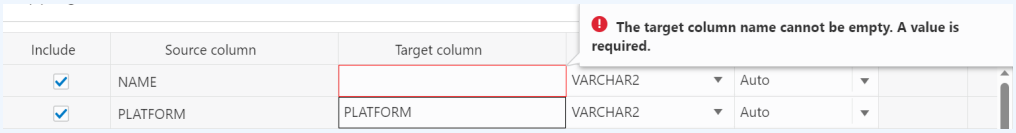
- Clear the content from the search field and click the magnifier icon to view the bulk edit updates in the Mapping table.



Specify Mappings

If you select the **Create Table** or the **Drop Table and Create New Table** option, then in the **Mapping** section either accept the default values for the target columns or specify different values. For the target column, enter a name for the column.

> **Note:**
>
> You will receive a tooltip error message with the exact reason for the error message when you complete editing a cell. The mapping grid cell will be highlighted with red to indicate an invalid value that must be fixed. The highlight is removed after you fix the invalid value. For example, you can view the following tooltip error message when the target column name is not filled in.



For the **Insert into Table** or **Replace Data** options, select a target column from the drop-down list of existing columns.

For the **Merge into Table** option, for each source column, select a target column form the drop-down list. You must specify at least one column as a key column. To specify a column as a key column, select the **Merge Key** check box for the column. Merge keys are one or more columns that uniquely identify each row in the table. Merge keys must not contain any null values. For loading tables with primary keys, this option automatically enables the selection of primary key columns as the merge keys.

Preview

To view the data in the source table, in the settings pane select the **Source Table** tab. The source preview displays the data in the table.

Table

For all options except **Create Table**, to view the existing data in the target table, in the settings pane select the **Target Table** tab. The target preview displays the data in the table before you run the data load job.

SQL

The SQL tab displays the SQL commands that will be run to complete this data load job.

> **Note:**
>
> You can see the SQL code even before the table is created.

Run the Data Load Job

When you have added all of the source tables for the job and specified the settings for each table, to run the job click the Start icon in the Data Load Cart menu bar. To stop the data load job, click the Stop icon.

Once the data load job starts, you can view the progress of the job in the Data Load dashboard.

When the data load job completes, the Data Load dashboard page displays the results of the job. At the top of the header of the table load, you can view the name of the table along with the total columns present in the table.

Click **Report** to view a report of the total rows loaded and failed for a specific table.You can view the name of the table, the time the table was loaded and the time taken to process the load.

At the header of the Table Load, you can view the name of the table with total number of columns loaded. When you expand the Table Load you can view the options you used to load the source data with the count of rows loaded.

To view information about an item in the job, click the Settings icon in the item. The settings pane has the same **Settings**, **Source**, **Table**, **SQL**, **Job Report** and **Data Definition** tabs as the settings pane before running the job, except that the target preview now contains the data loaded by the data load job. To close the settings pane, click **Close**.

To view a log of the load operation, click the Log icon. You can save the log, clear it, or refresh it. Click **OK** to dismiss the log.

The list of tables on the Data Load / Explore page contains any new tables created. The target tables for the **Insert into Table**, **Replace Data**, **Drop Table and Create New Table**, and **Merge into Table** options contain the loaded data.

Fixing a data load job

After your data load job, you might see errors that you want to correct, or upon inspection, realize that you wanted to name a column differently. In such cases, you will view a warning sign on the Table load. Click the **Reload** icon to reload source with suggested fixes. Click **Actions** icon on the Table load and select **Edit** to make any changes to the data load job (i.e., change a column name).

Click **Apply** to apply any changes. Click **Close** to return to the Database Actions page.

- Load Data from Oracle and Non-Oracle Databases using Database Links
  This section of the document describes how to create database links and then transfer data from various non-Oracle databases into your Autonomous Database using Data Studio.

Load Data from Oracle and Non-Oracle Databases using Database Links

This section of the document describes how to create database links and then transfer data from various non-Oracle databases into your Autonomous Database using Data Studio.

To load data from Oracle and Non-Oracle Databases using Database Links, you must first create a Database Link.

Run the following steps to load data from Oracle databases, or non-Oracle databases into Autonomous Database using Database Links:

1. Log in to your Database Actions instance. On the **Development** menu, click **SQL**. You can view the SQL Worksheet.
2. **Create Credentials and Database Links to access the Autonomous Database.**

   You will use the `DBMS_CLOUD.CREATE_CREDENTIAL` procedure to create credentials where the `username` and `password` values you specify are the credentials for the target database.

   The `DBMS_CLOUD.CREATE_CREDENTIAL` procedure stores the credentials in an encrypted format.

   After you create credentials, you will create database links to load data in Data Studio from non-Oracle Databases using the Data Load tool.

   Following are the sample codes you will run to create credentials and then use the same credentials to create the database links:

   a. **Create Database Links from MySQL**:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'MYSQLPE_CRED',
    username => 'admin',
    password => <enter password here> );

  DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
    db_link_name => 'MYSQLPE_DBLINK',
    hostname => 'scottmysql.scottprivate.com',
    port => '3306',
    service_name => 'scott',
    ssl_server_cert_dn => NULL,
    credential_name => 'MYSQLPE_CRED',
    private_target => TRUE,
    gateway_params => JSON_OBJECT('db_type' value 'MYSQL'));
END;
/
```

   b. **Create Database Links from Microsoft Azure**:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'AZURE_CRED',
    username => 'oracle',
    password => <enter password here> );

  DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
    db_link_name => 'AZURE_DBLINK',
    hostname => 'scott-server.database.windows.net',
    port => '1433',
    service_name => 'scottdb',
    ssl_server_cert_dn => NULL,
    credential_name => 'AZURE_CRED',
    gateway_params => JSON_OBJECT('db_type' value 'AZURE'));
```

```
END;
/
```

c. Create Database Links from Snowflake:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'SNOWFLAKE_CRED',
    username => 'SCOTT',
    password => <enter password here> );

  DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
    db_link_name => 'SNOWFLAKE_DBLINK',
    hostname => 'https://abcdefg-hij12345.snowflakecomputing.com',
    port => '443',
    service_name => 'SCOTTDB',
    credential_name => 'SNOWFLAKE_CRED',
    gateway_params => JSON_OBJECT('db_type' value 'SNOWFLAKE'),
    directory_name => NULL,
    ssl_server_cert_dn => NULL);
END;
/
```

d. Create Database Links from Oracle Cloud Object Storage using Native OCI Credentials

```
BEGIN
  DBMS_CLOUD.GET_OBJECT(
    object_uri => 'https://objectstorage.us-ashburn-1.oraclecloud.com/n/
abcdefgh12ij/b/scott-bucket/o/cwallet.sso',
    credential_name => 'OCI_NATIVE_CRED',
    directory_name => 'REG_WALLET_DIR');

  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'ADBS_CRED',
    username => 'ADMIN',
    password => <enter password here> );

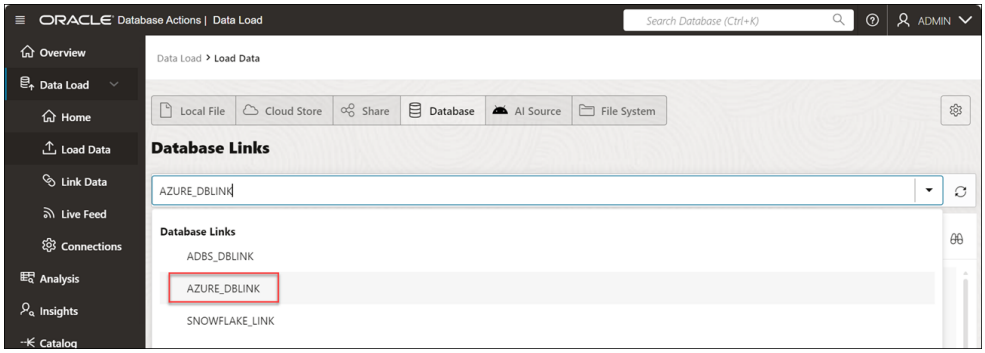
  DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
    db_link_name => 'ADBS_DBLINK',
    hostname => 'adb.us-ashburn-1.oraclecloud.com',
    port => '1522',
    service_name =>
'ab1cdefgh2i3jkl_scottadbs_high.adb.oraclecloud.com',
    credential_name => 'ADBS_CRED',
    directory_name => 'REG_WALLET_DIR');
END;
/
```

For more details on the parameters and their values, see the Create_Database_Link procedure.

> **Note:**
>
> - You require `ADMIN` privileges to run the `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` procedure.
> - You can use a vault secret credential for the target database credential in a database link. See Use Vault Secret Credentials for more information.

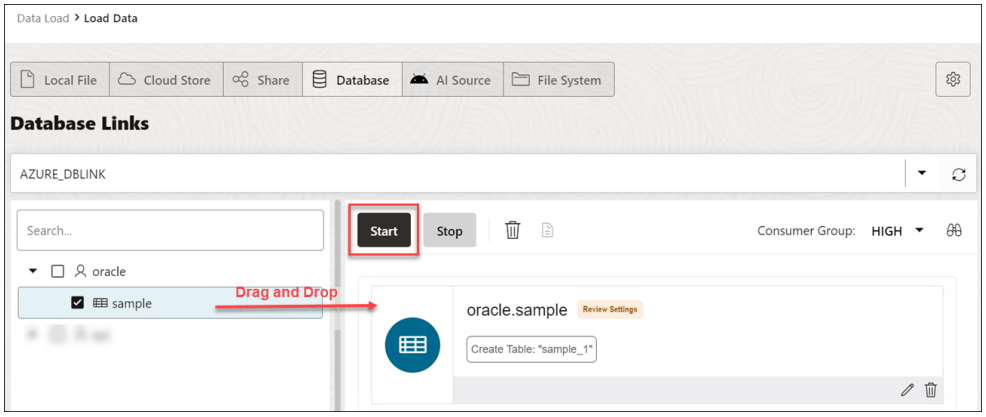
3. a. **Use the Data Load tool to LOAD data from Oracle and non-Oracle Databases:** You can use the database link you created to load access data on Data Studio. After you create the database links, you can view and use the list of database links you created in the Data Load tool and access data from Oracle and non-Oracle databases.

   - Click **Data Load** under the Data Studio suite of tools and select the **Load Data** card.
   - Click the **Database** tab and select the database link from the list of available database links you created. In this example, we will use `AZURE_DBLINK`.

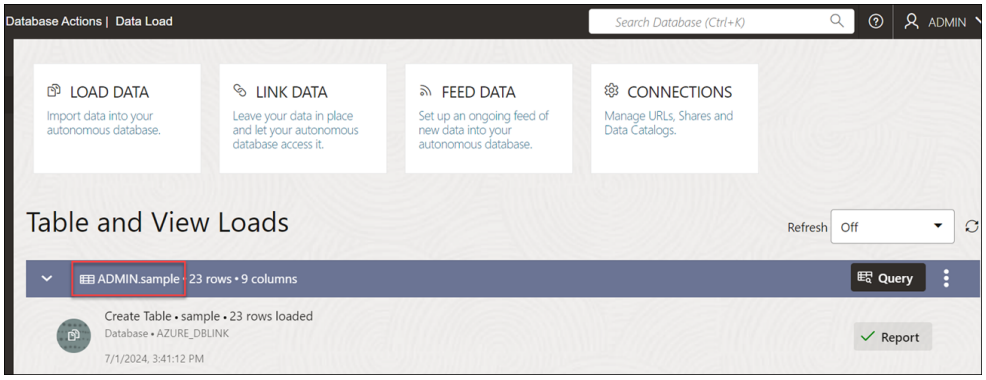


> **Note:**
>
> You cannot view the database link you create from **MySQL** since it resides within a Virtual Cloud Network (VCN), and the Autonomous Database would need to be either located in or connected to the Private Subnet within that VCN.

   - Drag the `sample` table from the database and drop it in the data load cart to load the data in your autonomous database.

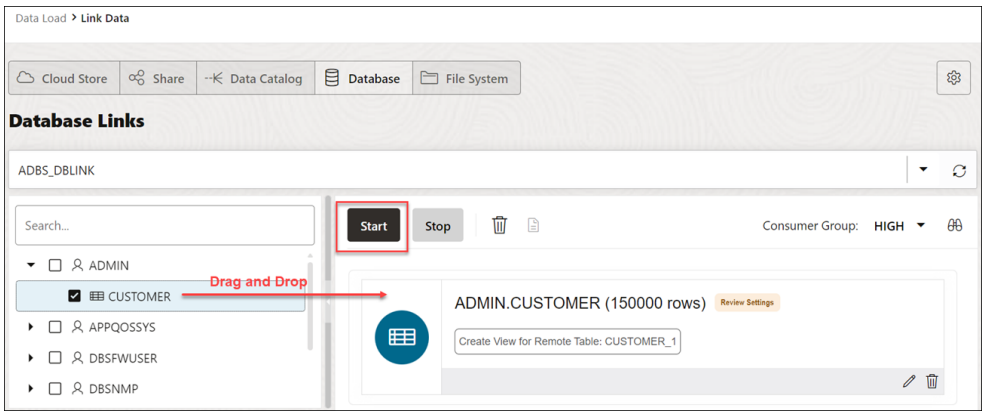


- Click **Start** to start the loading process.
- After the table is loaded in the tool, you can view the result in the Data Load Dashboard under the **Table and View Loads** section.

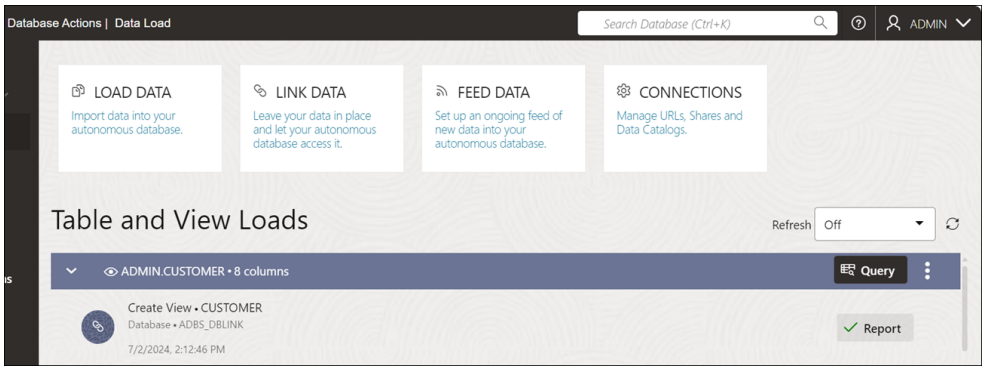


b. Use the Data Load tool to LINK data from Oracle and non-Oracle Databases: Alternatively, you can use the Data Load to link data that creates a view of the data by directly accessing cloud databases via the Database Link.

- From the **Data Load** menu under Data Studio, select **Link Data**.
- Click the **Database** tab. In this example, we will use the `ADBS_DBLINK` link.
- Drag the `customer` table from the database and drop it in the data link cart to link the data in your autonomous database.



- Click **Start** to start linking the data in Data Studio.
- After the data is linked, you will view the results of the data link job in the Data Load dashboard under the **Table and View Loads** section.



You have successfully loaded the data from Oracle and non-Oracle databases using Database Links in Data Studio.

Loading Data from Cloud Storage

You can load data from a cloud store to a table in your Autonomous Database.

You can load files in these file formats: AVRO, CSV, JSON, GeoJSON, ORC, Delimited TXT, XLSX, PRQ, GZ, GNU ZIP and Tab-Separated Values. For information on supported file formats, see Format Specifications for JSON and AVRO Files. The Data Load tool supports loading tables from multiple worksheets XLSX files when the file is in a cloud store.

You can set filters on the data for a table to load only the specified data. For example, to limit the files to only those that are CSV files, enter *.CSV in the file extension filter.

Configure and run a data load job from the Load Cloud Object page. To open that page:

1. Open the Database Actions and select **Data Load**.
2. Select **LOAD DATA** and **Cloud Store** .

On the left side of the page is a *navigator pane*, where you choose a cloud store connection and the folders or files containing the data. On the right of the page is the data load *cart*, where you stage the files and folders for the data load job. You can set options for the data load job before running it. The Autonomous Database comes with predefined CPU/IO shares assigned to different consumer groups. You can set the consumer group to either low, medium or high while executing a data load job depending on your workload.

- Manage Cloud Storage Links for Data Load Jobs
  Before you can load data from a cloud store, you must establish a connection to the cloud store you want to use. You can select cloud store location from the cloud store locations field.
- Prepare the Data Load Job
- Add Files or Folders for the Data Load Job
- Enter Details for the Data Load Job
  Enter the details about the data load job in the Load Data from Cloud Storage pane.

- Run the Data Load Job
  Once you've added data sources to the data load cart and entered details about the data load job, you can run the job.
- View Details About the Data Load Job After It Is Run
- View the Table Resulting from the Data Load Job

Manage Cloud Storage Links for Data Load Jobs

Before you can load data from a cloud store, you must establish a connection to the cloud store you want to use. You can select cloud store location from the cloud store locations field.

On the Load Data page when you select Cloud Store:

1. Click the **Create Cloud Store locations** menu besides the cloud store locations text field. This opens an Add Cloud Store Location dialog box. See Managing Connections to add cloud store location.

See Managing Connections.

To return to the Load Cloud Object page, click **Data Load** in the breadcrumbs at the top of the page and then navigate back to the page.

Prepare the Data Load Job

As you'll see below in Enter Details for the Data Load Job, the first decision you'll make when configuring a data load job is how to load the source data into a new or existing table in the database. The choices are:

- Create a table and insert data loaded from the source into the new table.
- Insert data loaded from the source into an existing table.
- Delete all data in an existing table and then insert new data from the source into the table.
- Drop a table, create a new table, and then insert data loaded from the source into the new table.
- Merge data from the source into a table by updating existing rows in the table and inserting new rows into the table.

You may have to adjust your source data or your target table so that the source data loads correctly into the external target table. The number, order, and data types of columns in the source must match those in the target. Consider:

- If you're creating a new table or if the columns in your source exactly match the columns in an existing target, you don't have to do any special preparation.
- If the columns in your source don't match the columns in an existing target, you must edit your source files or target table so they do match.
- If you're loading multiple files, you must make sure that:
  - All the source files are of the same type, for example, CSV, JSON, etc.
  - The number, order, and data types of the columns in all the source files match (and that they match the target, if you're loading into an existing table).
- If you want to partition by date:
  - The source file must contain data where the data type is date or timestamp.
  - You must load a folder containing two or more data sources.
  - The names of the files in the folder must indicate a date or dates, for example, `MAR-1999.csv` or `2017-04-21.xlsx`.

Add Files or Folders for the Data Load Job

Add files from the cloud store to the data load cart, where you can edit the details of the data load job. To add the files:

1. From the list at the top of the navigator pane on the left, select the bucket from the drop-down with your source data.

   The list shows links that were established on the Manage Cloud Storage page. If you haven't yet registered the cloud store you want to use, click the **Create Cloud Store Location** button at the top of the page and register a connection. See Manage Cloud Storage Links for Data Load Jobs, above.

2. Drag one or more items from the file navigator on the left and drop them into the cart on the right.

   - You can add files, folders, or both. A card is added to the cart for each file or folder you drag into it. The card lists the name of the source file or folder and a proposed name for the target table.
   - If you add a folder that contains multiple files, all the files must be of the same type, that is, CSV, TXT, etc.

     When you add the folder to the cart, a prompt is displayed that asks if you want to load all the objects from the multiple source files into a single target table. Click **Yes** to continue or **No** to cancel.
   - When you add multiple individual files or multiple folders to the cart, the data represented by each card will be loaded into a separate table, but all the items in the cart will be processed as part of the same data load job. The multiple files you load must have same file extension.
   - You can add files or folders from a different bucket, but if you do that, you're prompted to remove all files that are already in the cart before proceeding. To select files from a different bucket, select the bucket from the drop-down list in the navigator pane on the left and then add the file(s), as described above.
   - You can drop files or folders into the data load cart and then navigate away from the Data Load Object page. When you return to the page, those items remain on the page.

You can remove items from the cart before running the data load job:

- To remove an item from the cart, click the **Remove** icon on the card for the item.
- To remove all items from the cart, click **Remove All** in the data link cart menu bar at the top of the pane besides the **Start** and **Stop** icon.

Enter Details for the Data Load Job

Enter the details about the data load job in the Load Data from Cloud Storage pane.

On the card in the data link cart, click the **Actions** icon and select the **Settings** to open the Load Data from Cloud Storage pane for that job. The pane contains:

- Settings Tab - Table Section
- Settings Tab - Properties Section
- Settings Tab - Mapping Section
- Preview Tab
- Table Tab
- Close Button - Save and Close the Pane

Settings Tab - Table Section

Set details about the target table in the **Table** section.

- **Option:** Select an item from the **Option** list to specify how the data should be loaded into a new or existing table. The processing options are:
  - **Create Table**: Creates a table and inserts the data into the new table. When you select this option, the **Name** field on the **Settings** tab is filled with a default name, based on the name of the source file or folder. You can change it if you want.
  - **Insert into Table**: Inserts data loaded from the source into an existing table. When you select this option, the **Name** field on the **Settings** tab presents a list of the tables in the current schema. Select the table into which you want to insert the data.
  - **Replace Data**: Deletes all data in the existing table and then inserts new data from the source into the table. When you select this option, the **Name** field on the **Settings** tab presents a list of the tables in the current schema. Select the table you want to use.
  - **Drop Table and Create New Table**: Drops the table (if it already exists), creates a new table, and then inserts the new data into the table. When you select this option, the **Name** field on the **Settings** tab presents a list of the tables in the current schema. Select the table you want to use.
  - **Merge into Table**: Updates existing rows and inserts new rows in the table. When you select this option, the **Name** field on the **Settings** tab presents a list of the tables in the current schema. Select the table you want to use.
- Select a different schema from the **Schema** drop-down to create your target table in another schema.

> **Note:**
>
> The **Schema** drop-down is available only if you have `PDB_DBA` role grant to you.
>
> To grant yourself a `PDB_DBA` role, you must log into your Database Actions instance and enter the following command in the **SQL** worksheet area displayed in the **SQL** tab under **Development** tools present in the Launchpad:
>
> ```
> Grant PDB_DBA to USER;
> ```
>
> You can now view the **Schema** drop-down. It is only available for **Create Table** and **Drop Table and Create New Table** options.

- **Name:** The name of the target table.
- **Partition Column:**

  List Partitions and Date-based partitions are the different types of partitions available in data loading.

  List partitioning is required when you specifically want to map rows to partitions based on discrete values.

  To partition according to a specific column, click the **Partition Column** drop-down list and select the column you want to use for the partitioning.

  You will have N files per partition value, all partitioned by the partition column you select.

> **Note:**
>
> – For linked files (from external tables) there is also a requirement that for each file, the list partitioning column can contain only a single distinct value across all of the rows.
>
> – If a file is list partitioned, the partitioning key can only consist of a single column of the table.

Date-based partitioning is available when you load a folder containing two or more data sources that contain date or timestamp data.

To partition according to date, click the **Partition Column** drop-down list and select the **DATE** or **TIMESTAMP** column you want to use for the partitioning.

Settings Tab - Properties Section

Specify options to control how the source data is interpreted, previewed, and processed. These options vary, depending on the type of source data.

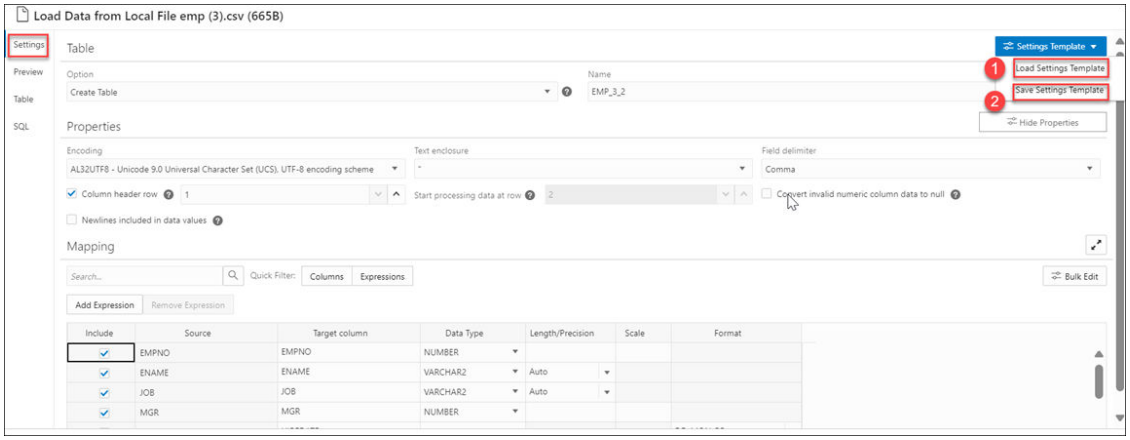
- **Encoding:** Select a character encoding type from the list. This option is available when the loaded file is in plain text format (CSV, TSV, or TXT). The default encoding type is UTF-8.
- **Text enclosure:** Select the character for enclosing text: **"** (double-quote character), **'** (single-quote character) or **None**. This option is visible only when the selected file is in plain text format (CSV, TSV, or TXT).
- **Field delimiter:** Select the delimiter character used to separate columns in the source. For example, if the source file uses semicolons to delimit the columns, select **Semicolon** from this list. The default is **Comma**. This option is visible only when the selected file is in plain text format (CSV, TSV, or TXT).
- **Column header row**: Select the **Column header row** checkbox to use the column names form the source table in the target table.
  - By selecting this option you can indicate what row in the file contains column names. The rows in the **Mapping** section, below, are filled with those names (and with the existing data types, unless you change them).
  - If you deselect this option, the first row is processed as data. To specify column names manually, enter a name for each target column in the **Mapping** section. (You will also have to enter data types.)
- **Start processing data at row**: Specifies the number of rows to skip when loading the source data into the target:
  - If you select the **Column header row** option under **Source column name** (see below) and if you enter a number greater than 0 in the **Start processing data at row** field, then that number of rows after the first row are not loaded into the target.
  - If you deselect the **Column header row** option under **Source column name**, and if you enter a number greater than 0 in the **Start processing data at row** field, then that number of rows including the first row are not loaded into the target.
- **Numeric column**: Select the **Convert invalid data to null** checkbox to convert an invalid numeric column value into a null value.
- **Newlines included in data values**: Select this option if there are newline characters or returns to the beginning of the current line without advancing downward in the data fields.

Selecting this option will increase the time taken to process the load. If you do not select this option when loading the data, the rows with newlines in the fields will be rejected. You can view the rejected row in the Job Report panel.

Settings Template

The save settings feature saves the configuration set in the Cart settings in the form of a JSON file. When opening the Settings template, you have the following options:

1. **Load Settings Template**: Loads a settings template from your local system.
2. **Save Settings Template**: Saves the current existing settings template.



You can use the **Load Settings Template** if you want to use an existing customized template present in your local.

1. From the Settings Template in the Settings tab of the Load Data page, select **Load Settings Template**.
2. You will see a Load Settings Template wizard, click the Settings Template JSON to load a JSON file from your system.
3. Clicking the Settings template JSON will open your local system. Click **OK** to load the JSON file.
4. After you load the JSON file, you can view the updates applied automatically to the settings tab which matches the JSON settings template you load from your local.

You can use the **Save Settings Template** to save the existing current Settings template.

1. From the Settings Template in the Settings tab of the Load Data page, select **Save Settings Template**.
2. The Template file editor appears. Click the Template File name and name the new template.
3. Click **OK** to finish saving the new name of the existing template.
4. You can test the configuration of the new template.

Bulk Edit Settings

You can use the Bulk edit settings to update all the columns at once from the mapping table. Use it to apply changes to the selection currently displayed in the results pane. You can search for the values of the column you want to edit in the search field and click the magnifier icon. The mapping table will display the results of the search. Select the Bulk Edit setting to update the column. The Bulk Edit setting allows you to:

- Update values of all the fields in a group.

- Find and replace, Add Prefix and Add suffix to target column name.
- Include the column(s) for loading data to the target table.
- Exclude the column(s) for loading data to the target table.

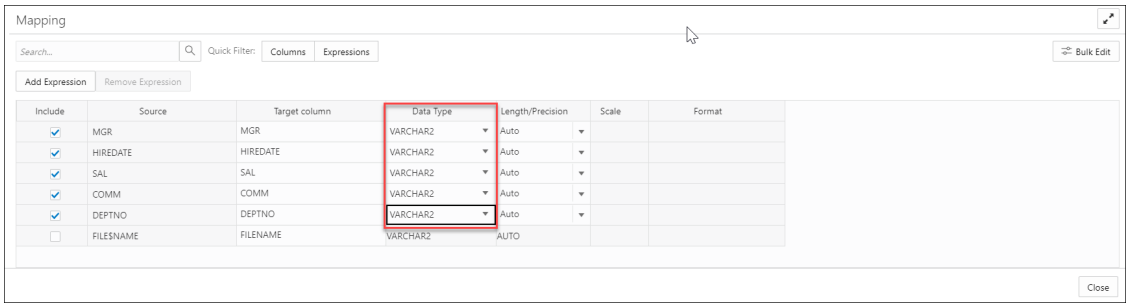
Searching the Column

The Bulk Edit setting updates the columns returned by the search field. The search box besides the Bulk Edit setting icon filters the list of columns you wish to update in a bulk. As soon as you start typing in the search field, the tool returns the field values which contains the letters you type. You can remove the filter by deleting all the content from the search box and clicking the magnifier icon that appears next to the search box.

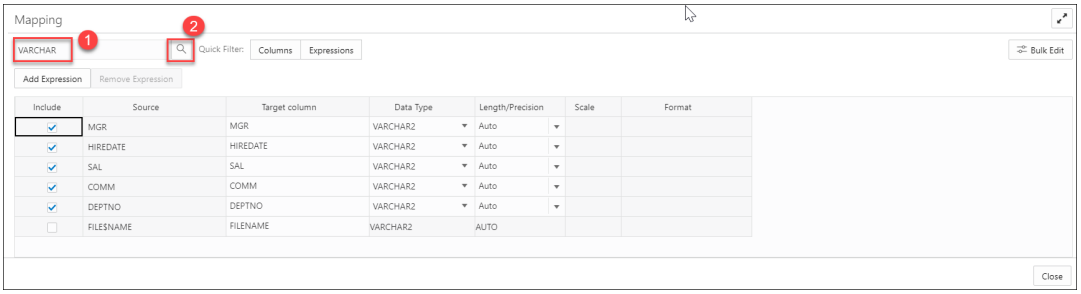
The Bulk Edit setting enables you to update the values of the following columns for all the searches returned by the search field:

- Data Type
- Target Column name
- Include Columns for loading
- Exclude columns for loading

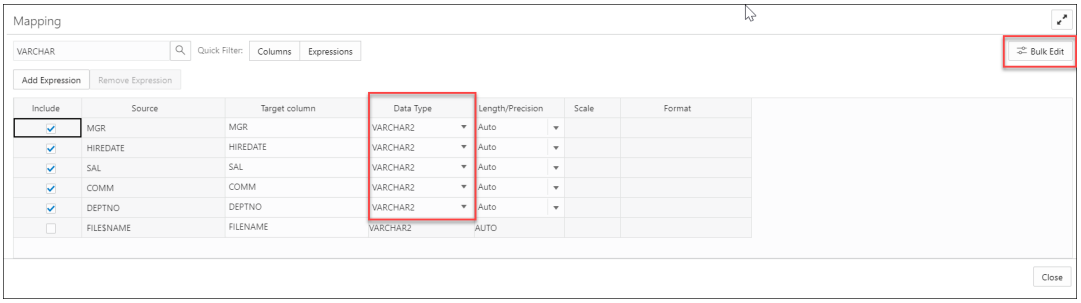
Consider changing the Data Type of first five rows from VARCHAR to NUMBER in the mapping table.



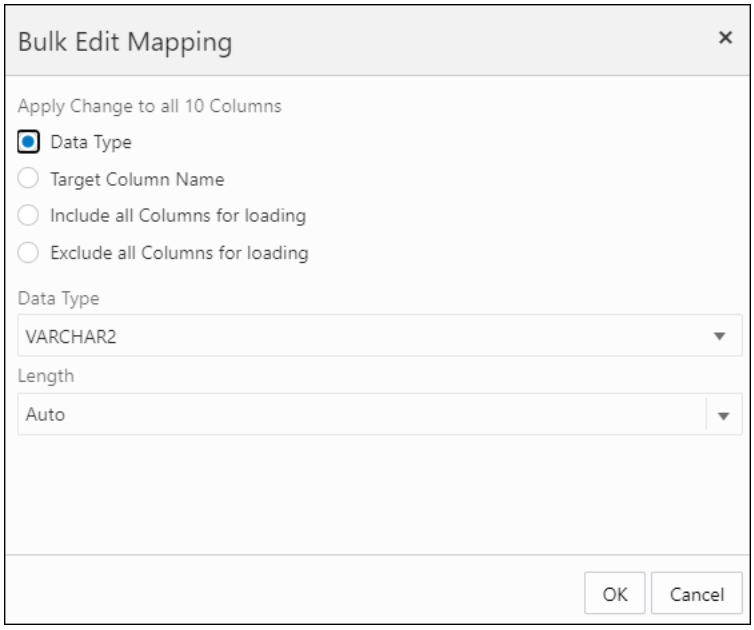
- Type `VARCHAR` in the search field and click on the magnifier icon besides the search field.



- The search will return the rows having `VARCHAR` as the Data Type. Select the **Bulk Edit** icon besides the magnifier icon.



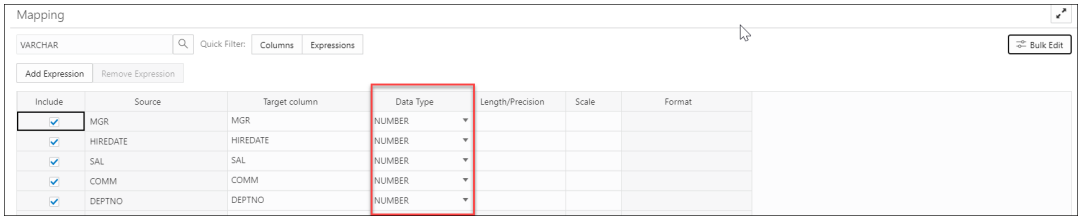
- Selecting the Bulk Edit icon opens the Bulk Edit Mapping dialog box.



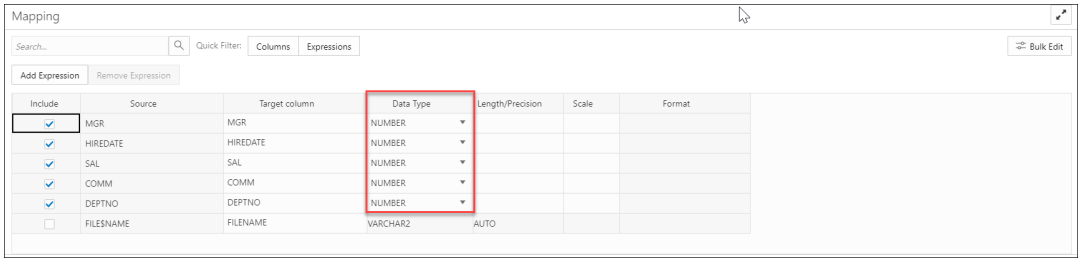
- Click **Data Type** and select `NUMBER` from the Data Type drop-down. Click **OK**.



- The Data Type of the selected rows changes from `VARCHAR` to `NUMBER`.



- Clear the content from the search field and click the magnifier icon to view the bulk edit updates in the Mapping table.



Settings Tab - Mapping Section

The settings in the **Mapping** section control how data from the source files are loaded into the rows of the target database table. For each row, the data from the column listed under **Source column** will be loaded into the column listed under **Target column**.

As mentioned above, the contents of the **Mapping** table change according to what processing option you chose in the **Table** section and which properties you set in the **Properties** section.

You can filter the results in the mapping table with **Quick Filter** field that enables you to filter out **Columns** or **Expressions**.

Select **Add Expression** to perform **Sentiment Analysis** or **Key Phrase extraction** with the source data. See Use OCI Language Service Capabilities in Data Studio for more details.

- **Include**: This check box ensures that the row you select is loaded from the source column to the target column.
- **Source column**: Lists the columns from the source file.

  If the **Column header row** option under **Properties** is selected, **Source column** shows the names of the columns in the source file. If the **Column header row** option is not selected, generic names like **COLUMN_1**, **COLUMN_2**, etc., are used. This field is always read only.

  You can view two source columns `FILE$NAME` and `SYSTIMESTAMP`. The `FILE$NAME` column enables you to locate the source file containing a particular data record. For example, you load a source file that contains a list of files. The file names in the file list refer to the department names across the organization. For instance, a *finance.txt* file contains data from the Finance department. In the mapping, you can use string data types to extract the department name from the output of the file name column. You can use the extracted department name to process data differently for each department.

  The `SYSTIMESTAMP` column allows us to view the current timestamp in the database.

> **Note:**
>
> – `FILE$NAME` and `SYSTIMESTAMP` source columns are not included by default. You must check the **Include** check box and run the load for the target table to display these two columns.
>
> – When you are creating a livefeed, the `FILE$NAME` and `SYSTIMESTAMP` source columns appear in the Mapping table by default.

- **Target column:** Lists the columns in the target table. Accept, select, or enter a column in the target table. You need to make sure that the target column is not empty. The target

column name cannot have duplicate name as another target column. The target column length must not be beyond 128 bytes. 128 byte limit is a database limit.

The contents of this column differ, depending on what you selected for the table processing **Option** and whether you selected for the **Column header row** option.

- If (1) the processing option is **Create Table** or **Drop Table and Create New Table** and (2) the **Column header row** option *is* selected, then the **Target column** uses the names of the columns in the source file. You can change the name of a target column by replacing the provided name with a new one.
- If (1) the processing option is **Create Table** or **Drop Table and Create New Table** and (2) the **Column header row** option is *not* selected, then generic names like **COLUMN_1**, **COLUMN_2**, etc., are used. You can change the name of a target column by replacing the provided name with a new one.
- If (1) the processing option is **Insert into Table**, **Replace Data**, or **Merge Into Table** and (2) the **Column header row** option *is* selected, then the **Target column** has a drop-down list of all the columns in the target table, with their data types. By default, the column with the name corresponding to the source column is selected, but you can select a different one from the list.
- If (1) the processing option is **Insert into Table**, **Replace Data**, or **Merge Into Table** and (2) the **Column header row** option is *not* selected, then the **Target column** has a drop-down list of all the columns in the target table, with their data types. Select a column from the list to use as the target column.

> **Note:**
>
> If you're loading multiple files from a folder in a single data load job, only the first file will be shown in the **Mapping** section. However, as long as the column names and data types match, the data from all source files will be loaded.

- **Data Type:** Lists the data type to use for data in that column. This column is displayed only for **Create Table** or **Drop Table and Create New Table**. The contents change depending on whether the **Get from file header** option is selected.
  - If the **Get from file header** option *is* selected, **Data type** shows the data types of the columns in the source file (for **Create Table**) or in the existing table (for **Drop Table and Create New Table**). If you want to change the data type for the target, click the name and select a different one from the list.
  - If the **Column header row** option is *not* selected, **Data type** shows all available data types. Select the data type to use for the target column from the list.
- **Length/Precision** (Optional): For columns where the **Data Type** is **NUMBER**, enter the length/precision for the numbers in the column. Precision is the number of significant digits in a number. Precision can range from 1 to 38.

  For columns where Data Type is VARCHAR2, the **Auto** value in **Length/Precision** field enables the Auto Size feature.

  With the Auto-Size column Width feature, you can automatically size any column to fit the largest value in the column. Select **Auto** from the **Length/Precision** drop-down values or pick a value from the drop-down list.
- **Scale** (Optional): For columns where the **Data Type** is **NUMBER**, enter the scale for the numbers in the column. Scale is the number of digits to the right (positive) or left (negative) of the decimal point. Scale can range from ranges from -84 to 127.

- **Format:** If the data type in the **Data type** column is **DATE** or one of the **TIMESTAMP** types, select a format for that type from the from the **Format** drop-down list.
- **Merge Key:** This option is used only for the processing option **Merge into Table**.

  For the **Merge into Table** option, you must specify at least one column to use as a key column. Merge keys are one or more columns that uniquely identify each row in the table. To specify a key column, select the **Merge Key** checkbox for the column. Merge keys must not contain any null values. For loading tables with primary keys, this option automatically enables the selection of primary key columns as the merge keys.

Preview Tab

The **Preview** tab displays the source data in tabular form in the **Load Preview** menu. The display reflects the settings you chose in the **Properties** section.

The **File** menu in the **Preview** tab displays the source data with all the columns with their names. You can copy the text.

Table Tab

The **Table** tab displays what the target table is expected to look like after the data has been loaded. If you chose the **Create Table** processing option, no table is shown.

SQL Tab

The **SQL** tab displays the SQL commands that will be run to complete this data load job.

> **Note:**
>
> You can see the SQL code even before the table is created.

Close Button - Save and Close the Pane

After entering all the details for the data load job, click **Close** at the bottom of the page. This saves the details you entered and returns you to the Load Data from Cloud Storage pane. To close the page without saving your entries, press **Escape**.

Run the Data Load Job

Once you've added data sources to the data load cart and entered details about the data load job, you can run the job.

To run the job:

1. If you haven't already done so, click the **Close** button in the **Load Data from Cloud Storage** pane to save your settings and close the pane. If any of the settings are invalid, an error message reports the problem. Fix the problem and click **Close**.
2. Click **Start** in the data load cart menu bar. To stop the data load job, click **Stop**.

   When you have added all of the sources for the job and specified the settings for each source, to run the job click the Start icon in the Data Load Cart menu bar. In the Run Data Load Job dialog box, click **Start**. To stop the data load job, click the Stop icon.

   Once the data load job starts, you can view the progress of the job in the Data Load dashboard.

When the data load job completes, the Data Load dashboard page displays the results of the job. At the top of the header of the table load, you can view the name of the table along with the total columns present in the table.

Click **Job Report** to view the total number of rows processed successfully and the count of rejected rows. You can also view the Start time. The SQL pane of the Job Report displays the equivalent SQL code of the job.

To view information about an item in the job, click the **Actions** icon on the Table Load.

To view a log of the load operation, click the Logging icon. You can save the log, clear it, or refresh it. Click **OK** to dismiss the log.

View Details About the Data Load Job After It Is Run

To view details about the data load job after it is run, click the **Actions** icon on the **Table and View Loads** section of the Data Load dashboard.

You can view Table details, Gather Statistics, Register to Cloud Link, Create Analytic View, Export Data to Cloud, Edit Table and delete the table.

When the data load job completes, the Data Load dashboard page displays the results of the job. At the top of the header of the table load, you can view the name of the table along with the total columns present in the table.

Click **Report** to view the total number of rows processed successfully and the count of rejected rows. You can also view the Start time. The SQL pane of the Report displays the equivalent SQL code of the job.

After your data load job, you might see errors that you want to correct, or upon inspection, realize that you wanted to name a column differently. In such cases, you will view a warning sign on the Table load. Click the **Reload** icon to reload source with suggested fixes. Click **Actions** icon on the Table load and select **Edit** to make any changes to the data load job (i.e., change a column name).

View the Table Resulting from the Data Load Job

To view the new tables or tables modified by the data load job, you can:

1. Fix your data load job. After your data load job, you might see errors that you want to correct, or upon inspection, realize that you wanted to name a column differently. In such cases, click the **Reload** option on the selected Table Load to reload cards from your recent cart and edit them as you did before your first attempt. The Reload icon reloads the source data with the fixes suggested by the tool. Click the **Actions** icon on the Table header, click **Table** and select **Edit** to make any changes to the data load job (i.e., change a column name).
2. You can use the **Edit** icon as mentioned in the above paragraph to review the table.

Linking Data

You can link to data in remote databases or in cloud storage buckets.

When you link to data in a remote database or in cloud storage, the target object produced is an external table or a view. When you select that target table or view on the Data Load Jobs page, the source preview for the object shows the current data in the source object.

Linking to columns in a remote database or to files in cloud storage can be useful when the source data is being continually updated. For example, if you have a database table with columns that are updated for each new sales transaction, and you have run a data load job that links columns of the source table to targets in your Oracle Autonomous Database, then those targets have the new sales data as it is updated in the source.

- Linking to Other Databases
  To link to data in tables in another database from your Oracle Autonomous Database, on the Data Load Dashboard, click **LINK DATA**, then click **Database**. Select a database link from the drop-down list. Drag one or more tables from the list of database tables and drop them in the Data Load Cart.
- Linking to Objects in Cloud Storage
  When you create a link to files in a cloud store bucket from you Oracle Autonomous database, you create an external table that links to the files in the cloud store.
- Linking to File System
  You can link to file system directories from your Autonomous Database.
- Linking to Share
  You can subscribe to a Share Provider for linking data.
- Linking to Data Catalog
  You can link your Data Catalogs to Autonomous Database.

Linking to File System

You can link to file system directories from your Autonomous Database.

On the Data Load dashboard page, click **Link Data** and select **File System**. On the left side of the page is a navigator pane, where you choose a directory with the folders or files containing the data. On the right of the page is the data link cart, where you stage the files and folders for the data link job. You can set options for the data link job before running it. The Autonomous Database comes with predefined CPU/IO shares assigned to different consumer groups. You can set the consumer group to either low, medium or high while executing a data load job depending on your workload.

Drag and drop the files or folders from the directory you select to add them to linking job.

Enter Details for the Data Link Job

Enter the details about the data link job in the Link Data from Directory dialog box. For more details, refer to the Linking to Objects in Cloud Storage chapter.

Linking to Share

You can subscribe to a Share Provider for linking data.

On the Data Load Dashboard page, click **Link Data** and select **Share**. For more details, refer to the Loading data from a Share chapter.

Linking to Data Catalog

You can link your Data Catalogs to Autonomous Database.

Click **OCI Data Catalog** in the Link Data page under the Data Load menu to register a Data Catalog for linking data. Refer to the Register Data Catalog chapter for more details on this.

Linking to Other Databases

To link to data in tables in another database from your Oracle Autonomous Database, on the Data Load Dashboard, click **LINK DATA**, then click **Database**. Select a database link from the drop-down list. Drag one or more tables from the list of database tables and drop them in the Data Load Cart.

Each table appears as an item in the Data Load Cart. The item shows the name of the table and the number of rows in it, and the name of the table that is the target for the data load.

To remove a table from the Data Load Cart, select the Remove (trash can) icon for the item. To remove all tables from the cart, click the Remove All (trash can) icon in the Data Load Cart menu bar.

To add a remote database to the list of databases, create a database link to the remote database. For information on creating a database link, see Database Links in *Oracle® Database Database Administrator's Guide*.

The databases available to you appear in the drop-down list of the database navigation pane of the Link Tables page.

Specify the Target, Create Filters, and View Mappings

To specify settings for the data load job, preview the data in the source or the target, and see statistics about the data, click the Settings (pencil) icon for the item in the Data Load Cart.

In the settings pane, on the Settings tab, the **Source** field displays the name of the table and the number of rows in the table.

Preview Source Data

To view the data in the source table, in the settings pane select the **Source Preview** tab. The source preview displays the data in the table.

View Statistics

To view statistics about the source table, in the settings pane select the **Statistics** tab. It may take a moment for the statistics to appear. The statistics include the size of the table, the number of rows and columns, the column names, data types, number of distinct values, and other information. Below the details about the columns is a bar graph that displays the top unique values for the selected column.

Preview Target Data

To view the data in the target view, in the settings pane select the **Target Preview** tab. The target preview displays the data in the target view. If the view does not yet exist, then the target data is the same as the source data, regardless of any filters set for the data load job.

Run the Data Load Job

When you have added all of the source tables for the job and specified the settings for each table, to run the job click the Start icon in the Data Load Cart menu bar. In the Run Data Load Job dialog box, click **Start**. To stop the data load job, click the Stop icon.

At the top of the page, the **Status** shows the number of items for which the load has completed over the number of items in the job, and the total time elapsed for the job. When the data load job completes, the Link Tables page displays the results of the job.

To view information about an item in the job, click the Settings (circled i) icon in the item. The settings pane has the same **Settings**, **Preview**, and **Target Preview** tabs as the settings pane before running the job, except that the **Target Preview** now contains the data loaded by the data load job. To close the settings pane, click **Close**.

To view a log of the load operation, click the Logging icon. You can save the log, clear it, or refresh it. Click **OK** to dismiss the log.

The list of views on the Data Load page contains any new views created. A preexisting view that was the target for the data load job now contains the loaded data.

On the Database links page, click the **Reload Cart** button. Clicking the button enables you to add the data sources to the data load the cart again.

Click **Done** to return to the Database Actions page.

Linking to Objects in Cloud Storage

When you create a link to files in a cloud store bucket from you Oracle Autonomous database, you create an external table that links to the files in the cloud store.

You can link to files in these file formats: AVRO, CSV, JSON, GeoJSON, Parquet, ORC, Delimited TXT. For information on supported file formats, see Format Specifications for JSON, AVRO, and XML Files.

Configure and run a data link job from the Link Cloud Object page. To open that page:

1. Open the Database Actions launchpad , click **Data Studio** tab and select the **Data Load** menu. See, The Data Load Page.
2. Select **LINK DATA** and **CLOUD STORE**.

On the left side of the page is a *navigator pane*, where you choose a cloud store connection and the folders or files containing the data. On the right of the page is the data load *cart*, where you stage the files and folders for the data link job. You can set options for the data link job before running it. The Autonomous Database comes with predefined CPU/IO shares assigned to different consumer groups. You can set the consumer group to either low, medium or high while executing a data load job depending on your workload.

- Manage Cloud Storage Links for Data Link Jobs
  Before you can link to data in a cloud store, you must establish a connection to the cloud store you want to use.
- Prepare the Data Link Job
- Add Files or Folders for the Data Link Job
- Enter Details for the Data Link Job
  Enter the details about the data link job in the Link Data from Cloud Storage pane.
- Run the Data Link Job
  Once you've added data sources to the data link cart and entered details about the data link job, you can run the job.
- View Details About the Data Link Job After It Is Run
  You can view the progress of the job on the Data Load dashboard.
- View the Table Resulting from the Data Link Job
  After running a data link job, you can view the table created by the data link job on the Data Load dashboard.

Manage Cloud Storage Links for Data Link Jobs

Before you can link to data in a cloud store, you must establish a connection to the cloud store you want to use.

On the Link Cloud Object page:

1. Click the Manage cloud store icon besides the field where you enter the cloud store location. Select **+ Create Cloud Store Location**.
2. Enter your information in the **Add Cloud Store Location** pane. See to add cloud storage location.

See Managing Connections.

To return to the Link Cloud Object page, click **Data Load** in the breadcrumbs at the top of the page and then navigate back to the page.

Prepare the Data Link Job

You may have to adjust your source data or your target table so that the source data links correctly to the external target table. Consider:

- If you're linking to multiple files, you must make sure that:
  - All the source files are of the same type, for example, CSV, JSON, etc.
  - The number, order, and data types of the columns in all the source files match.
- If you want to partition by date:
  - The source file must contain data where the data type is date or timestamp.
  - You must load a folder containing two or more data sources.
  - The names of the files in the folder must indicate a date or dates, for example, `MAR-1999.csv` or `2017-04-21.xlsx`.

Add Files or Folders for the Data Link Job

Add files from the cloud store to the data link cart, where you can edit the details of the data link job. To add the files:

1. From the list at the top of the navigator pane on the left, select the bucket with your source data.

   The list shows links that were established on the Manage Cloud Storage page. If you haven't yet registered the cloud store you want to use, click the **Connections** button under the Data Load menu in Data Studio suite of tools and register a connection.

2. Drag one or more items from the file navigator on the left and drop them into the cart on the right.
   - You can add files, folders, or both. A card is added to the cart for each file or folder you drag into it. The card lists the name of the source file or folder and a proposed name for the target table.
   - If you add a folder that contains multiple files, all the files must be of the same type, that is, CSV, TXT, etc.

     When you add the folder to the cart, a prompt is displayed that asks if you want to load all the objects from the multiple source files into a single target table. Click **Yes** to continue or **No** to cancel.
   - When you add multiple individual files or multiple folders to the cart, the data represented by each card will be loaded into a separate table, but all the items in the cart will be processed as part of the same data load job.
   - You can add files or folders from a different bucket, but if you do that, you're prompted to remove all files that are already in the cart before proceeding. To select files from a different bucket, select the bucket from the drop-down list in the navigator pane on the left and then add the file(s), as described above.
   - You can drop files or folders into the data load cart and then navigate away from the Data Link Object page. When you return to the page, those items remain on the page, but you may receive a message, "Remove All Data Link Items. Changing to another Cloud storage location requires all items to be removed from the data load job. Do you wish to continue?" Click **Yes** to remove the items from the cart. Click **No** to keep the items in the cart. Then you can continue to work.

You can remove items from the cart before running the data link job:

- To remove an item from the cart, select **Remove** on the card for the item Data Link cart menu bar at the top of the pane.

- To remove all items from the cart, click **Remove All** in the data link cart menu bar at the top of the pane.

Enter Details for the Data Link Job

Enter the details about the data link job in the Link Data from Cloud Storage pane.

On the card in the data link cart, click **Settings** to open the Link Data from Cloud Storage pane for that job. The pane contains:

- Settings Tab - Table Section
- Settings Tab - Properties Section
- Settings Tab - Mapping Section
- File Tab
- Table Tab
- SQL Tab
- Close Button - Save and Close the Pane

Settings Tab - Table Section

Set details about the target table in the **Table** section.

- **Name:** The name of the target table.
- **Partition Column:**

  List Partitions and Date-based partitions are the different types of partitions available in data linking.

  List partitioning is required when you specifically want to map rows to partitions based on discrete values.

  To partition according to a specific column, click the **Partition Column** drop-down list and select the column you want to use for the partitioning.

  You will have N files per partition value, all partitioned by the partition column you select.

  > **Note:**
  >
  > – For linked files (from external tables) there is also a requirement that for each file, the list partitioning column can contain only a single distinct value across all of the rows.
  >
  > – If a file is list partitioned, the partitioning key can only consist of a single column of the table.

  Date-based partitioning is available when you link a folder containing two or more data sources that have columns that contain date or timestamp data.

  To partition according to date, click the **Partition Column** drop-down list and select the **DATE** or **TIMESTAMP** column you want to use for the partitioning.

- **Validation Type:** Validation examines the source files, optional partitioning information, and report rows that do not match the format options specified. Select **None** for no validation; select **Sample** to perform validation based on a sample of the data; or select **Full** to perform validation based on all the data.

- **Use Wildcard:** This check box enables use of wildcard characters in search condition to retrieve specific group of files that matches the filter criteria.
  You can use a wildcard character, such as an asterisk (*) that searches, filters, and specifies groups of files that detect and add new files to the external table.

  For example, if you enter file*, then file01, file02, file03, and so on are considered to match the keyword. The asterisk (*) matches zero or more characters of the possibilities, to the keyword.

> **Note:**
>
> The wildcard support is incompatible with partitioning. The validation of source file fails if you use wildcards with partitioned data.

Settings Tab - Properties Section

Specify options to control how the source data is interpreted, previewed, and processed. These options vary, depending on the type of source data.

- **Encoding:** Select a character encoding type from the list. This option is available when the linked file is in plain text format (CSV, TSV, or TXT). The default encoding type is UTF-8.
- **Text enclosure:** Select the character for enclosing text: **"** (double-quote character), **'** (single-quote character) or **None**. This option is visible only when the selected file is in plain text format (CSV, TSV, or TXT).
- **Field delimiter:** Select the delimiter character used to separate columns in the source. For example, if the source file uses semicolons to delimit the columns, select **Semicolon** from this list. The default is **Comma**. This option is visible only when the selected file is in plain text format (CSV, TSV, or TXT).
- **Start processing data at row:** Specifies the number of rows to skip when linking the source data to the target external table:
  - If you select the **Column header row** option under **Source column name** (see below) and if you enter a number greater than 0 in the **Start processing data at row** field, then that number of rows after the first row are not linked to the target.
  - If you deselect the **Column header row** option under **Source column name**, and if you enter a number greater than 0 in the **Start processing data at row** field, then that number of rows including the first row are not linked to the target.
- **Source column name:** Select the **Column header row** checkbox to use the column names form the source table in the target table.
  - If you select this option, the first row in the file is processed as column names. The rows in the **Mapping** section, below, are filled with those names (and with the existing data types, unless you change them).
  - If you deselect this option, the first row is processed as data. To specify column names manually, enter a name for each target column in the **Mapping** section. (You will also have to enter data types.)
- **Numeric column**: Select the **Convert invalid data to null** checkbox to convert an invalid numeric column value into a null value.
- **Newlines included in data values**: Select this option if there are newline characters or returns to the beginning of the current line without advancing downward in the data fields. Selecting this option will increase the time taken to process the load. If you do not select

this option when loading the data, the rows with newlines in the fields will be rejected. You can view the rejected row in the Job Report panel.

Settings Tab - Mapping Section

The settings in the **Mapping** section control how data from the source files are linked to the rows of the target external table. For each row, the data from the column listed under **Source column** will be linked to the column listed under **Target column**.

- **Source column**: Lists the columns from the source file.

  If the **Column header row** option under **Properties** is selected, **Source column** shows the names of the columns in the source file. If the **Column header row** option is not selected, generic names like **COLUMN_1**, **COLUMN_2**, etc., are used. This field is always read only.

  You can view two source columns `FILE$NAME` and `SYSTIMESTAMP`. The `FILE$NAME` column enables you to locate the source file containing a particular data record. For example, you load a source file that contains a list of files. The file names in the file list refer to the department names across the organization. For instance, a *finance.txt* file contains data from the Finance department. In the mapping, you can use string data types to extract the department name from the output of the file name column. You can use the extracted department name to process data differently for each department.

  The `SYSTIMESTAMP` column allows us to view the current timestamp in the database.

  > **Note:**
  >
  > – `FILE$NAME` and `SYSTIMESTAMP` source columns are not included by default. You must check the **Include** check box and run the load for the target table to display these two columns.
  >
  > – When you are creating a livefeed, the `FILE$NAME` and `SYSTIMESTAMP` source columns appear in the Mapping table by default.

- **Target column:** Lists the columns in the target table.
  – If the **Column header row** option *is* selected, then the **Target column** uses the names of the columns in the source file. You can change the name of a target column by replacing the provided name with a new one. You need to make sure that the target column is not empty. Target column name must not be a duplicate of another target column. The target column name cannot have duplicate name as another target column. The target column length must not be beyond 128 bytes. 128 byte limit is a database limit.
  – If the **Column header row** option is *not* selected, then generic names like **COLUMN_1**, **COLUMN_2**, etc., are used. You can change the name of a target column by replacing the provided name with a new one.

  > **Note:**
  >
  > If you're linking multiple files from a folder in a single data link job, only the first file will be shown in the **Mapping** section. However, as long as the column names and data types match, the data from all source files will be linked.

- **Data Type:** Lists the data type to use for data in that column. The contents change depending on whether the **Get from file header** option is selected.
  - If the **Column header row** option *is* selected, **Data type** shows the data types of the columns in the source file. If you want to change the data type for the target, click the name and select a different one from the list.
  - If the **Column header row** option is *not* selected, **Data type** shows all available data types. Select the data type to use for the target column from the list.
- **Length/Precision** (Optional): For columns where the **Data Type** is **NUMBER**, enter the length/precision for the numbers in the column. Precision is the number of significant digits in a number. Precision can range from 1 to 38.

  For columns where Data Type is VARCHAR2, the **Auto** value in **Length/Precision** field enables the Auto Size feature.

  With the Auto-Size column Width feature, you can automatically size any column to fit the largest value in the column. Select **Auto** from the **Length/Precision** drop-down values or pick a value from the drop-down list.
- **Scale** (Optional): For columns where the **Data Type** is **NUMBER**, enter the scale for the numbers in the column. Scale is the number of digits to the right (positive) or left (negative) of the decimal point. Scale can range from ranges from -84 to 127.
- **Format:** If the data type in the **Data type** column is **DATE** or one of the **TIMESTAMP** types, select a format for that type from the from the **Format** drop-down list.

Preview Tab

The **Load Preview** menu in the **Preview** tab displays the source data in tabular form. The display reflects the settings you chose in the **Properties** section. The **File** menu displays source data with the column names.

If you dragged a folder containing multiple files into the data link cart and then clicked  **Settings** for that card, the **Preview** pane includes a **Preview Object (File)** drop-down list at the top of the pane that lists all the files in the folder. Select the source file you want to preview from that list.

Table Tab

The **Table** tab displays what the target table is expected to look like after the data has been linked.

SQL Tab

The **SQL** tab displays the SQL commands that will be run to complete this data link job.

> **Note:**
>
> You can see the SQL code even before the table is created.

Close Button - Save and Close the Pane

After entering all the details for the data link job, click **Close** at the bottom of the page. This saves the details you entered and returns you to the Link Data from Cloud Storage pane.

Run the Data Link Job

Once you've added data sources to the data link cart and entered details about the data link job, you can run the job.

To run the job:

1. If you haven't already done so, click the **Close** button in the **Link Data from Cloud Storage** pane to save your settings and close the pane. If any of the settings are invalid, an error message reports the problem. Fix the problem and click **Close**.

2. Click  **Start** in the data link cart menu bar. To stop the data link job, click  **Stop**.

   When the data link job completes, the **Data Load Dashboard** page displays the results of the job under **Table and View Loads** section.

   Once the data link job starts, you can view the progress of the job in the Data Load dashboard.

View Details About the Data Link Job After It Is Run

You can view the progress of the job on the Data Load dashboard.

When the data load job completes, the Data Load dashboard page displays the results of the job. At the top of the header of the table load, you can view the name of the table along with the total columns present in the table.

Click **Job Report** to view the total number of rows processed successfully and the count of rejected rows. You can also view the Start time. The SQL pane of the Job Report displays the equivalent SQL code of the job.

To view information about an item in the job, click the **Actions** icon on the Table Load.

To view a log of the load operation, click the Logging icon. You can save the log, clear it, or refresh it. Click **OK** to dismiss the log.

View the Table Resulting from the Data Link Job

After running a data link job, you can view the table created by the data link job on the Data Load dashboard.

Fix your data load job. After your data load job, you might see errors that you want to correct, or upon inspection, realize that you wanted to name a column differently. In such cases, click the **Reload** option on the selected Table Load to reload cards from your recent cart and edit them as you did before your first attempt. The Reload icon reloads the source data with the fixes suggested by the tool. Click the **Actions** icon on the Table header, click **Table** and select **Edit** to make any changes to the data load job (i.e., change a column name).

Feeding Data

You can run a live table feed on demand, on a schedule, or as the result of a notification.

A Live Table Feed automates loading of data into a table in your database. Files automatically load as they appear in object storage and the Live Table Feed system ensures that files are only loaded once. The loading can happen manually, by a schedule, or even by notifications delivered directly from Object Storage.

The bucket can contain files in these formats: AVRO, CSV, JSON, GeoJSON, Parquet, ORC, Delimited TXT. All of the files must have the same column signature.

About the Live Feed Page

On the Database Actions - Data Load Dashboard page select **FEED DATA** to display the Live Feed page. On this page, you can:

- Manage Cloud Storage Connections for Live Table Feeds
- Create a Live Table Feed Object
- List, Filter, and Sort Live Table Feed Objects
- Find and View Live Table Feed Objects
- Run a Live Table Feed
- Delete a Live Table Feed

Manage Cloud Storage Connections for Live Table Feeds

Before you create a live table feed, you must establish a connection to the cloud store you want to use:

1. Click **Connections** under the Data Load menu. For instructions, see Managing Connections.

Create a Live Table Feed Object

To create a live table feed object,

1. On the Live Feed page, click the **+ Create Live Table Feed** button to display the Live Feed Settings pane. Enter information on the Data Source tab as follows:
   - **Cloud Store Location**: Select the Cloud Store Location from the drop-down. Select the cloud connection for the bucket containing the file you want to use for feeding data.

     On the **Basic** mode, you can view the following options:
   - **Folders**:
     Select the folder containing the file(s) you want to use for feeding data on the object store. Select **Entire Bucket** to load all the files in your bucket into your table. The folders are listed and organized in the drop-down based on how you create folders or directories and store your files. For example, you can create a sales folder to store *sales1.csv* and *sales2.csv* files.
   - **Extensions**:
     Enter an extension to limit the live table feed to only those files in the bucket that match the extension. For example, to limit the files to only those that are CSV files, select **CSV**.

     On the **Advanced** mode, you can view the following options:
   - **Object Filter (glob):** Enter a file glob to limit the live table feed to only those files in the bucket that match the glob. For example, to limit the files to only those that are CSV files, enter ***.CSV**.
   - Under the Live Feed File Preview section, you can view a preview of the file you select in the previous step.

   Click **Next** to progress to the **Table Settings** tab.

   On the **Option** field, select any of the two available options:
   - **Load Table**: This option appends the feed into the target table. It adds new rows to your table for every file it views.
   - **Merge into Table**: This option merges live feed data into the target table. Click **Merge Key** in the Mapping section when the row already exists, and the loaded key matches the newly added feed. This avoids insertion of a new row.
   - **Load Collection**: You can define your livefeed over JSON files with this option. By selecting this option, you will not view the mapping section and Add Expression

section. The tool loads the source data into a JSON collection which you can extract later.

Target Table Name: Accept the default name or enter a different name. This is the name of the target table that data from the live feed will be loaded into in your Autonomous Database instance. If the table does not exist, Live Feed will attempt to guess the correct columns. You can pre-create the table in which you wish the Live Feed to load into. This is for higher accuracy.

The Table Settings tab specifies options to control how the source data is interpreted, previewed, and processed. These options vary, depending on the type of source data.

- **Encoding:** Select a character encoding type from the list. This option is available when the linked file is in plain text format (CSV, TSV, or TXT). The default encoding type is UTF-8.
- **Text enclosure:** Select the character for enclosing text: **"** (double-quote character),**'** (single-quote ch **None**. This option is visible only when the selected file is in plain text format (CSV, TSV, or TXT).
- **Field delimiter:** Select the delimiter character used to separate columns in the source. For example, if the source file uses semicolons to delimit the columns, select **Semicolon** from this list. The default is **Comma**. This option is visible only when the selected file is in plain text format (CSV, TSV, or TXT).
- **Start processing data at row:** Specifies the number of rows to skip when linking the source data to the target external table:
  If you select the **Column header row** option under **Source column name** (see below) and if you enter a number greater than 0 in the **Start processing data at row**field, then that number of rows after the first row are not linked to the target.

  If you deselect the **Column header row** option under **Source column name**, and if you enter a number greater than 0 in the **Start processing data at row** field, then that number of rows including the first row are not linked to the target.

  Column header row: Select the **Column header row** checkbox to use the column names form the source table in the target table.

  If you select this option, the first row in the file is processed as column names. The rows in the **Mapping** section, below, are filled with those names (and with the existing data types, unless you change them).

  If you deselect this option, the first row is processed as data. To specify column names manually, enter a name for each target column in the **Mapping** section. (You will also have to enter data types).

  Select the **Convert invalid data to null** checkbox to convert an invalid numeric column value into a null value.

  Newlines included in data values: Select this option if there are newline characters or returns to the beginning of the current line without advancing downward in the data fields. Selecting this option will increase the time taken to process the load. If you do not select this option when loading the data, the rows with newlines in the fields will be rejected. You can view the rejected row in the Job Report panel.
- Edit or update the table settings in the **Mapping** section: In this pane, the mapping of the source to target columns is displayed.
  The contents of the **Mapping** table change according to what processing option you chose in the **Table** section and which properties you set in the **Properties** section.

  You can filter the results in the mapping table with **Quick Filter** field that enables you to filter out **Columns** or **Expressions**.

Select **Add Expression** to perform **Sentiment Analysis** or **Key Phrase extraction** or **Language Detection** or **Text Translation**with the source data. See Use OCI Language Service Capabilities in Data Studio for more details.

- Select the Include check box at the beginning of a row to add the column to the target table.
- Select or enter values for column attributes such as Target Column Name, Column Type, Precision, Scale, Default, Primary Key and Nullable.
- You need to review the suggested data type and if needed, modify it by entering the data type directly into the target cell.
  Review the generated mapping table code based on the selections made in the previous screens.

  Click **Merge Key** in the Mapping section when the row already exists, and the loaded key matches the newly added feed. This avoids insertion of a new row. This option shows up when you select the **Merge into Table** option.

Click **Next** to progress to the Preview tab.

- The Preview Pane displays the changes you make to the table.
- Click **Next** to progress to the Live Feed Settings tab.

On the Live Feed Settings tab, specify the following field values:

- **Live Table Feed Name**: Accept the default name or enter a different name to identify this live table feed.
- **Enable for Notification**: Select this option so that new or changed data in the data source will be loaded based on an Oracle Cloud Infrastructure *notification*. When you select this option, you can avoid delays that might occur when polling is initiated on a schedule (that is, if you selected the live table feed **Scheduled** option).

  When you select the **Enable for Notification** option, you must also:

  - Configure your object store bucket to emit notifications
  - Create a Notifications service subscription topic
  - Create an Events service rule
  - Copy the notification URL
  - Create a Notifications service subscription
  - Confirm that notifications are allowed

  For complete instructions, see Creating a Notification-Based Live Table Feed.
- **Enable for Scheduling:** Select this option to set up a schedule for running the live table feed object; that is, to poll the data source on a regular basis:
  - In the time interval fields, enter a number, and select a time type and the days on which to poll the bucket for new or changed files. For example, to poll every two hours on Monday, Wednesday, and Friday, enter **2**, select **Hours**. You can select **All Days**, **Monday to Friday**, **Sunday to Thursday**, or **Custom** from **Week Days** drop-down. The Custom field enables you to select **Monday**, **Tuesday**, **Wednesday**, **Thursday** and **Friday** in the appropriate fields.
  - Select a start and end date with start and end time. If you don't select a start date, the current time and date are used as the start date. The end date is optional. However, without an end date, the live feed will continue to poll.

Select a consumer group from the drop-down, namely, low, medium and high.

2. Click **Create** to create the Live Table feed object.

   Show code: Select this option to view the PL/SQL code equivalent of the Create Live Table Feed wizard. You can copy and execute this PL/SQL code in the worksheet to perform the same action that occurs when you click **Create** in the Create Live Table Feed wizard.

List, Filter, and Sort Live Table Feed Objects

When you open the Live Feed page, existing live table feed objects are displayed as cards on the page. They are identified as LIVE_TABLE_FEED entities.

To filter live table feed objects:

1. Click the search field at the top of the page to display filter options. By default, the live table feed objects from the current user's schema are shown. As soon as you start typing in the search field, the feed tool returns the values which contain the letters you type. You can remove the filter by deleting the content from search box and clicking the cross icon that appears next to the search box.
2. To include objects from other schemas, select the drop-down next to the search field, under **Schema**. To remove a schema from the filter list, deselect the box next to its name.
3. To show objects from all available schemas, select **All** from the **Schema** drop-down.

To sort live table feed objects

1. Click the **Sort by** button at the top right of the page.
2. Select a sorting option. To sort ascending, click the icon with the up arrow. To sort descending, click the icon with the down arrow.

Find and View Live Table Feed Objects

To search for available live table feed entities in the selected schemas, enter a value in the search field at the top of the page and press **Enter**. The display then includes only the entities whose names contain the characters in the search field. To clear the search field, click the Clear search results (X) icon in the search field.

To remove a schema or sorting value from the selected filters, deselect the schema or sorting value in the filter panel, or click the Remove filter (X) icon for the schema or sorting value above the display of live table feed objects. To close the filter panel, click the Hide filter panel (X) icon in the panel.

To refresh the display of live table feeds, click the Refresh icon at the top of the page.

Edit a Live Table Feed Object

To edit details of a live table feed object,

1. On the Live Feed page, find the card for the live table feed whose details you want to edit.
2. Click the Actions icon (three dots) on the card and select **Edit Live Table Feed**. You can edit the following options:
   - Enter information on the Data Source tab as follows:
   - **Cloud Store Location**: Select the Cloud Store Location from the drop-down. Select the cloud connection for the bucket containing the file you want to use for feeding data.
     - On the **Basic** mode, you can view the following options:
     - **Folders**:

Select the folder containing the file you want to use for feeding data on the object store. Select **Entire Bucket** to upload all the files to your bucket. The folders are listed and organized in the drop-down based on how you create folders or directories and store your files. For example, you can create a sales folder to store *sales1.csv* and *sales2.csv* files.

- **Extensions**:
  Enter an extension to limit the live table feed to only those files in the bucket that match the extension. For example, to limit the files to only those that are CSV files, select **CSV**.

  On the **Advanced** mode, you can view the following options:

- **Object Filter (glob):** Enter a file glob to limit the live table feed to only those files in the bucket that match the glob. For example, to limit the files to only those that are CSV files, enter ***.CSV**.

- On the **Live Feed Settings** tab, edit the following fields:
- **Enable for Notification**: Select this option so that new or changed data in the data source will be loaded based on an Oracle Cloud Infrastructure *notification*. When you select this option, you can prevent any delays that might occur when polling is initiated on a schedule (that is, the live table feed **Scheduled** option).

  When you select the **Enable for Notification** option, you must also:

  - Copy the live table feeds notification URL
  - Configure your cloud store to emit notifications
  - Configure Oracle Cloud Infrastructure to route events to the endpoint used for the live table feed.
  - Create a rule.
  - Create a subscription.
  - Confirm that notifications are allowed at the live feed service.

  For complete instructions, see Creating a Notification-Based Live Table Feed.

- **Scheduled:** Select this option to set up a schedule for running the live table feed object; that is, to poll the data source on a regular bases:
  - In the time interval fields, enter a number, and select a time type and the days on which to poll the bucket for new or changed files. For example, to poll every two hours on Monday, Wednesday, and Friday, enter `2`, select **Hours**. You can select **All Days**, **Monday to Friday**, **Sunday to Thursday**, or **Custom** from **Week Days** drop-down. The Custom field enables you to select **Monday**, **Tuesday**, **Wednesday**, **Thursday** and **Friday** in the appropriate fields.
  - Select a start and end date with start and end time.

3. Click **Save**.

Run a Live Table Feed

You can run a live table feed on demand, on a schedule, or as the result of a notification.

To run a live table feed on demand:

1. On the Live Feed page, find the card for the live table feed you want to run.
2. Click the Actions icon (three dots) on the card and select **Run Live Table Feed Immediately (Once)**.

To run a live table feed on a schedule:

You can set a schedule for running live table feeds on the **Create Live Table Feed** pane (when creating a new table feed) or the **Edit Live Table Feed** pane (when editing an existing table feed). See Create a Live Table Feed Object or Edit a Live Table Feed Object.

To run a live table feed as the result of a notification:

See Creating a Notification-Based Live Table Feed.

Select the **Scheduled** check box to display the schedule options and then set the schedule by selecting the options you want.

To view live table feed run details:

1. On the Live Feed page, find the card for the live table feed whose run details you want to see.
2. Click the Actions icon (three dots) on the card and select **Live Table Feed Run Details**. The **Objects** tab on the Live Table Feed Run Details pane displays information about the jobs, such as when the run occurred, the objects involved in the run, the table owner, the table Name, status of the live feed, the rows loaded and the rows rejected, and other details. Click the **All** tab to view more details, such as the event type.

Delete a Live Table Feed

1. On the Live Feed page, find the card for the live table feed job you want to delete.
2. Click the Actions icon (three dots) on the card and select **Delete Live Table Feed**.

- Creating a Notification-Based Live Table Feed
  You can load data through a live table feed based on an Oracle Cloud Infrastructure *notification*.
- Creating a Notification-Based Live Table Feed using Amazon Simple Storage Service (S3)
  You can integrate Amazon Simple Storage Service (S3) and Oracle Cloud Infrastructure (OCI) to automate the process of live feed notifications when storage objects it is observing have updates. The following section provides instructions for creating event notifications in your Amazon S3 bucket where your data files are stored.
- Creating a Notification-Based Live Table Feed using Microsoft Azure
  A notification-based Live Table Feed is an interface between Oracle Cloud Infrastructure and a third-party cloud message queuing service such as Azure Event Grid.
- Send Email using Live Feed Tool
  You can configure the Live Feed tool to send you email notifications when certain specific live feed events occur.

Creating a Notification-Based Live Table Feed

You can load data through a live table feed based on an Oracle Cloud Infrastructure *notification*.

In addition to being able to run a live table feed on demand or on a schedule, as described in Feeding Data, you can also run a feed as the result of a notification. When data in the source bucket is changed, a notification is sent which triggers a run of the table feed. With a notification-based live table feed, you can avoid any delay that might come from running on-demand or scheduled live table feed jobs.

> **Note:**
>
> - Notification-based live table feeds aren't available on the Oracle Cloud Infrastructure free tier. You must be on a paid tenancy with appropriate permissions on your account to use this feature.
> - Notification-based live table feeds aren't available on Oracle Autonomous Data Warehouse Databases (ADW) that are configured using a private endpoint.

To create a notification-based live table feed:

- Step 1: Configure your object store bucket to emit notifications
- Step 2: Create a Notifications service subscription topic
- Step 3: Create an Events service rule
- Step 4: Create and configure a live table feed to use notifications, and copy the notification URL
- Step 5: Create a Notifications service subscription
- Step 6: Confirm that the endpoint can receive notifications

> **Tip:**
>
> To complete those steps, you will alternate between Oracle Cloud Infrastructure Console pages and Oracle Database Actions pages. You may find it convenient to open the Cloud Console in one browser page or tab and Database Actions in another, so it's easy to move back and forth.

Step 1: Configure your object store bucket to emit notifications

Where: Oracle Cloud Infrastructure Console: Object Storage & Archive Storage - Buckets page

Configure the bucket containing your source data so that it will emit notifications when the data changes. You can set this option when you create a bucket or you can set it in an existing bucket.

1. Open the Cloud Console navigation menu and click **Storage**. Under **Object Storage and Archive Storage**, click **Buckets**.
2. **If you're creating a new bucket:**
   a. On the Buckets page, click the **Create Bucket** button to create a new bucket, as described in Managing Buckets. In the **Create Bucket** wizard, select the **Emit Object Events** option, along with the other options for your new bucket.
   b. Click **Create**.

   If you're using an existing bucket:
   a. On the Buckets page, click the name of the bucket you want to use, as described in Managing Buckets.
   b. On the Bucket Details page, click the **Edit** link next to **Emit Object Events**.
   c. Select the **Emit Objects Events** check box, and then click **Save Changes**.

Step 2: Create a Notifications service subscription topic

Where: Oracle Cloud Infrastructure Console: Notifications - Topics page

1. Open the Cloud Console navigation menu and click **Developer Services**. Under **Application Integration**, click **Notifications**.
2. Click **Create Topic**, enter a name and optional description, and then click **Create**.

Step 3: Create an Events service rule

Where: Oracle Cloud Infrastructure Console: Events - Rules page

1. Open the Cloud Console navigation menu and click **Observability & Management**. Under **Events Service**, click **Rules**.
2. Click **Create Rule**, and fill out the Create Rule page as described in Managing Rules for Events.
   - Under **Rule Conditions**, select:
     - **Condition:** Event Type
     - **Service Name:** Object Storage
     - **Event Type:** Object - Create
   - Under **Actions**, select:
     - **Action Type:** Notifications
     - **Notifications Compartment:** Select the compartment to use for the notifications.
     - **Topic:** Select the name of the topic you created above, in Step 2: Create a Notifications service subscription topic.
3. Click **Create Rule**.

Step 4: Create and configure a live table feed to use notifications, and copy the notification URL

Where: Database Actions: Live Feeds page

You can configure a new or an existing live table feed to use notifications:

1. Go to the Database Actions Live Feeds page, as described in Feeding Data.
2. Create or edit a live table feed object, as described in Create a Live Table Feed Object or Edit a Live Table Feed Object. Select the **Enable for Notification** option
3. Click **Create** or **Save**.
4. Click the Actions (three vertical dots) icon on the card for your live feed, and select **Show Confirmation URL**.
5. In the **Notification URL** dialog box, click the **Copy** icon to copy the URL to the clipboard. You may want to copy it to a temporary file, so you can retrieve it later. You'll use this URL in the next step, Step 5: Create a Notifications service subscription.

Step 5: Create a Notifications service subscription

Where: Oracle Cloud Infrastructure Console: Notifications - Subscriptions page

1. Return to the Oracle Cloud Infrastructure Console. Open the navigation menu and click **Developer Services**. Under **Application Integration**, click **Notifications**.

2. On the Notifications page, click the **Subscriptions** tab (on the left side of the page), the status will be **Active**.
3. Click **Create Subscription** and fill in the **Create Subscription** page:
   - **Subscription topic:** Select the subscription topic you created in Step 2: Create a Notifications service subscription topic.
   - **Protocol:** HTTPS (Custom URL)
   - **URL:** Paste in the URL you copied in Step 4: Create and configure a live table feed to use notifications, and copy the notification URL.
   - Click **Create**. The subscription will be listed in the **Subscriptions** table in a state of "Pending."

Step 6: Confirm that the endpoint can receive notifications

Where: Database Actions: Live Feeds page

1. Return to the Database Actions Live Feeds page and find the card for the live table feed you are configuring for a notification-based feed.
2. You will view the live feed card in an Active status.

You will receive email notifications when certain specific live feed events occur.

See Send Email using Live Feed Tool for more details.

Once you finish the above steps, any new files uploaded to the bucket will automatically be loaded into the live table feed table.

Creating a Notification-Based Live Table Feed using Amazon Simple Storage Service (S3)

You can integrate Amazon Simple Storage Service (S3) and Oracle Cloud Infrastructure (OCI) to automate the process of live feed notifications when storage objects it is observing have updates. The following section provides instructions for creating event notifications in your Amazon S3 bucket where your data files are stored.

> **Tip:**
>
> To complete these steps, you will need to alternate between Amazon Web Services (AWS) Management console and Oracle Database Actions pages. You may find it convenient to open the Amazon Web Services in one browser page or tab and Database Actions in another, so it is easy to move back and forth.

To create a notification- based live feed with Amazon S3 as cloud storage you must:

- Step 1: Create your object store bucket in Amazon S3
- Step 2: Create Access Keys
- Step 3: Add an OCI Cloud Storage using Amazon S3
- Step 4: Create and configure a live table feed to use notifications and copy the notification URL
- Step 5: Create a notifications service subscription topic
- Step 6: Enable and configure event notifications using the Amazon S3 console
- Step 7: Create a notifications service subscription

- Step 8: Confirm that the endpoint can receive notifications

Step 1: Create your object store bucket in Amazon S3

Where: Amazon Web Services (AWS) Management console

Configure and create your bucket containing source data so that it emits notifications when the data changes.

1. Log in to AWS Management console and open the Amazon S3 console.
2. On the home page click the **Create Bucket** icon.
3. In **Bucket name**, enter a valid name for your bucket. For example: testbucket. After you create the bucket, you cannot change its name.
4. In **Region**, select the Amazon Web Services (AWS) Region from the dropdown. For example: us-west-2
5. In Bucket settings for Block Public Access, select the **Block Public Access** settings that you want to apply to the bucket. It is recommended to keep all settings enabled unless you know that you need to turn any of them off.
6. Select **Advanced settings**, and accept all the default options if you want to enable S3 Object Lock. This step is optional.
7. Select **Create bucket**.

Step 2: Create Access Keys

Where: AWS Management console

To access Amazon Simple Notification Service (SNS), you must have credentials that Amazon Web Services (AWS) can use to validate your requests. These credentials must have permissions to access Amazon SNS topics. The following steps provide you details on steps to create access keys using AWS Identity and Access Management (IAM) for security purposes.

1. Log in to AWS Management console and open Amazon Identity and Access Management (IAM) console.
2. On the navigation menu, select **Users**.
3. Select your **user name**.
4. In the Security Credentials tab, select **Create access key**.
5. Copy the **Access key ID** and **Secret access key** in the display. Paste them in a clipboard.
6. To download the keys, select **Download.csv file** icon. This way you can store the file in a secure location.

Step 3: Add an Amazon S3 Cloud Storage Link

Where: Database Actions: Manage Cloud page

Before you create a live table feed, you must establish a connection to the cloud store you want to use.

1. Click the **Manage Cloud Store** button at the top of the page to go to the Manage Cloud page. For further instructions on adding source files residing in cloud storage provided by Amazon S3, refer to *Create an Amazon S3 Cloud Storage Link* topic in Managing Connections.

> **Note:**
>
> Paste the Access key ID and Secret access key generated in the previous step (Step 2: Create Access Keys) to their respective text fields in the **Add Cloud Storage** page.

Step 4: Create and configure a live table feed to use notifications, and copy the notification URL

Where: Database Actions: Live Feeds page

Creating a live table feed enables you to load data in real time from external storage sources to your table in ADB. External storage you use include as Oracle Object Store, AWS S3 or Microsoft Azure containers.

You can configure a new or an existing live table feed to use notifications:

1. Go to the Database Actions Live Feeds page, as described in Feeding Data.
2. Create or edit a live table feed object, as described in Create a Live Table Feed Object or Edit a Live Table Feed Object. Select the **Enable for Notification** option
3. Click **Create** or **Save**.
4. Click the **Actions** (three vertical dots) icon on the card for your live feed, and select **Show Notification URL**.
5. In the **Notification URL** dialog box, click the **Copy** icon to copy the URL to the clipboard. You may want to copy it to a temporary file, so you can retrieve it later. You will use this URL in the subsequent step (Step 7: Create a notifications service subscription).

Step 5: Create a notifications service subscription topic

Where: Amazon Simple Notification Service (SNS) console

You receive Amazon S3 notifications using Amazon Simple Notification Service (Amazon SNS) topic. You need to add a notification configuration to your bucket using an Amazon SNS topic. SNS topics are shared locations which are used to send notifications of various events that happen in AWS buckets.

During creation, you select a topic name and topic type. After creating a topic, you cannot change the topic type or name. All other configuration choices are optional during topic creation, which you can edit later.

To access any AWS service, you must first create an AWS account.

Navigate to the AWS Management console, and then select **Create an AWS Account**.

Follow the instructions as provided in the Amazon SNS link to create your first IAM administrator user and group. Now you can log in to any of the AWS services as an IAM user.

1. Log in to Amazon SNS console as an IAM user.
2. On the **Topics** page, select **Create topic**.
3. Specify the following fields on the **Create topic** page, in the **Details** section.
   - **Type:**Standard (Standard or FIFO)
   - **Name:** notify-topic. For a FIFO topic, add fifo to the end of the name.
   - **Display Name:** This field is optional.

4. Expand the **Encryption** section and select **Disable encryption**.
5. Expand the **Access policy** section and configure additional access permissions, if required. By default, only the topic owner can publish or subscribe to the topic. This step is optional. Edit the JSON format of the policy based on the topic details you enter. Here is a sample of Access policy in JSON format.

```
{ "Version": "2008-10-17",
"Id": "__default_policy_ID",
"Statement":[
{"Sid": "__default_statement_ID",
"Effect": "Allow",
"Principal": {"AWS": "*"
},"Action": [
        "SNS:Publish",
        "SNS:RemovePermission",
        "SNS:SetTopicAttributes",
        "SNS:DeleteTopic",
        "SNS:ListSubscriptionsByTopic",
        "SNS:GetTopicAttributes",
        "SNS:AddPermission",
        "SNS:Subscribe"
      ],
      "Resource": "arn:aws:sns:us-west-2:555555555555:notify-topic", //us-
west-2 is the region
      "Condition": {
        "StringEquals": {
          "AWS:SourceOwner": "555555555555"
        }
      }
    },
    {
      "Sid": "s3_policy", //This field accepts string values
      "Effect": "Allow",
      "Principal": {
        "Service": "s3.amazonaws.com"
      },
      "Action": [
        "SNS:Publish"
      ],
      "Resource": "arn:aws:sns:us-west-2:555555555555:notify-topic", //
notify-topic is the topic name
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "555555555555" //This is the Account ID
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:s3:*:*:testbucket /*testbucket is the
s3 bucket name. You will get notifications only when file is uploaded to
this
    bucket.*/
"
        }
      }
    }
```

```
    ]
}
```

6. Expand the **Delivery retry policy (HTTP/S)** section to configure how Amazon SNS retries failed message delivery attempts. This step is optional.
7. Expand the **Delivery status logging** section to configure how Amazon SNS logs the delivery of messages to CloudWatch. This step is optional.
8. Expand **Tags** section to add metadata tags to the topic. This step is optional.
9. Select **Create topic**.
10. The topic's Name, ARN (Amazon Resource Name), and Topic owner's AWS account ID are displayed in the Details section.
11. Copy the topic ARN to the clipboard.

Step 6: Enable and configure event notifications using the Amazon S3 console

Where: Amazon S3 Management console

You can enable Amazon S3 bucket events to send a notification message to a destination whenever those events occur. You configure event notifications for your S3 bucket to notify OCI when there is an update or new data available to load. The following steps explain the procedure to be followed in Amazon S3 console to enable event notifications.

1. Log in to Amazon S3 Management console and sign in as an IAM (Amazon Identity and Access Management) user.
2. In the **Buckets** list, select the name of the bucket i.e. testbucket. This is the bucket that you had created in Step 1: Create your object store bucket in Amazon S3.
3. Select **Properties** icon.
4. Navigate to the **Event Notifications** section and select **Create event notification** icon.
5. In the **General configuration** section, specify the following values for event notification.
   - **Event name:** bucket-notification
   - **Prefix:** This value is to filter event notifications by prefix. It is an optional value. This is added to filter event activity.
   - **Suffix:** This value is to filter event notifications by suffix. It is an optional value. This is added to filter event activity.
6. In the **Event types** section, select one or more event types that you want to receive notifications for. If you are unsure of what event types to pick, then select the **All object create events** option.
7. In the **Destination** section, select **SNS Topic** as the event notification destination.

> **Note:**
>
> Before you can publish event notifications, you must grant the Amazon S3 the necessary permissions to call the relevant API. This is so that it can publish notifications to a Lambda function or an SNS topic.

8. After you select SNS topic as the event notification destination, select the SNS topic i.e. notify-topic from dropdown. This is the topic you created in Step 5: Create a notifications service subscription topic.

9. Select **Save changes**.

Step 7: Create a notifications service subscription

Where: Amazon SNS console

Every Amazon SNS topic has a set of subscriptions. Once a message is published to a topic, SNS handles distributing the message to all its subscribers. The subscribers can be AWS Lambda functions, HTTP(S) endpoints, email addresses and mobile phone numbers capable of receiving SMS messages.

Amazon SNS matches the topic to a list of subscribers who have subscribed to that topic and delivers the message to each of those subscribers.

1. Log in to Amazon SNS console.
2. In the left navigation pane, select **Subscriptions**.
3. Select **Create subscription** on the subscriptions page.
4. In the **Details** section of the **Create subscription** page, specify the following values.
   - **Topic ARN:** Paste the ARN value copied from previous step (Step 5: Create a notifications service subscription topic).
   - **Protocol:** HTTPS
   - **Endpoint:** Paste the endpoint value you copied while creating the live table feed in previous step (Step 4: Create and configure a live table feed to use notifications and copy the notification URL).
5. Expand the **Subscription filter policy** section to configure a filter policy. This step is optional.
6. Expand the **Redrive policy (dead-letter queue)** section to configure a dead-letter queue for the subscription. This step is optional.
7. Select **Create subscription**.

> **Note:**
>
> HTTP(S) endpoints, email addresses, and AWS resources in other AWS accounts require confirmation of the subscription before they can receive messages.

Step 8: Confirm that the endpoint can receive notifications

Where: Database Actions: Live Feeds page

1. Return to the Database Actions Live Feeds page and find the card for the live table feed you are configuring for a notification-based feed.
2. Click the **Actions** (three vertical dots) icon on the card, and select **Show Confirmation URL**.
3. In the **Confirmation URL** dialog box, click the link to confirm the URL. This does not close this dialog box. If the link is successful, a message is displayed that confirms the subscription is active.
4. Return to the **Confirmation URL** dialog box and select the Check only when the cloud store confirmation process is complete check box, and click **OK**.

Once you finish the above steps, any new files uploaded to the bucket will automatically be loaded into the live table feed table.

For more information on how to enable and configure event notifications using the Amazon S3 console, refer Enabling and configuring event notifications using the Amazon S3 console.

Creating a Notification-Based Live Table Feed using Microsoft Azure

A notification-based Live Table Feed is an interface between Oracle Cloud Infrastructure and a third-party cloud message queuing service such as Azure Event Grid.

The following section explains the procedure to generate automatic live feed messages using Microsoft (MS) Azure as the cloud storage. When there is an update in the container and the notification conditions are met, a log message is generated and displayed in the live feed in Oracle Cloud Infrastructure.

To create a notification-based live feed with Microsoft Azure as cloud storage you must:

- Step 1: Create a resource group in Microsoft Azure
- Step 2: Create a storage account in Microsoft Azure
- Step 3: Create Access Keys
- Step 4: Create a container
- Step 5: Add cloud storage using Microsoft Azure cloud store
- Step 6: Create and configure a live table feed to use notifications and copy the notification URL
- Step 7: Enable Event Resource Provider
- Step 8: Create Event subscription
- Step 9: Confirm that the endpoint can receive notifications

> **Tip:**
>
> To complete the steps above, you will need to alternate between Microsoft Azure portal and Oracle Database Actions pages. You may find it convenient to open the Microsoft Azure portal in one browser page or tab and Database Actions in another, so it is easy to move back and forth.

Step 1: Create a resource group in Microsoft Azure

Where: Microsoft Azure Portal

Resource groups are logical containers where you can manage Azure resources like storage accounts. Resource groups are created so you can deploy, update and delete them as a group. You can create a resource group by following these steps:

1. On the Azure portal, click the **Resource groups** button.
2. Select **Add**.
3. Enter the following values:
   - **Subscription:** Select your Azure subscription, such as Microsoft Azure Enterprise.
   - **Resource group:** Enter a new resource group name, such as resource-group.
   - **Region:** Select your location, such as US west.

4. Click **Review+create**.
5. Click **Create**. It takes a few seconds to create a resource group.

Step 2: Create a storage account in Microsoft Azure

Where: Microsoft Azure Portal

An Azure storage account contains all your storage data objects like blobs, tables, disks etc. You can create a storage account inside the resource group. It provides a unique namespace for your data. To create a storage account, do the following:

1. From the left portal menu, select **Storage accounts** to display a list of your storage accounts.
2. On the Storage accounts page, click the **Create** icon.
3. On the Basic tab, provide the following information for your storage account.
   - **Subscription:** Microsoft Azure Enterprise
   - **Resource group:** resource-group
   - **Storage account name:** teststorage
   - **Region:** Select your location, such as US west.
   - **Redundancy:** Locally-redundant storage (LRS)
4. You can select **Review+create** to accept the default options and proceed to validate the account.
5. After the validation passes, you can proceed to click on **Create storage** account. In case the validation fails, the portal indicates which settings must be modified.

Step 3: Create access Keys

Where: Microsoft Azure Portal

You must grant Microsoft Azure the permissions necessary to obtain access keys on your storage locations. The access keys specific to the storage account are generated automatically after the storage account is created in the previous step. The following steps describes the procedure to create access keys.

1. In **Security+Networking** , select **Access keys**. Your account access keys appear with the complete connection string for each key.
2. Select **Show keys** to show your access keys and connection string for each key and to copy values.
3. Copy the connection string value under key1. This value will be pasted in **Azure storage account access key** text field of next step (Step 5: Add cloud storage using Microsoft Azure cloud store).
4. Copy the storage account name i.e. teststorage and paste it in **Azure storage account name** text field of next step (Step 5: Add cloud storage using Microsoft Azure cloud store).
5. Test the credentials to see if it works or not.

Step 4: Create a container

Where: Microsoft Azure Portal

A container is a location (also known as buckets in Amazon S3 and OCI) which holds Azure Blob (Binary large object) storage. Follow these steps to create a container.

1. Navigate to your new storage account in the Azure portal.
2. In the left menu for the storage account, scroll to the **Data storage** section, then select **Containers**.
3. Click the **+Container** icon.
4. Enter the name for your new container. The container name must be lowercase, must start with a letter or number, and can include only letters, numbers, and the dash character.
5. Set the level of **Public Access Level** to Private. The default level is Private.
6. Select **Create** to create the container.

Step 5: Add cloud storage using Microsoft Azure cloud store

Where: Database Actions: Manage Cloud page

1. Click the **Manage Cloud Store** button at the top of the page to go to the Manage Cloud page. For further instructions on adding source files residing in cloud storage provided by Microsoft Azure cloud storage, refer to *Create an Microsoft Azure Cloud Storage Link* topic in Managing Connections section.

> **Note:**
>
> Paste the connection string value under key 1 of previous step Step 3: Create Access Keys in **Azure storage account access key** text field of **Add Cloud Storage** page. Also paste the storage account name generated in the previous step Step 3: Create Access Keys in **Azure storage account name** text field of **Add Cloud Storage** page.

Step 6: Create and configure a live table feed to use notifications and copy the notification URL

Where: Database Actions: Live Feeds page

The Live table feed object enables data to be loaded from Microsoft Azure cloud storage with no polling delay. This object creates an integration between Oracle Cloud Interface and Microsoft Azure.

You can configure a new or an existing live table feed to use notifications:

1. Go to the Database Actions Live Feeds page, as described in Feeding Data.
2. Create or edit a live table feed object, as described in Create a Live Table Feed Object or Edit a Live Table Feed Object. Select the **Enable for Notification** option
3. Click **Create** or **Save**.
4. Click the **Actions** (three vertical dots) icon on the card for your live feed, and select **Show Confirmation URL**.
5. In the **Notification URL** dialog box, click the **Copy** icon to copy the URL to the clipboard. You may want to copy it to a temporary file, so you can retrieve it later. You will use this URL in the subsequent step, (Step 8: Create Event subscription).

Step 7: Enable Event Resource Provider

Where: Microsoft Azure Portal

If this is the first time you are using Event Grid, you must enable Event Grid resource provider.

1. Select **Subscriptions** on the left menu.
2. Select the subscription you are using for Event Grid i.e. Microsoft Azure Enterprise.
3. On the left menu, under Settings, select **Resource Providers**.
4. Search **Microsoft.EventGrid**.
5. Select **Register**.

It takes a minute for the registration to finish.

Step 8: Create Event subscription

Where: Microsoft Azure Portal

You create an Event Subscription by configuring the subscription and specifying the endpoint that will receive the notifications.

1. Select the storage account you created in Step 2: Create a storage account in Microsoft Azure.
2. Select the **Events** icon in the left navigation pane.
3. Click on **+Event Subscription**.

The **Create Event Subscription** window appears.

1. Specify the following fields in **Event Subscription** details section:
   - **Name:** Eventssub. This is the name of the Event Subscription we create.
   - **Event Schema:** Event Grid Schema
2. Specify the following fields in **Topic Details** section:
   - **Topic Type:** Storage account
   - **System Topic Name:** eventtopic.
3. Specify the following fields in **Event Types** section:
   - **Event Type:** MicrosoftStorage.BlobCreated
4. Specify the following fields in **Endpoint Details** section:
   - **Endpoint Type:** Web Hook
   - **Endpoint:** Paste the Notification URL copied in Step 6: Create and configure a live table feed to use notifications and copy the notification URL.
5. Select **Create**.

This way Microsoft Azure creates a system topic first and then the Event Subscription for the topic.

Step 9: Confirm that the endpoint can receive notifications

Where: Database Actions: Live Feeds page

1. Return to the Database Actions Live Feeds page and find the card for the live table feed you are configuring for a notification-based feed as created in Step 6: Create and configure a live table feed to use notifications and copy the notification URL.
2. Click the **Actions** (three vertical dots) icon on the card, and select **Show Confirmation URL**.

3. In the **Confirmation URL** dialog box, click the link to confirm the URL. This does not close this dialog box. If the link is successful, a message is displayed that confirms the subscription is active.

> **Note:**
>
> The **Confirmation URL** link expires after a few minutes. You must make sure that you click the link before it expires.

4. Return to the **Confirmation URL** dialog box and select the **Check only when the cloud store confirmation process is complete** check box, and click **OK**.

Once you finish the above steps, upload a new file to the Microsoft Azure container you created in Step 4: Create a container.

1. Navigate to the container you created.
2. Select the **Container** to show a list of blobs it contains.
3. Select **Upload** button to open your local repository and browse the file you need to upload as a block blob.
4. Select the **Upload** button to upload the blob.
5. You can now view the new blob listed within the container.
6. Return to the Database Actions Live Feeds page and find the card for the live table feed you are configuring for a notification-based feed.
7. Click the **Actions** (three vertical dots) icon on the card, and select **Live table feed Run Details**.

You should be able to view logs for the blob uploaded to the Live Feed table from the Microsoft Azure storage in **Live table feed Run Details** window.

For more details on how to create a topic and subscription in Azure portal, refer to Azure Event Grid Notifications.

Send Email using Live Feed Tool

You can configure the Live Feed tool to send you email notifications when certain specific live feed events occur.

Prerequisite

1. Before subscribing to notifications, make sure that you have first created a live feed.
2. **Configure the mail service**:
   For the live feed tool to send notifications, you need to configure the Simple Mail Transfer Protocol (SMTP) Settings in the Data Studio Settings.

   Refer to the Send Email on Autonomous Database chapter for more details.

> **Note:**
>
> The Live Feed tool uses the UTL_SMTP package in the back end to provide an interface. While UTL_SMTP itself does not inherently support encryption like OAuth2, which is required by some providers (e.g., Microsoft 365), using a secure connection (SSL/TLS) can encrypt the communication between your application and the SMTP server. However, UTL_SMTP does not natively support OAuth2, which is becoming increasingly common for authentication with large email providers.

The following rules apply while configuration:

- The SMTP server port must be 25 or higher. The only ports likely to be used here are 587 (most common for modern setups, typically with StartTLS), 465 (Secure Sockets Layer) or 25 (no encryption).
- The SMTP server host name must not be blank.
- The E-mail address field must not be blank.
- Typically, clients will need to provide authentication credentials to access their SMTP server. Additionally, most setups will likely use StartTLS for encryption. However, the specific requirements can vary depending on how the company or provider has configured their system.

You can configure the email notifications in the following locations:

Livefeed page

Data Load Dashboard

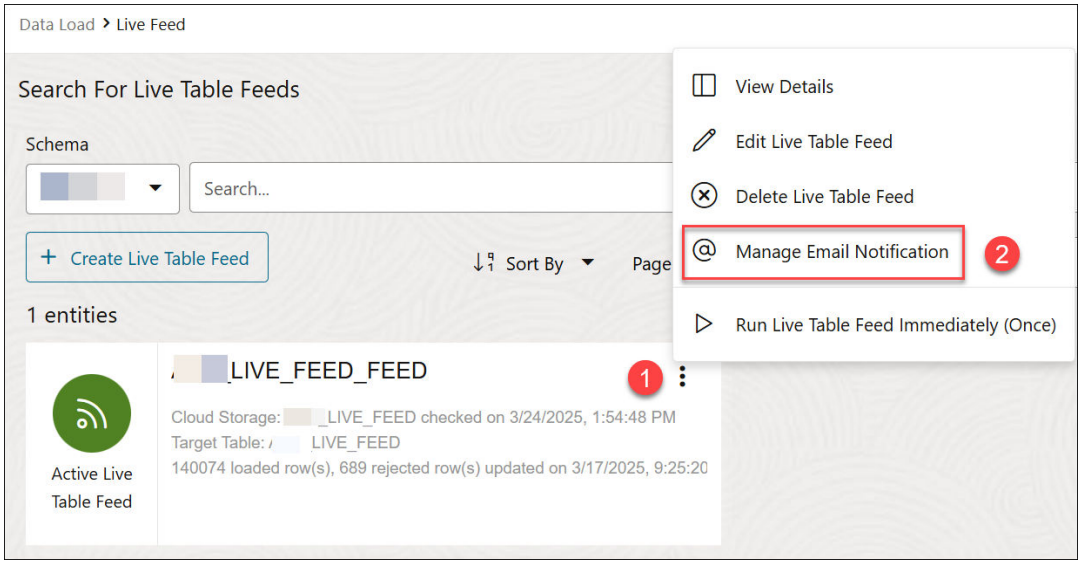
Catalog page

Configure Email Notifications on the Live Feed tool

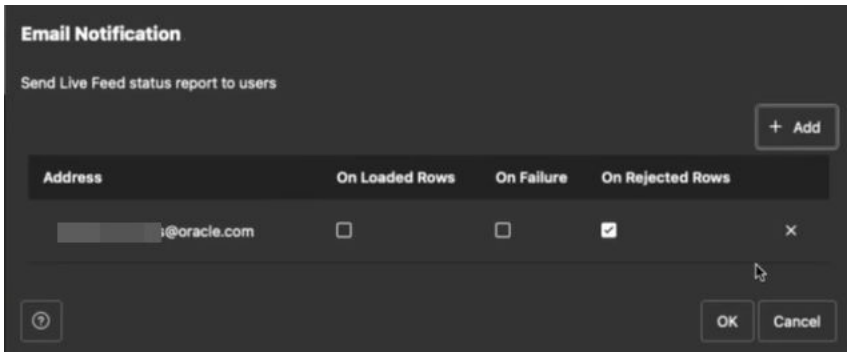
After you have created a Live Table Feed using the Live Feed tool, navigate to the Live Feed tool .

You will view the Live Table Feed card on the Live Feed tool page.

1. Click the **Actions** icon on the card and select **Manage Email Notification**.



2. On the Email Notification dialog, select **+Add** to add your email address where you will receive the notification.

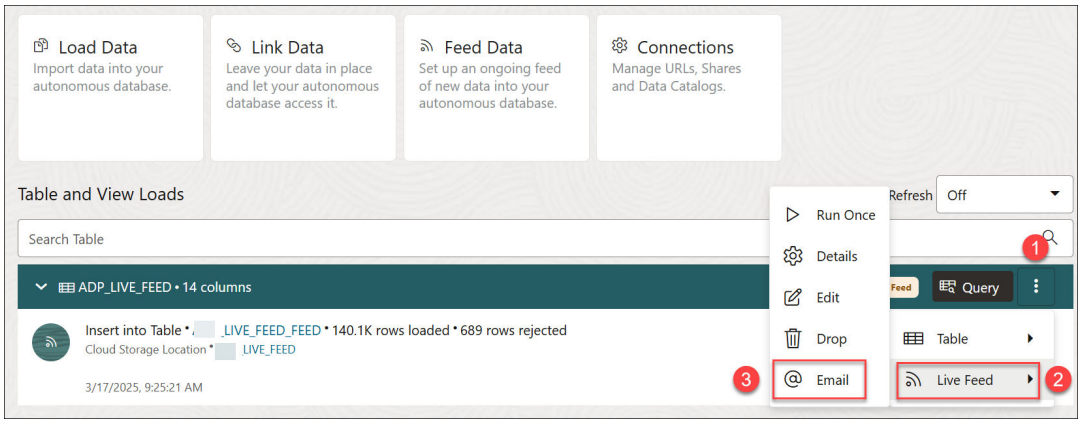


3. Enter the email address.
4. Select any of the available options to receive email notifications based on the following conditions:
   - **On Loaded Rows**: Select this option if you want to be notified in case the Live Feed successfully imports new data into your table.
   - **On Failure**: Select this option if you want to be notified in case the Live Feed fails to run due to an error. For example, if the storage location for your files is missing, or your access credentials are no longer valid.
   - **On Rejected Rows**: Select this option if you want to be notified in case the Live Feed encounters and rejects incorrect data. For instance, when there is a data type mismatch, such as attempting to import the word `Blue` into a column that must consist of only numbers.
5. Click **OK**.

Configure Email Notifications on the Data Load Dashboard

Once you have created a Live Feed, you will view the status of the feed on the Data Load Dashboard.

1. Click the **Actions** menu on the table header, click **Live Feed** and select **Email**.



2. On the Email Notification dialog, select **+Add** to add your email address where you will receive the notification.
3. Enter the email address.
4. Click **OK**.

Configure Email Notifications on the Catalog Tool

You can view the Live Feed entity on the Catalog page after you create it.

1. Select **Manage Email Notifications** the Live Feed entity.
2. **Continue with the wizard**: Progress to the next steps in the **Manage Email Notifications** wizard by following the same process you used previously.

After you have configured the email notifications, the subscriber with the email address specified in the Email Management notification dialog will now receive notifications according to the settings you configure during the live feed creation process.

DBMS_LIVE_FEED Package Reference

This chapter provides information about the packages you use with the Data Load Tool in Data Studio. The `DBMS_LIVE_FEED` topic also covers the procedures included in the `DBMS_LIVE_FEED` package.

Topics:

Summary of `DBMS_LIVE_FEED` Subprograms

- `DBMS_LIVE_FEED` Procedures
- Summary of DBMS_LIVE_FEED Subprograms
  The `DBMS_LIVE_FEED` package simplifies common operations like running a live table feed on demand, on a schedule, or as the result of a notification driven updates to a table based on new objects in object store buckets.

Summary of DBMS_LIVE_FEED Subprograms

The `DBMS_LIVE_FEED` package simplifies common operations like running a live table feed on demand, on a schedule, or as the result of a notification driven updates to a table based on new objects in object store buckets.

The following table lists the `DBMS_LIVE_FEED` subprograms and briefly describes them.

DBMS_LIVE_FEED Package Subprograms

Subprogram	Description
create_for_storage_link	This procedure creates a live feed for storage link.
create_for_directory	This procedure creates a live feed for directory.
create_for_url	This procedure creates a live feed from URL.
modify_live_feed	This procedure modifies a live feed.
drop_live_feed	This procedure drops a live feed.
load_new_content	This procedure loads new content in an existing Live Feed.
get_notification_key	This procedure retrieves Notification Key from a Live Feed.
get_confirmation_url	This procedure receives confirmation URL from a Live feed.
get_notification_state	This procedure receives notification state from a Live Feed.

Create For Storage Link Procedure

This procedure creates a live table feed from a cloud storage link. A cloud storage link is a named association between an OCI bucket URI, and a local credential name.

Syntax

```
procedure create_for_storage_link(live_feed_name     IN VARCHAR2,
                                  table_name         IN VARCHAR2,
                                  storage_link_name  IN VARCHAR2,
                                  schema             IN VARCHAR2
DEFAULT NULL,
                                  table_schema       IN VARCHAR2
DEFAULT NULL,
                                  column_list        IN CLOB
DEFAULT NULL,
                                  object_type        IN VARCHAR2
DEFAULT NULL,
                                  object_format      IN VARCHAR2
DEFAULT NULL,
                                  object_filter      IN VARCHAR2
DEFAULT NULL,
                                  object_filter_type IN VARCHAR2
DEFAULT NULL,
                                  scheduled          IN BOOLEAN
DEFAULT NULL,
                                  scheduler_interval IN VARCHAR2
DEFAULT NULL,
                                  notifications      IN BOOLEAN
DEFAULT NULL,
                                  job_class          IN VARCHAR2
```

```
DEFAULT NULL,
                                                start_date          IN TIMESTAMP
WITH TIME ZONE DEFAULT NULL,
                                                end_date            IN TIMESTAMP
WITH TIME ZONE DEFAULT NULL);
```

DBMS_ LIVE_FEED Parameters

If you are accessing `DBMS_LIVE_FEED` procedures, use the following parameters.

Parameter	Description
`live_feed_name`	The name of the Live feed which you want to create.
`storage_link_name`	The name of the cloud storage link. The name of the link should follow standard Oracle naming conventions.
`table_name`	The name of the Autonomous Database target table into which the live feed will load data from Cloud Object Storage.
`Object_filter`	The name of the target table that data from the live feed will be loaded into in your Autonomous Database instance to access the Cloud Object Storage.
`object_filter_type`	The regular expression to limit the live table feed to only those files in the bucket that match the expression. . It could be a regular expression or a glob, depending on the `object_filter_type` parameter.
`scheduled`	The default is `TRUE`. If `TRUE`, you can set up a schedule for running the live table feed object.
`notifications`	The default is `FALSE`. When `TRUE`, the Live Feed tool allows your livefeed to be notified by object storage that there are new files to be loaded.
`schema`	The name of the schema. The name of the schema of the livefeed. Only the current schema is supported at the moment.
`table_schema`	The schema of the target table.
`object_type`	The type of the object. The type of files you can load are CSV, JSON and Parquet.
`object_format`	This is one of the SQL Loader options. It accepts format received by `DBMS_CLOUD.COPY_DATA` procedure.
`scheduler_interval`	This parameter displays the interval at which the livefeed runs using the `DBMS_SCHEDULER` syntax.
`job_class`	The Scheduler job class. The values are `TPURGENT`, `TP`, `HIGH`, `MEDIUM` and `LOW`.
`start_date`	The starting date for the live feed job.
`end_date`	The ending date of the Live feed job.

> **Note:**
>
> The parameters of all the procedures described below are listed in the above table.

Example

In this example, a live feed named `MYLIVEFEED` is created on the given URL.

```
SQL> BEGIN
   dbms_live_feed.create_for_storage_link(
     live_feed_name => 'MYLIVEFEED',
     storage_link_name => 'DATATOOLS_LF_BUCKET',
     table_name => 'MALTESE',
     object_filter => 'fndcalday?.csv',
     object_filter_type => 'GLOB',
     scheduled => false,
     notifications => false);
end;
/;
PL/SQL procedure successfully completed.
```

Create For Directory Procedure

This procedure creates a live table feed from a directory.

Syntax

```
procedure create_for_directory(live_feed_name     IN VARCHAR2,
                               table_name         IN VARCHAR2,
                               directory_name     IN VARCHAR2,
                               schema             IN VARCHAR2 DEFAULT NULL,
                               table_schema       IN VARCHAR2 DEFAULT NULL,
                               column_list        IN CLOB     DEFAULT NULL,
                               object_type        IN VARCHAR2 DEFAULT NULL,
                               object_format      IN VARCHAR2 DEFAULT NULL,
                               object_filter      IN VARCHAR2 DEFAULT NULL,
                               object_filter_type IN VARCHAR2 DEFAULT NULL,
                               scheduled          IN BOOLEAN  DEFAULT NULL,
                               scheduler_interval IN VARCHAR2 DEFAULT NULL,
                               notifications      IN BOOLEAN  DEFAULT NULL,
                               job_class          IN VARCHAR2 DEFAULT NULL,
                               start_date         IN TIMESTAMP WITH TIME
ZONE DEFAULT NULL,
                               end_date           IN TIMESTAMP WITH TIME
ZONE DEFAULT NULL);
```

Create For URL Procedure

This procedure creates a live table feed from a URL.

Syntax

```
 procedure create_for_url(live_feed_name     IN VARCHAR2,
                          table_name         IN VARCHAR2,
                          url                IN VARCHAR2,
                          credential_name    IN VARCHAR2,
                          schema             IN VARCHAR2 DEFAULT NULL,
                          table_schema       IN VARCHAR2 DEFAULT NULL,
                          column_list        IN CLOB     DEFAULT NULL,
```

```
                            object_type        IN VARCHAR2 DEFAULT NULL,
                            object_format      IN VARCHAR2 DEFAULT NULL,
                            object_filter      IN VARCHAR2 DEFAULT NULL,
                            object_filter_type IN VARCHAR2 DEFAULT NULL,
                            scheduled          IN BOOLEAN  DEFAULT NULL,
                            scheduler_interval IN VARCHAR2 DEFAULT NULL,
                            notifications      IN BOOLEAN  DEFAULT NULL,
                            job_class          IN VARCHAR2 DEFAULT NULL,
                            start_date         IN TIMESTAMP WITH TIME ZONE
DEFAULT NULL,
                            end_date           IN TIMESTAMP WITH TIME ZONE
DEFAULT NULL);
```

Modify Live Feed Procedure

This procedure modifies a live table feed.

Syntax

```
procedure modify_live_feed(live_feed_name     IN VARCHAR2,
                           schema               IN VARCHAR2 DEFAULT NULL,
                           scheduled            IN BOOLEAN  DEFAULT NULL,
                           scheduler_interval IN VARCHAR2 DEFAULT NULL,
                           notifications      IN BOOLEAN  DEFAULT NULL,
                           job_class          IN VARCHAR2 DEFAULT NULL,
                           start_date         IN TIMESTAMP WITH TIME ZONE
DEFAULT NULL,
                           end_date           IN TIMESTAMP WITH TIME ZONE
DEFAULT NULL);
```

Example

In this example, a live feed named `MYLIVEFEED` is modified.

```
begin
   dbms_live_feed.modify_live_feed(
     live_feed_name => 'MYLIVEFEED',
     scheduled => true,
     scheduler_interval => 'FREQ=daily;INTERVAL=1',
     notifications => true);
end;
/
```

Drop Live Feed Procedure

This procedure drops a live table feed.

Syntax

```
procedure drop_live_feed(live_feed_name     IN VARCHAR2,
                                 schema             IN VARCHAR2 DEFAULT
NULL);
```

Example

In this example, a live feed named `MYLIVEFEED` is dropped.

```
begin
   dbms_live_feed.drop_live_feed(live_feed_name => 'MYLIVEFEED');
end;
/
```

Load New Content Procedure

This procedure loads new content in the feed.

Syntax

```
 procedure load_new_content(live_feed_name IN VARCHAR2,
                                    schema         IN VARCHAR2 DEFAULT NULL);
```

Get Notification Key Procedure

This procedure receives Live Table Feed Notification Key. It is used for a live feed with notifications enabled to construct a notification URL to provide a cloud service. Follow the instructions in the Creating a Notification-Based Live Table Feed chapter to create a notification based live feed.

Syntax

```
function get_notification_key(live_feed_name IN VARCHAR2,
                                      schema         IN VARCHAR2 DEFAULT
NULL)
    return VARCHAR2;
```

Get Confirmation URL

This procedure receives confirmation URL which verifies that the endpoint can receive notifications.

You must visit this URL to enable notifications you receive from a cloud service. It is not possible to confirm the enabling of notifications automatically due to database firewall restrictions. Click this URL to enable notifications else you will not receive any notifications.

Syntax

```
function get_confirmation_url(live_feed_name IN VARCHAR2,
                                      schema         IN VARCHAR2 DEFAULT
NULL)
    return VARCHAR2;
```

Get Notification State

This procedure receives the state of notification of the Live Feed job execution.

Syntax

```
 function get_notification_state(live_feed_name IN VARCHAR2,
                                        schema         IN VARCHAR2 DEFAULT
```

```
NULL)
    return VARCHAR2;
```

Augment Data using Table AI Assist Tool

The Data Studio suite of tools now features an AI-powered Table Assist Tool to effortlessly enhance your local data tables with the power of AI. You can use natural language prompts to augment and correct your information with this tool.

Topic

- The Table AI Assist Tool
  The **Table AI Assist Tool** enables you to add, remove, and rename columns in a table using AI assistance. It enables you to augment or correct data found in tables in your local schema using natural language AI prompts. For example, prepare data for analytics by adding columns for days of the week or reporting quarters or deriving the distance between two sets of geographical co-ordinates, and more.

The Table AI Assist Tool

The **Table AI Assist Tool** enables you to add, remove, and rename columns in a table using AI assistance. It enables you to augment or correct data found in tables in your local schema using natural language AI prompts. For example, prepare data for analytics by adding columns for days of the week or reporting quarters or deriving the distance between two sets of geographical co-ordinates, and more.

You can add columns for day, month, quarter and year using data in a date column or add a distance column from two columns with geographical co-ordinates.

The **Table AI Assist** tool enables you to prepare and enhance data for analytics, for example by deriving separate columns from a single date column for time-based reporting, or by calculating the distance between two geographical co-ordinates and adding that as a new column. The Table AI Assist tool works by using your configured AI profile to generate SQL to prepare your data. The output of the tool can then be a new view on top of your source table, or a new table.

Prerequisite:

See Data Studio Settings to set and configure your AI profile.

The following topics describe the **Table AI Assist** page and how to use it:

- Prepare Data Using Table AI Assist Tool
  Data Augmentation helps you to add data to your source tables from your target tables. The Table AI Assist allows you to add, rename, and remove columns in a table using natural language prompts.
- View Table Entity Details
  You can view the entity details of the selected entity by just clicking the **Actions** icon on the entity. This displays a list of icons that lets you view the entity's details. The icons vary from entity to entity.

Prepare Data Using Table AI Assist Tool

Data Augmentation helps you to add data to your source tables from your target tables. The Table AI Assist allows you to add, rename, and remove columns in a table using natural language prompts.

The **Table AI Assist** tool uses an AI service to generate SQL. You can choose to use the generated SQL to alter the table, create a view, or create a SELECT statement. It enables you to try suggested prompts generated by Artificial Intelligence (AI) and receive a query that is run where you can preview the data.

What is a Recipe?

A recipe is a set of steps. Each step lets you add, replace, remove and rename columns without making any changes to your base table. You add as many steps as you want to a recipe. You can save a recipe to your schema or export it as JSON to a file. You can return back to your saved recipe and keep working on it.

Different modes of Recipe
The following are different modes of a recipe:

- Create or Replace View
- Create a new table
- Alter table

Once you are satisfied with a recipe you can use it to Create or replace a View, Create a new Table or Alter the Table.

You can choose how to use the recipe at any time. The choices are different modes of recipe.

Let's say you have a table with four columns, and then add two new columns that are derived from the original data. The **Table AI Assist** tool can then:

- Alter the table to add the two new columns.
- Create either a new table, or a view on the table, with any of the four original columns and the two new columns.

Consider a table with `EMP` details. The table contains columns such as `HIREDATE`, `ENAME`, `JOB`, `MGR`, `SAL`, `COMM` and `DEPTNO`.

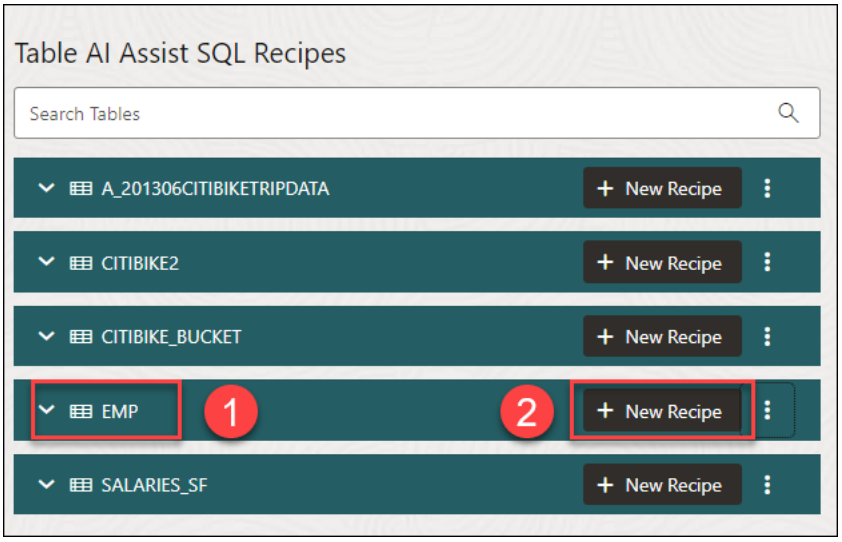
Using Table AI Assist, you can use natural prompts such as:

- add month in format year-month from `order_date`
- add year from `order_date`

The **Table AI Assist** tool generates SQL expressions for each prompt and gathers all changes into a recipe.

The following article demonstrates augmenting data using the **Table AI Assist** menu.

1. **Select the table to be modified.**
   a. From Launchpad, select Table AI Assist menu from the Data Studio suite of tools.
   b. On the **Table AI Assist** page, you can either search for the existing table from the search field at the top of the page or view it on the list of tables. In this example, let us select the ***EMP*** table.
2. **Create a new Recipe.**
   a. Click **+ New Recipe** to create a new recipe.



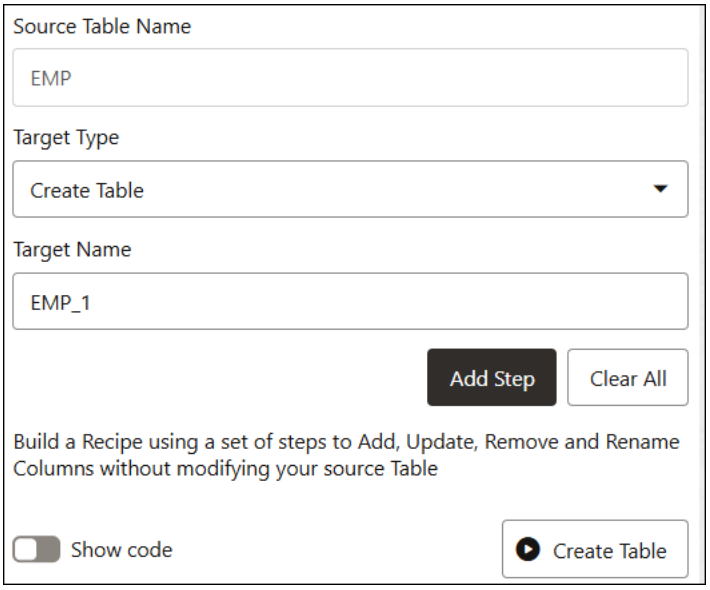
You can build a Recipe using a set of steps to Add, Update, Remove and Rename Columns without modifying your source Table.

b. **Select the Target Type and Target Name for viewing the changes you make to the source data.**
On the right panel of the Table AI Assist page, specify the following field values:

- **Source Table Name**: This is the name of the source table. In this example, it is *EMP*.
- **Target Type**: You can select one of the following options:
  - **Create View**: This option creates a view using steps in the recipe and updates the column values using column expressions in the view. Views contain only selected columns and expressions. Views do not store data.
  - **Create Table**: This option creates a new table using steps in the recipe. Data is stored in all columns, including columns added by using the recipe. It stores the updated column values.
  - **Alter Table**: This option adds, removes, and renames columns in the source table and stores the updated column values using steps in the recipe. The new columns are added as virtual columns using expressions (they do not store data).
  - **Query**: This option creates a query using steps in the recipe.

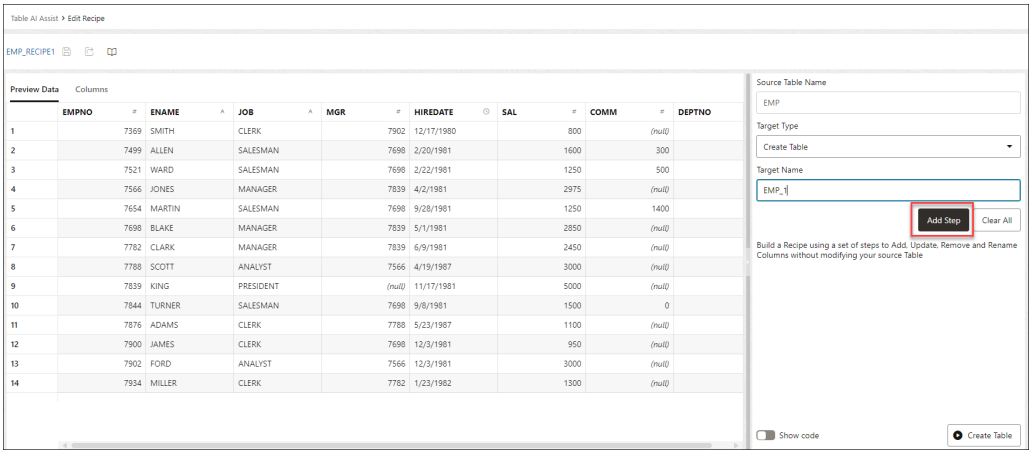
  In this example, you will **Create Table**.
- **Target Name**: Enter the name of the newly created table or view. It is a table or a view depends on the selection in **Target Type** drop-down field.

  In this example, you will view *EMP_1*.



3. **Add Steps to build the recipe that is comprised of various column actions.**

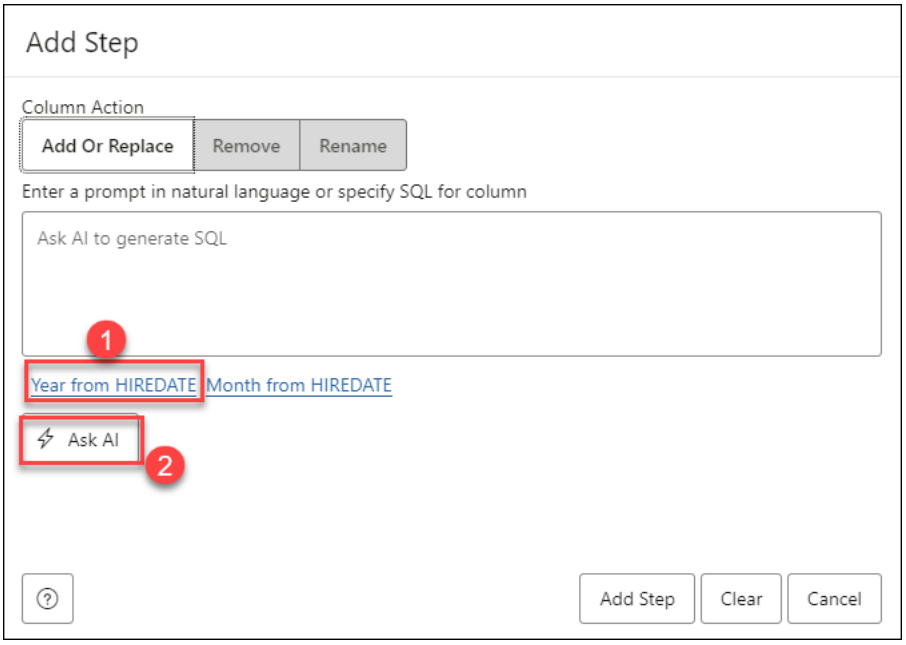
   a. Click **Add Step** to update the source tabular data.



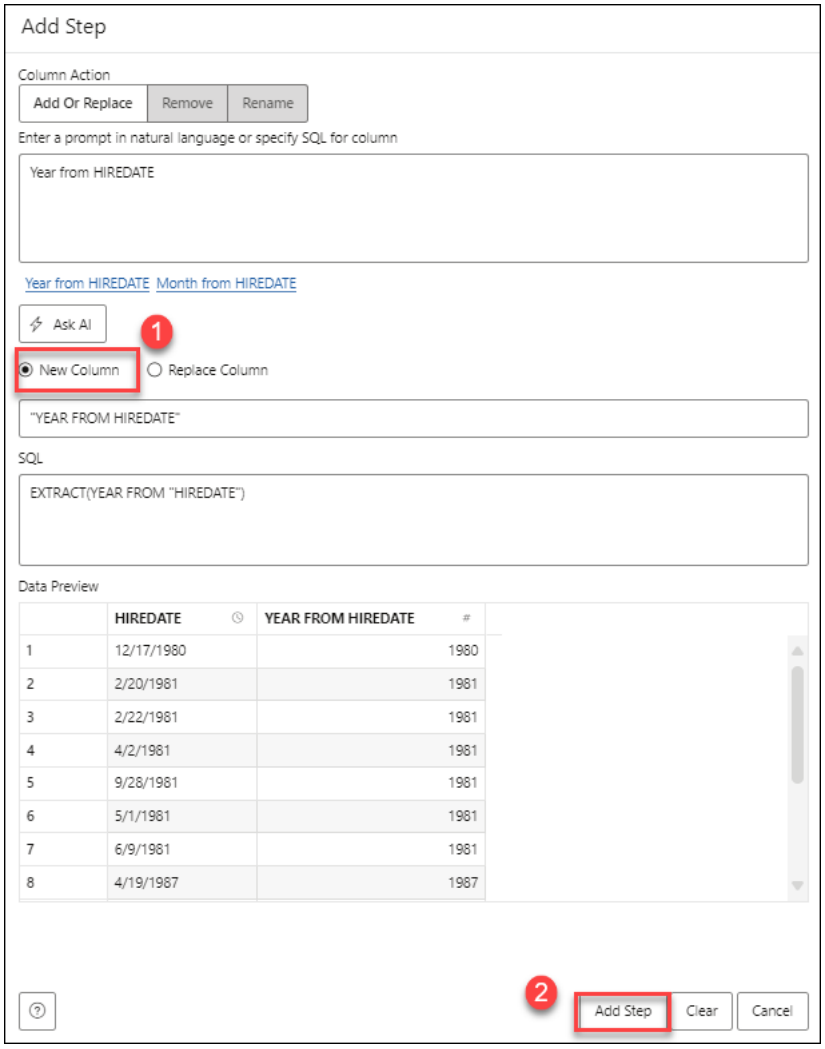
   b. Specify the following fields in the Add Step dialog box:
   On the Column Action, you can select one of the following available actions:

   - **Add or Replace**: You can add or replace the data in the source column based on the SQL code the tool generates in the prompt field. See Add a Column or Replace a Column. For example, if you replace `SALES_DATE` with `YEAR`, (assume `SALES_DATE` as a date data type) with prompt such, you will get a `SALE_DATE` column with years.
   - **Remove**: You can remove a column from the source data. See Remove a Column.
   - **Rename**: You can rename a source column. See Rename a Column.

   The fields generated when you add steps varies on the column action you select.



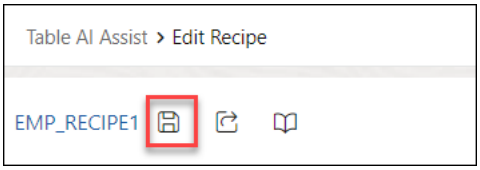
- Select **Year from `HIREDATE`**.
- Click **Ask AI**.
- Click **New Column** to create a new column and store the **Year from `HIREDATE`** value.



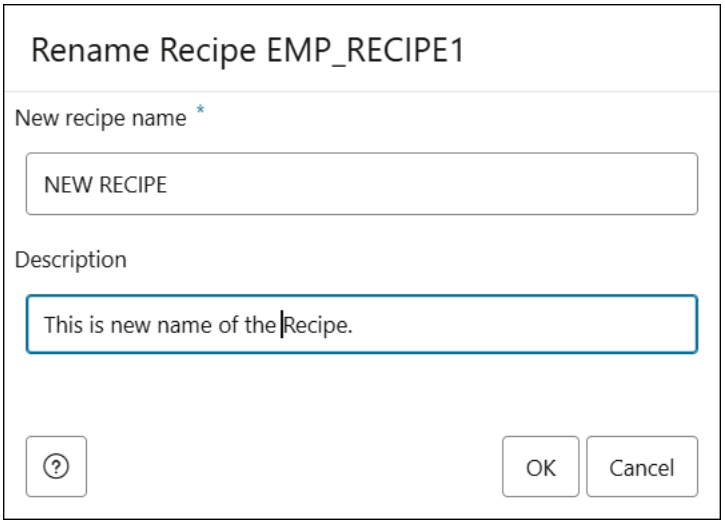
- Click **Add Step** to create a recipe.

4. **Save Recipe**.

   a. Click **Save Recipe** on the left panel of the Table AI assist page to save this recipe.



   b. Click the default Recipe name besides the **Save Recipe** to rename the Recipe. On clicking the Recipe name, you will view a **Rename Recipe** dialog.



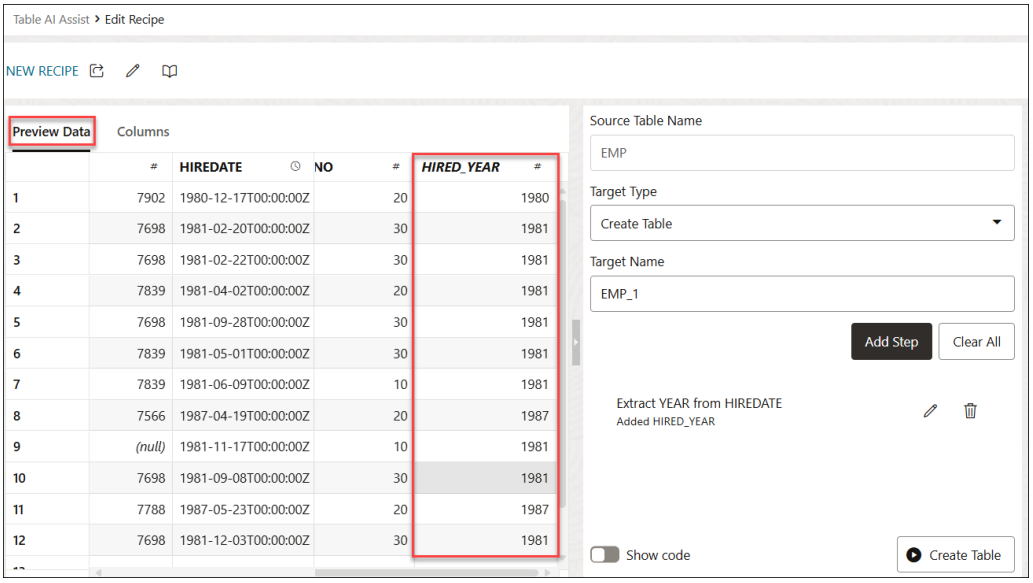
Specify the following field values on the **Rename Recipe** dialog:

- **New recipe name**: Enter the recipe's name that you created. For example, *NEW RECIPE*.
- **Description**: Enter any description. This field is optional.

5. **View the results**

   a. You can view the newly added recipe in the main recipe panel on the Table AI Assist page.

   b. You can progressively keep adding steps, modifying the source data, storing it on the target table or target view and view the updated data in the **Preview Data** tab.



You can view the newly added column header in *italics*.

6. **Create a new target table, a new target view or Alter the target table with modified column source data.**

   a. After you have modified the source data, you can create a new table, create a new view or Alter table based on the option you select in the **Target Type** drop-down.

b. Select **Create Table**, **Create View** or **Alter Table** on the side panel of the Table AI Assist page based on the **Target Type** you select.
The tool creates the target data upon confirmation.

c. You will receive a successful notification message that says the creation of the target type is successful.

After you have created the recipe, you can view it under the **Table AI Assist SQL Recipes** on the Table AI Assist page.

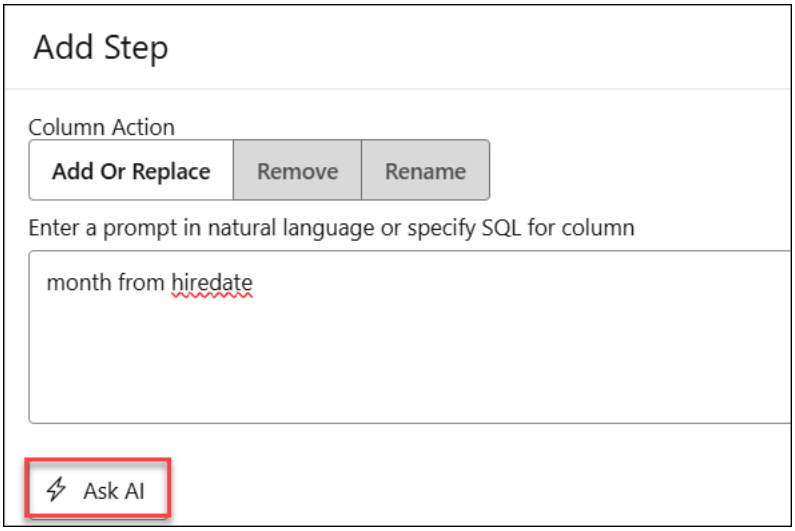
Edit Recipe

You can edit an existing recipe from the Table AI Assist tool.

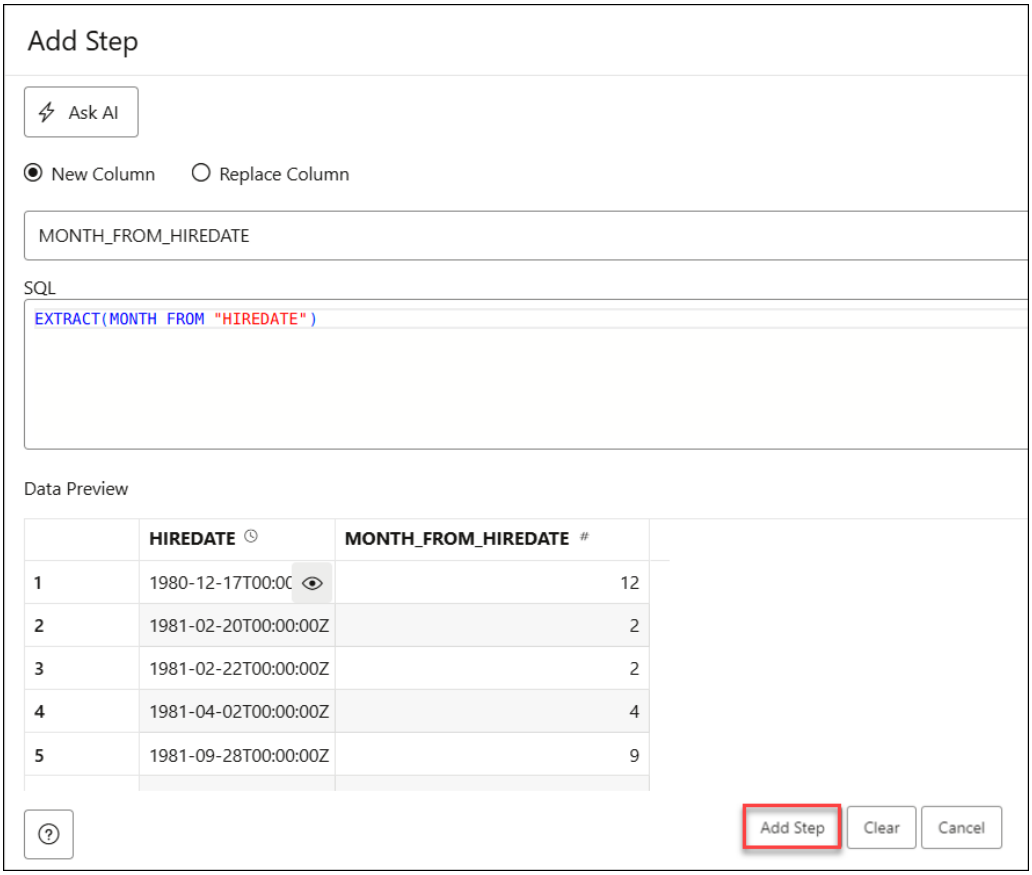
1. On the Table AI Assist page, click on the **Edit Recipe** (pencil) icon besides the recipe name. This opens the Edit Recipe page.
2. You can add steps or delete the existing steps to edit the existing recipe.
Let's perform another operation of adding a new column that consists of "month" extracted from the hire date.
   - Click **Add Step**. In the **Enter a prompt in natural language or specify SQL for column** text field, provide some natural language prompt of what you want to do such as

     ```
     month from hiredate
     ```

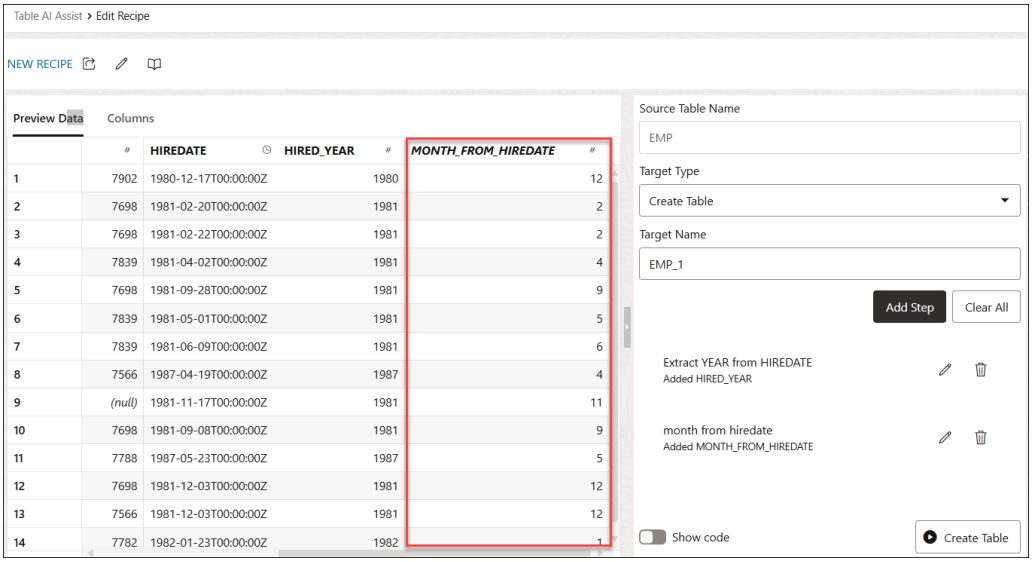
   - Click **Ask AI**.



   - The SQL expression that was used to create this column is displayed in the **SQL** field.



Scroll to the right of the table to view the newly added ***MONTH_FROM_HIREDATE*** column is displayed in the **Preview Data** section.



3. Click **Save Recipe** to save the updated recipe.

Export Recipe

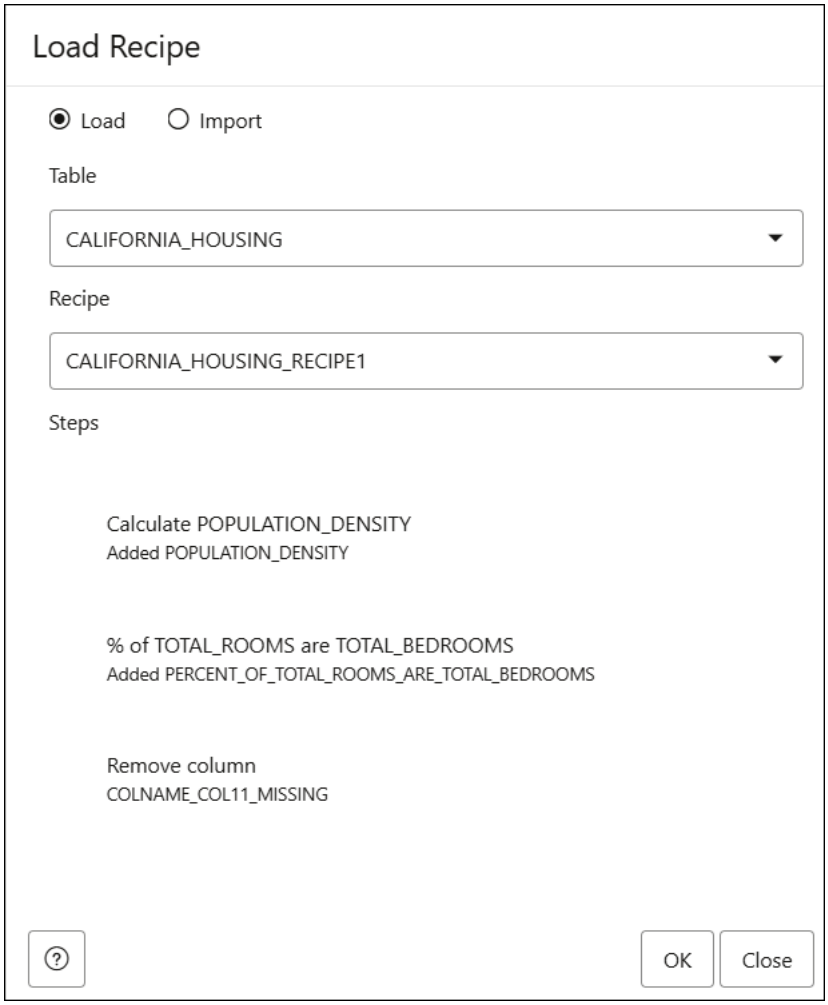
You can export the specific recipe.

1. On the **Edit Recipe** page.
2. Click the Export Recipe icon besides the Recipe name.
3. Select the **Export** icon on the to generate the JSON file for download.

Load or Import Recipe

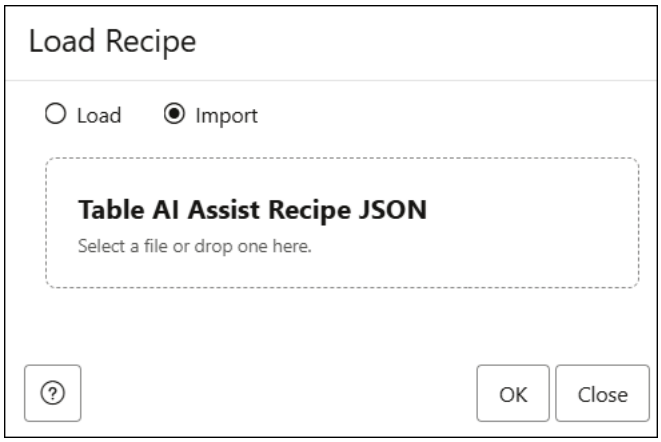
You can load or import recipe.

1. Click **Load Recipe** on the Edit Recipe page. Selecting Load Recipe opens a **Load Recipe** Dialog which displays the following options:
   - Load
     - Select **Load** to load an existing recipe.

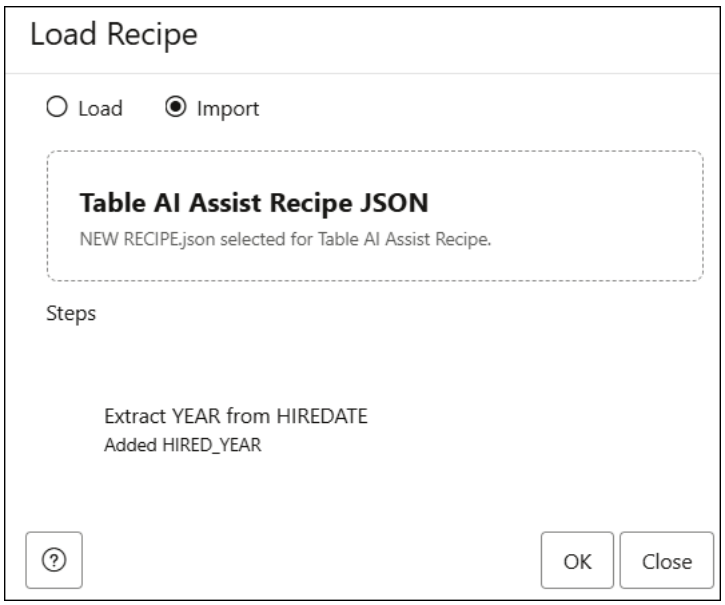


     - Specify the following field values:
       * **Table**: Select the table from the drop-down field you want to load the recipe from.
       * **Recipe**: Select the recipe you want to load from the drop-down.
       * Under Steps, you can view the Steps involved in creating the recipe.
       * Click **OK** to load the recipe.

- Import
  - Select **Import** to import a recipe file you download. This file is in JSON format.



  - Drop an existing Table AI Assist Recipe file (in JSON format) you have exported.



  - Click **OK** to import the recipe. Click **Cancel** to cancel the ongoing process of loading or importing.

  You can share a recipe with another user by importing it. For example, if you and another user both load the same table, and you use the Table AI Assist tool to set up a bunch of new and interesting columns, you can import and share the recipe with the other user so they can view the modified data.

- Add a Column
  You can view the following fields when you select **Add or Replace** from Column Action.
- Replace a Column
  You can view the following fields when you select **Add or Replace** from Column Action.
- Remove a Column
  You can view the following fields when you select **Remove** from the Column Action field.

- Rename a Column
  You can view the following fields when you select **Rename** from the Column Action field.
- View Statistical Details about the Table you Load after you Create a Recipe
  After you have created a recipe using the Table AI Assist tool, you can view additional information such as statistics about the table you create. You can generate statistics that measure the data distribution and storage characteristics of the table.

Add a Column

You can view the following fields when you select **Add or Replace** from Column Action.

1. **Enter a prompt in natural language or specify SQL for column**: Select any one of the prompts suggested by the tool.

   > **Note:**
   >
   > The Data Studio suite of tools uses a Large Language Model (LLM) when using prompts to generate SQL.

2. Click **Ask AI**.
   The **Ask AI** button sends the prompt, along with the table metadata (not data) to the Large Language Model (LLM) that is selected in your Data Studio settings. The LLM then attempts to generate SQL that will perform the request of the prompt.
3. Select **New Column** to create a new column that consists of the updated new column. After you select **New Column**, you can view a text box where you specify the name of the new column.

   The SQL textbox displays the SQL code the tool generates to display the action performed by the column action you selected in the previous step.
4. Under **Data Preview**, you can view the existing old column and the generated new column.

After you have decided the column action, click **Add Step** to add this step in the main recipe panel.

Clicking **Add Step** closes the Add Step dialog and opens the Table AI assist page.

Replace a Column

You can view the following fields when you select **Add or Replace** from Column Action.

1. **Enter a prompt in natural language or specify SQL for column**: Select any one of the prompts suggested by the tool.
2. Click **Ask AI**.
   The **Ask AI** button to send the prompt with table metadata and asks the Large Language Model (LLM) to generate a SQL.
3. Click **Replace Column**.
   Select a column you want to replace from the **Column Name to Replace** drop-down field.

   The SQL textbox displays the equivalent SQL code the tool generates to display the action performed by the column action you selected in the previous step.
4. Under **Data Preview**, you can view the existing old column and the updated new column.
5. Save the newly created recipe.

> **Note:**
>
> By default, the **Table AI Assist** tool does not replace the original source table. It only affects the new view or table that is generated by the **Table AI Assist** tool.

After you have decided the column action, click **Add Step** to add this step in the main recipe panel.

Clicking **Add Step** closes the Add Step dialog and opens the Table AI assist page.

Remove a Column

You can view the following fields when you select **Remove** from the Column Action field.

Select the column which you want to remove from the **Column Name** to remove drop-down.

After you have decided the column action, click **Add Step** to add this step in the main recipe panel.

Clicking **Add Step** closes the Add Step dialog and opens the Table AI assist page.

> **Note:**
>
> By default, the **Table AI Assist** tool does not replace the original source table. It only affects the new view or table that is generated by the **Table AI Assist** tool.

Rename a Column

You can view the following fields when you select **Rename** from the Column Action field.

1. Select the column name you want to change the name of from the **Source Column Name** drop-down.
2. Enter the new name for the column you want to use.
3. After you have decided the column action, click **Add Step** to add this step in the main recipe panel.
4. Clicking **Add Step** closes the Add Step dialog and opens the Table AI assist page.

View Statistical Details about the Table you Load after you Create a Recipe

After you have created a recipe using the Table AI Assist tool, you can view additional information such as statistics about the table you create. You can generate statistics that measure the data distribution and storage characteristics of the table.

1. Select the table whose details you want to view.
   a. On the Table AI Assist page, you can either search for the existing table from the search field at the top of the page or view it on the list of tables.
   b. Click **Edit Recipe** on any of the recipes that were generated from the table.
2. View information about the tabular data
   a. You can preview the tabular data.
   b. Click **Columns** tab to view the following information about the table:

- **Table Size (Bytes)**: Specifies the size of the table in bytes.
- **Number of Rows**: Displays the number of rows in the table.
- **Number of Columns**: Displays the number of columns in the table.
- **Compressed Size**: Displays the size of compressed table in bytes.
- **Statistics Gathered on**: Displays date and time of last statistics gathered.

c. Click **Gather Statistics** to collect statistics periodically for tables where the statistics become stale over time because of changing data volumes or changes in column values.

3. **Column Statistics**:

- Number of distinct values (NDV) in column
- Number of nulls in column
- Data distribution (histogram)

The above information is displayed in the statistics details of the table.

View Table Entity Details

You can view the entity details of the selected entity by just clicking the **Actions** icon on the entity. This displays a list of icons that lets you view the entity's details. The icons vary from entity to entity.

On the **Table AI Assist** page, click the Action icon on the table header to view the details of a table. You can view the following options:

- **View Details**: Select **View Details** to view details about the table. Refer to Viewing Entity Details for more detailed information.
- **Gather Statistics**: Select this option to display new statistics after you modify the table's structure. Refer to Gathering Statistics for Tables for more detailed information.
- **Add to Share**: Select this option to share the table. Refer to Provide Share for more detailed information.
- **Register to Cloud Link**: Select this option to register the table for remote access for a selected audience you define. See Registering Cloud Links to Access Data section for more information.
- **Publish to Marketplace**: Select this option to publish the selected table to Data Marketplace. Refer to Publish Listings to Marketplace for more detailed information.
- **Create Analytic View**: Select this option to create an Analytic View from the selected associated table. See Creating Analytic Views.
- **Query**: Select this option to analyze the selected table. See The Data Analysis Tool for more detailed information.
- **Export Data to Cloud**: Select Export Data to Cloud to export data to a cloud object store.
- **Edit**: Select **Edit** to edit the properties of the table. Refer to Editing Tables for more information.
- **Drop**: Select **Drop** to delete the table.

Transform Data Using Data Studio

Data Transforms is an easy-to-use graphical user interface that you can use to design graphical data transformations for data integration with Autonomous Database.

- The Data Transforms Page
  Data Transforms allows you to design data transformations in the form of data loads, data flows, and workflows, without requiring you to write any code.

The Data Transforms Page

Data Transforms allows you to design data transformations in the form of data loads, data flows, and workflows, without requiring you to write any code.

Data loads provide a convenient way of loading data into Autonomous Database, data flows define how the data is moved and transformed between different systems, while the workflows define the sequence in which the data flows are executed.

After your data flows and workflows are ready, you can execute the mappings immediately or schedule to execute them at a later time. Oracle Data Transforms run-time agent orchestrates the execution of jobs. On execution, Oracle Data Transforms generates the code for you.

You can launch Data Transforms in any of the following ways:

- **Oracle Cloud Marketplace:** Create a Data Transforms instance from Oracle Cloud Marketplace. Data Transforms is available as a separate listing on Marketplace called Data Integrator: Web Edition.
- **Database Actions Data Studio:** Navigate to Database Actions Data Studio Page, and click **Data Transforms** in the Database Actions page.
  If you have already registered a Data Transforms instance from OCI Marketplace with the Autonomous Database, the **Data Transforms** card on the Database Actions page will continue to take you to your Marketplace instance.

> **Note:**
>
> If you do not see the Data Transforms card then your database user is missing the required DATA_TRANSFORM_USER role.

Access to the standard set of Data Transforms features may depend on where you launch Data Transforms from. In this documentation, certain topics could include any of the following badges to indicate features that may or may not be available for use:

- APPLIES TO: ● Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.
- APPLIES TO: ● Data Transforms instance that is registered with Autonomous Database.
- APPLIES TO: ● Data Transforms that is part of the suite of data tools built into Oracle Autonomous Database.

Topics

- Access Oracle Data Transforms From Data Studio
- Data Transforms Notes
- Work with Connections
  Connections help you to connect Data Transforms to various technologies reachable from your OCI network.
- Create and Run Data Loads
  A data load allows you to load multiple data entities from a source connection to a target connection.

- Work with Data Entities
  A Data Entity is a tabular representation of a data structure.
- Work with Projects
  A project is the top-level container, which can include multiple folders to organize your data flows or work flows into logical groups.
- About Data Flows
  A data flow defines how the data is moved and transformed between different systems.
- Schedule Data Flows or Workflows
  You can schedule workflows and data flows to run at specified time interval.
- Monitor Status of Data Loads, Data Flows, and Workflows
  When you run a data load, data flow, or workflow Oracle Data Transforms runs jobs in the background to complete the request. You can view the status of the job in the panel on the bottom right of the Data Load Details, the Data Flow Editor, and the Workflow Editor page.
- Introduction to Workflows
  A workflow is made up of multiple flows organized in a sequence in which they must be executed.
- Create and Use Variables
  A variable is an object that stores a single value, which can be a string, a number or a date.
- Use Machine Learning (ML) Models
  Data Transforms supports the use of ML Model in a data flow. Learn how to create and use Machine Learning (ML) Models in data flows.
- Create and Manage Jobs
  When you execute a data load, data flow or workflow Oracle Data Transforms creates a job to complete the process in the background.
- Export and Import Objects
  You can move Data Transforms objects between environments. You require an Object Storage connection to store the exported file. You then need to use the same Object Storage connection to import the objects.
- Reference
  Get additional information about using Oracle Data Tranforms on an Autonomous Database.

Access Oracle Data Transforms From Data Studio

APPLIES TO: ✔ Data Transforms that is part of the suite of data tools built into Oracle Autonomous Database.

Data Transforms combines all the elements of data integration - data movement, data synchronization, data quality, and data management to ensure that information is timely, accurate, and consistent across complex systems.

To use the Data Transforms tool you must access Database Actions as the ADMIN user or have the DATA_TRANSFORM_USER role assigned. See Manage User Profiles with Autonomous Database for information on granting roles.

To access the Data Transforms tool:

1. Login as an ADMIN user or a user with the DATA_TRANSFORM_USER role assigned.
2. Click **Data Transforms** in the Database Actions page, or click the Selector icon and select **Data Transforms** from the Data Studio menu in the navigation pane.

> **Note:**
>
> Data Transforms is also available as a separate listing on OCI Marketplace. If you have already registered a Data Transforms instance from OCI Marketplace with the Autonomous Database, the Data Transforms card on the Database Actions page will continue to take you to your Marketplace instance. If you wish to use the embedded Data Transforms, then you must unregister the Marketplace instance. See Unregister the ODI Instance from Autonomous Database.

3. When you login to the Data Transforms tool for the first time, you need to provide the database user credentials to sign in.
   It takes approximately 1-2 minutes for the Data Transforms service to start. After the service starts, the Data Transforms home page opens.
4. From the left pane of the Home page, click the **Connections** tab to view the newly created Autonomous Database connection.
   4. Click the **Actions** icon next to the connection and select **Edit**.
5. In the **Update Connection** page, enter the database username and password to use the connection. Click **Update** to save the changes.

Data Transforms Notes

APPLIES TO: Data Transforms that is part of the suite of data tools built into Oracle Autonomous Database.

Notes for using Data Transforms.

- **Enable access to private data sources from Autonomous Database**: If your Autonomous Database is configured to use a Private Endpoint, then you can only access private data sources from clients in the same Virtual Cloud Network (VCN). See Enable Access to Private Data Sources from Autonomous Database for more information.

  Data Transforms jobs stuck in the Running status: If there are any Data Transforms jobs that are stuck in the Running status for an interminably long time, either stop the job immediately or delete the job and then rerun it. You may want to do this to avoid unwanted usage of resources for your tenancy. If the issue persists, file a service request at Oracle Cloud Support or contact your support representative.

Work with Connections

Connections help you to connect Data Transforms to various technologies reachable from your OCI network.

This section describes the generic steps to create a connection. The displayed connection detail options may vary depending on the selected connection type.

Apart from the connection types listed in Supported Connection Types you can create custom connectors, which you can use to connect Data Transforms to any JDBC supported data sources. See Create Custom Connectors.

To create a new connection:

1. From the left pane of the Home page, click the **Connections** tab.
   Connections page appears.
2. Click **Create Connection**.
   The **Create Connection** page slides in.

3. Do one of the following:
   - In the **Select Type** field, enter the name or part of the name of the connection type.
   - Select the type of connection that you wish to create.
     - Databases - Allows you to configure any connection type for supported database types.
     - Applications - Allows you to configure any connection type for supported applications.
     - Services - Allows you to configure any connection type for supported services.
4. After selecting the required connection type, click **Next**.
   The **Connection Name** field is pre-populated with a default name. You can edit this value.
5. For **Connection Details**, provide the connection details for the selected type such as:
   - Connection -
     - JDBC URL - The URL to connect to the data server. For example:
       `jdbc:weblogic:sqlserver://hostname:port[;property=value[;...]]`

       For connectors that use an autoREST driver that provides the model files along with the driver, specify the servername and other properties required to connect to that datasource. For example:

       `jdbc:weblogic:autorest://servername;[property=value[;...]]`
     - User - The user name, if required, for connecting to the server.
     - Password - The password for connecting to server.
   - Advanced Options
     - Array Fetch Size - When reading large volumes of data from a data server, Oracle Data Transforms fetches successive batches of records. This value is the number of rows (records read) requested by Oracle Data Transforms on each communication with the data server.
     - Batch Update Size - When writing large volumes of data into a data server, Oracle Data Transforms pushes successive batches of records. This value is the number of rows (records written) in a single Oracle Data Transforms INSERT command.

       > **Note:**
       >
       > Set Batch Update Size to 1000 or less for loading tables with BLOB data type columns.

     - Degree of Parallelism for Target - This value indicates the number of threads allowed for a loading task. The default value is 1. The maximum number of threads allowed is 99.

       > **Note:**
       >
       > Connection details are specific and the above options vary based on the selected connection type. For the default connection that is created during provisioning, only the User and Password fields are editable. All the other fields are disabled.

6. After providing all the required connection details, click **Test Connection** to test the connection.
   If the test connection fails, do one of the following:
   - Check whether the Autonomous Database from where you are accessing Data Transforms is configured to use a private endpoint.
     See Enable Access to Private Data Sources from Autonomous Database for more information.
   - If a connection between Autonomous Databases on a Private Endpoint fails with a "Mismatch with server cert DN" error, check whether you have specified the `ssl_server_dn_match=yes` property in the JDBC URL. See Troubleshoot Mismatch with Server Cert DN Error for more information.
7. Click **Create**.
   The new connection is created.

The newly created connections are displayed in the **Connections** page.

Click the **Actions** icon (⋮) next to the selected connection to perform the following operations:

- Select **Edit** to edit the provided connection details.
- Select **Test Connection** to test the created connection.
- Click **Export** to export the connection. See Export Objects.
- Select **Delete Schema** to delete schemas.
- Select **Delete Connection** to delete the created connection.

You can also search for the required connection to know its details based on the following filters:

- **Name** of the connection.
- **Technology** associated with the created connection.

Topics

- Supported Connection Types
  This topic lists the connection types that are supported for connecting to Data Transforms.
- Create Custom Connectors
- Create a Data Transforms Connection for Remote Data Load
  You can connect to an existing Data Transforms instance and run a data load remotely.
- Create a Delta Share Connection
  Databricks Delta Share is an open protocol for secure data sharing. Oracle Data Transforms integrates with Delta Share to load data to Oracle Autonomous Database. You can use the Delta Share connection to load data from Databricks or Oracle Data Share.
- Create an Oracle Business Intelligence Cloud Connector Connection
  Oracle Business Intelligence Cloud Connector (BICC) allows you to extract business data from a data source and load it into Autonomous Database.
- Create and use an Oracle Cloud Infrastructure Generative AI Connection
  Oracle Cloud Infrastructure (OCI) Generative AI enables organizations to automate text summarization and dynamic content generation. Data Transforms integrates with OCI Generative AI to support the use of embedding vectors in a data flow.

- Create an Oracle Financials Cloud Connection
  You can fetch real time transactional data from Oracle Financials Cloud REST endpoints, import the data entities into Data Transforms, and use them as a source in a data flow.
- Create and Use an Oracle NetSuite Connection
  You can use the Oracle NetSuite JDBC Driver or OAuth 2.0 authentication to connect to the Oracle NetSuite application. For Oracle NetSuite connections, Data Transforms allows you to load pre-built dataflows and workflows that you can run to transfer data from NetSuite to your target schema.
- Create an Oracle Object Storage Connection
  You can use Data Transforms to upload data from Oracle Object Storage to Autonomous Database.
- Create a REST Server Connection

Supported Connection Types

This topic lists the connection types that are supported for connecting to Data Transforms.

Government cloud regions are required to support FIPS compliance standards for data protection. Data Transforms in Government realms is FIPS 140-2 Level 1 compliant. Make sure the connections sources are FIPS compliant to ensure secure communication between the servers.

> **Note:**
>
> APPLIES TO: Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.
>
> - For the connectors that require driver installation, you need to copy the jar files to the `/u01/oracle/transforms_home/userlibs` directory before you add the connection.
> - Apart from the connection types listed here you can create custom connectors, which you can use to connect Data Transforms to any JDBC supported data sources. See Create Custom Connectors.

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Aha!	Application	Yes	Yes	Yes	No	**OAuth2**: `jdbc:weblogic:autorest:servername=https://{yourdomain}.aha.io; AuthenticationMethod=OAuth2; clientid=**********;clientsecret=**********; authuri=https://{yourdomain}.aha.io/oauth/authorize; tokenuri=https://{yourdomain}.aha.io/oauth/token; redirectUri=http://localhost;`	
Ahrefs	Application	Yes	Yes	Yes	No	**URLParameter**: `jdbc:weblogic:autorest:ServerName=https://apiv2.ahrefs.com;AuthenticationMethod=URLParameter;SecurityToken=xxxxxxxx;AuthParam=token;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Amazon Aurora	Database	Yes	Yes	Yes	Yes	• **userIDPassword**: `jdbc:weblogic:postgresql://servername:port;AuthenticationMethod=userIDPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **kerberos**: `jdbc:weblogic:postgresql://servername:port;AuthenticationMethod=kerberos;ServicePrincipalName=servicePrincipalName;[property=value[;...]]; User: [Ignored] Password: [Ignored]` • **EntraIDPassword**: `jdbc:weblogic:postgresql://servername:port;AuthenticationMethod=EntraIDPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **EntraIDServicePrincipal**: `jdbc:weblogic:postgresql://servername:port;AuthenticationMethod=EntraIDServicePrincipal; ServicePrincipal=your_service_principal ;Secret=your_client_secret; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Amazon EMR Hive	Database	Yes	Yes	Yes	No	• **delegationToken**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=delegationToken;DelegationToken=delegationToken;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **kerberos**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=kerberos;ServicePrincipal=servicePrincipal;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **userIdPassword**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=userIdPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **none**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=none;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Amazon Redshift	Database	Yes	Yes	Yes	Yes	`jdbc:weblogic:redshift://server_name:port_number;DatabaseName=database_name;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Apache Hive	Database	Yes	Yes	Yes	No	• **delegationToken**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=delegationToken;DelegationToken=delegationToken;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **kerberos**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=kerberos;ServicePrincipal=servicePrincipal;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **userIdPassword**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=userIdPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **none**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=none;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Base	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Apache Impala	Database	Yes	Yes	Yes	No	• **userIdPassword**: `jdbc:weblogic:impala://servername:port;DatabaseName=databaseName;AuthenticationMethod=userIdPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **kerberos**: `jdbc:weblogic:impala://servername:port;DatabaseName=databaseName;AuthenticationMethod=kerberos;ServicePrincipleName=serviceprincipalname;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **none**: `jdbc:weblogic:impala://servername:port;DatabaseName=databaseName;AuthenticationMethod=none;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Apache Spark SQL	Database	Yes	Yes	Yes	No	• **kerberos**: `jdbc:weblogic:sparksql://servername:port;DatabaseName=database;AuthenticationMethod=kerberos;ServicePrincipal=servicePrincipal; [property=value[;...]]; User: [Ignored] Password: [Ignored]` • **userIdPassword**: `jdbc:weblogic:sparksql://servername:port;DatabaseName=database;AuthenticationMethod=userIdPassword; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]`	
AWS S3	Database	Yes	Yes	Yes	No	**AWS**: `jdbc:weblogic:autorest:AuthenticationMethod=AWS;servicename=s3;accesskey=*******;secretkey=*******;`	
Azure Billing	Application	Yes	Yes	Yes	No	`jdbc:weblogic:autorest://servername; [property=value[;...]];`	
Azure Compute	Database	Yes	Yes	Yes	No	`jdbc:weblogic:autorest://servername; [property=value[;...]]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Azure Data Lake Storage	Database	Yes	Yes	Yes	No	**OAuth2**: `jdbc:weblogic:autorest: ServerName=https://{accountName}.dfs.core.windows.net;AuthenticationMethod=OAuth2;ClientID=xxxxxxxx;ClientSecret=xxxxxxxx; RefreshToken=xxxxxxxx;Scope=https://storage.azure.com/user_impersonation offline_access;TokenURI=https://login.microsoftonline.com/{tenantID}/oauth2/v2.0/token;AuthorizationURI=https://login.microsoftonline.com/{tenantID}/oauth2/v2.0/authorize;RedirectURI: http://localhost;`	
Azure Reserved VM Instances	Database	Yes	Yes	Yes	No	`jdbc:weblogic:autorest://servername; [property=value[;...]];`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Azure Resource Health	Application	Yes	Yes	Yes	No	**OAuth2**: `jdbc:weblogic:autorest:ServerName=https://management.azure.com;AuthenticationMethod=OAuth2;TenantId=xxxxxxxx;ClientID=xxxxxxxx;ClientSecret=xxxxxxxx;AccessToken=xxxxxxxx;RefreshToken=xxxxxxxx;Scope=https://management.azure.com/.default offline_access;TokenURI=https://login.microsoftonline.com/{tenantId}/oauth2/token;RedirectURI=http://localhost/;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Base	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Azure SQL Database	Database	Yes	Yes	Yes	Yes	• **activeDirectoryPassword**: `jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=activeDirectoryPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **ActiveDirectoryServicePrincipal**: `jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=ActiveDirectoryServicePrincipal;ServicePrincipal=clientID;Secret=clientSecret; [property=value[;...]]; User: [Ignored] Password: [Ignored]` • **auto**: `jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=auto; [property=value[;...]]; User: [Optional] Password: [Optional]` • **kerberos**: `jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=kerberos;ServicePrincipalName=service_principal_name;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
						[property=value[;...]]; User: [Ignored] Password: [Ignored] • **ntlm**: jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=ntlm; [property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **ntlmjava**: jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=ntlmjava; [property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **ntlm2java**: jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=ntlm2java; [property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **userIdPassword**: jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=userIdPassword;	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Base	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
						[property=value[;...]]; User: [Mandatory] Password: [Mandatory]	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Azure Synapse Analytics	Database	Yes	Yes	Yes	Yes	• **activeDirectoryPassword**: `jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=activeDirectoryPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **ActiveDirectoryServicePrincipal**: `jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=ActiveDirectoryServicePrincipal;ServicePrincipal=clientID;Secret=clientSecret; [property=value[;...]]; User: [Ignored] Password: [Ignored]` • **auto**: `jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=auto; [property=value[;...]]; User: [Optional] Password: [Optional]` • **kerberos**: `jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=kerberos;ServicePrincipalName=service_principal_name;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
						[property=value[;...]]; User: [Ignored] Password: [Ignored] • **ntlm**: jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=ntlm; [property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **ntlmjava**: jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=ntlmjava; [property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **ntlm2java**: jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=ntlm2java; [property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **userIdPassword**: jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=userIdPassword;	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
						[property=value[;...]]; User: [Mandatory] Password: [Mandatory]	
BigCommerce	Application	Yes	Yes	Yes	No	**HTTPHeader**: jdbc:weblogic:autorest:servername=https://api.bigcommerce.com/stores/{store_hash};AuthenticationMethod=HTTPHeader;securitytoken=**********;authheader=X-Auth-Token;	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Cassandra	Database	Yes	Yes	Yes	Yes	• **userIdPassword**: `jdbc:weblogic:cassandra://server:port;KeyspaceName=keyspace;AuthenticationMethod=userIdPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **kerberos**: `jdbc:weblogic:cassandra://server:port;KeyspaceName=keyspace;AuthenticationMethod=kerberos;ServicePrincipal=servicePrincipal;[property=value[;...]]; User: [Ignored] Password: [Ignored]` • **none**: `jdbc:weblogic:cassandra://server:port;KeyspaceName=keyspace;AuthenticationMethod=none;[property=value[;...]]; User: [Ignored] Password: [Ignored]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Cloudera CDH Hive	Database	Yes	Yes	Yes	No	• **delegationToken**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=delegationToken;DelegationToken=delegationToken;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **kerberos**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=kerberos;ServicePrincipal=servicePrincipal;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **userIdPassword**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=userIdPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **none**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=none;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Confluence Cloud	Database	Yes	Yes	Yes	No	**Basic**: `jdbc:weblogic:autorest:servername=https://<your-domain.atlassian.net>;authenticationmethod=Basic;`	
Data Transforms	Service	Yes	Yes	Yes	No	-	For instructions on connecting to an existing Data Transforms instance, see Create a Data Transforms Connection for Remote Data Load.

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
DataStax	Application	Yes	Yes	Yes	Yes	• **userIdPassword**: `jdbc:weblogic:cassandra://server:port;KeyspaceName=keyspace;AuthenticationMethod=userIdPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **kerberos**: `jdbc:weblogic:cassandra://server:port;KeyspaceName=keyspace;AuthenticationMethod=kerberos;ServicePrincipal=servicePrincipal;[property=value[;...]]; User: [Ignored] Password: [Ignored]` • **none**: `jdbc:weblogic:cassandra://server:port;KeyspaceName=keyspace;AuthenticationMethod=none;[property=value[;...]]; User: [Ignored] Password: [Ignored]`	`jdbc:weblogic:cassandra://server:port;KeyspaceName=keyspace;[property=value[;...]]` where `KeyspaceName` specifies the default name of the Cassandra keyspace to which the driver connects.

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data base	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Delta Share	Application	Yes	Yes	Yes	No	-	For instructions on creating a connection using Delta Share, see Create a Delta Share Connection

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
DocuSign	Database	Yes	Yes	Yes	No	• **userIdPassword**`jdbc:weblogic:cassandra://server:port;KeyspaceName=keyspace;AuthenticationMethod=userIdPassword; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **kerberos**: `jdbc:weblogic:cassandra://server:port;KeyspaceName=keyspace;AuthenticationMethod=kerberos;ServicePrincipal=servicePrincipal; [property=value[;...]]; User: [Ignored] Password: [Ignored]` • **none**: `jdbc:weblogic:cassandra://server:port;KeyspaceName=keyspace;AuthenticationMethod=none; [property=value[;...]]; User: [Ignored] Password: [Ignored]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
eBay	Application	Yes	Yes	Yes	No	**OAuth2**: `jdbc:weblogic:autorest:ServerName=https://api.eBay.com;AuthenticationMethod=OAuth2;ClientID=NameHere-APIUserA-BCD-1234e5f6g-h7i8j890;ClientSecret=xxxxxxxx;AccessToken=xxxxxxxx;RefreshToken=xxxxxxxx;Scope=https://api.eBay.com/oauth/api_scopeTokenURI=https://api.eBay.com/identity/v1/oauth2/tokenRedirectURI=http://localhost/;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
EnterpriseDB	Database	Yes	Yes	Yes	Yes	• **userIDPassword**: `jdbc:weblogic:postgresql://servername:port;AuthenticationMethod=userIDPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **kerberos**: `jdbc:weblogic:postgresql://servername:port;AuthenticationMethod=kerberos;ServicePrincipalName=servicePrincipalName;[property=value[;...]]; User: [Ignored] Password: [Ignored]` • **EntraIDPassword**: `jdbc:weblogic:postgresql://servername:port;AuthenticationMethod=EntraIDPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **EntraIDServicePrincipal**: `jdbc:weblogic:postgresql://servername:port;AuthenticationMethod=EntraIDServicePrincipal; ServicePrincipal=your_service_principal ;Secret=your_client_secret; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
FinancialForce	Application	Yes	Yes	Yes	Yes	• **userIDPassword**: `jdbc:weblogic:sforce://servername;AuthenticationMethod=userIDPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **oauth2.0**: `jdbc:weblogic:sforce://servername;AccessToken=your_access_token;AuthenticationMethod=oauth2.0; [property=value[;...]]; User: [Ignored] Password: [Ignored]`	
FourSquare	Application	Yes	Yes	Yes	No	• **URLParameter**: `jdbc:weblogic:autorest:servername=https://api.foursquare.com;AuthenticationMethod=URLParameter;securitytoken=**********; authparam=oauth_token;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data base	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Generic Rest	Application	Yes	Yes	Yes	No	-	For information about connecting to any REST service endpoint to create a connection, see Create a REST Server Connection.

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data base	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Generic Rest Config	Application	Yes	No	No	No	**OAuth2**: `jdbc:weblogic:autorest: ServerName=https://{accountName}.dfs.core.windows.net;AuthenticationMethod=OAuth2;ClientID=xxxxxxxx;ClientSecret=xxxxxxxx; RefreshToken=xxxxxxxx;Scope=https://storage.azure.com/user_impersonation offline_access;TokenURI=https://login.microsoftonline.com/{tenantID}/oauth2/v2.0/token;AuthorizationURI=https://login.microsoftonline.com/{tenantID}/oauth2/v2.0/authorize;RedirectURI: http://localhost;`	For information about connecting to any REST service endpoint to create a connection, see Create a REST Server Connection.
GitHub	Application	Yes	Yes	Yes	No	**OAuth2**: `jdbc:weblogic:autorest:clientId=123a4b567c8901234567d8901e234fg5;clientSecret=**********;authUri=https://github.com/login/oauth/authorize;AuthenticationMethod=OAuth2; tokenUri=https://github.com/login/oauth/access_token;redirectUri=http://localhost;authenticationMethod=OAuth2;refreshtoken=**********;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Google Ads	Application	Yes	No	No	Depends on the driver	`jdbc:weblogic:autorest://servername;[property=value[;...]];`	Requires driver installation

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Google AdSense	Application	Yes	Yes	Yes	No	**OAuth2**: `jdbc:weblogic:autorest:servername=https://adsense.googleapis.com/v2;clientId=123456789012-abc123def456ghi789jkl012mno345pq.apps.googleusercontent.com;clientSecret=**********;authUri=https://accounts.google.com/o/oauth2/auth;tokenUri=https://oauth2.googleapis.com/token;redirectUri=http://localhost;AuthenticationMethod=OAuth2;scope=https://www.googleapis.com/auth/adsense;refreshtoken=**********;`	
Google Analytics	Application	Yes	Yes	Yes	No	• **OAuth2**: `jdbc:weblogic:googleanalytics4:AddTables='{myTableDefinitionString}';ClientID=clientID;ClientSecret=client_secret;AuthenticationMethod=OAuth2;RefreshToken=refresh_token;Scope=scope; [property=value[;...]]; User: [Ignored] Password: [Ignored]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Google BigQuery	Database	Yes	Yes	Yes	No	• **oauth2**: `jdbc:weblogic:googlebigquery:Project=project;Dataset=dataset;AccessToken=accesstoken;RefreshToken=refreshtoken;ClientID=clientid;ClientSecret=clientsecret;AuthenticationMethod=oauth2; [property=value[;...]]; User: [Ignored] Password: [Ignored]` • **serviceaccount**: `jdbc:weblogic:googlebigquery:Project=project;Dataset=dataset;ServiceAccountEmail=serviceAccountEmail;ServiceAccountPrivateKey=serviceAccountPrivateKey.json; AuthenticationMethod=serviceaccount; [property=value[;...]]; User: [Ignored] Password: [Ignored]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Google Calendar	Application	Yes	Yes	Yes	No	**OAuth2**: jdbc:weblogic:autorest:servername=https://www.googleapis.com/calendar/v3; clientId=123456789012-abc123def456ghi789jkl012mno345pq.apps.googleusercontent.com; clientSecret=**********; authUri=https://accounts.google.com/o/oauth2/auth; tokenUri=https://accounts.google.com/o/oauth2/token; redirectUri=http://localhost; AuthenticationMethod=OAuth2; scope=https://www.googleapis.com/auth/calendar https://www.googleapis.com/auth/calendar.readonly https://www.googleapis.com/auth/calendar.eventshttps://www.googleapis.com/auth/calendar.events.readonly https://www.googleapis.com/auth/calendar.settings.readonly https://www.googleapis.com/auth/calendar.addons.execute; authuri=https://accounts.google.com/o/oauth2/auth;	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
						tokenuri=https://accounts.google.com/o/oauth2/token; refreshtoken=**********;	
Google Campaign Manager	Application	Yes	Yes	Yes	No	jdbc:weblogic:autorest://servername; [property=value[;...]];	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Google Contacts	Application	Yes	Yes	Yes	No	**OAuth2**: jdbc:weblogic:autorest:servername=people.googleapis.com/v1; clientId=123456789012-abc123def456ghi789jkl012mno345pq.apps.googleusercontent.com; clientSecret=**********; authUri=https://accounts.google.com/o/oauth2/auth; tokenUri=https://accounts.google.com/o/oauth2/token; redirectUri=http://localhost; AuthenticationMethod=OAuth2; scope=https://www.googleapis.com/auth/contacts.other.readonly https://www.googleapis.com/auth/contactshttps://www.googleapis.com/auth/contacts.readonly https://www.googleapis.com/auth/directory.readonlyhttps://www.googleapis.com/auth/profile.agerange.read https://www.googleapis.com/auth/profile.emails.readhttps://www.googleapis.com/auth/profile.language.read https://	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
						www.googleapis.com/auth/ user.addresses.readhttps:// www.googleapis.com/auth/user.birthday.read https://www.googleapis.com/auth/ user.emails.readhttps://www.googleapis.com/ auth/user.gender.read https:// www.googleapis.com/auth/ user.organization.readhttps:// www.googleapis.com/auth/ user.phonenumbers.read https:// www.googleapis.com/auth/ userinfo.emailhttps://www.googleapis.com/ auth/userinfo.profile; refreshtoken: **********;	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Google Drive	Database	Yes	Yes	Yes	No	**OAuth2**: `jdbc:weblogic:autorest:servername=https://www.googleapis.com/drive/v3;clientId=123456789012-abc123def456ghi789jkl.apps.googleusercontent.com;clientSecret=**********;authUri=https://accounts.google.com/o/oauth2/auth;tokenUri=https://accounts.google.com/o/oauth2/token;redirectUri=http://localhost;AuthenticationMethod=OAuth2;scope=https://www.googleapis.com/auth/drive;refreshtoken=**********;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Base	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Google Search Ads 360	Application	Yes	Yes	Yes	No	**OAuth2**: `jdbc:weblogic:autorest:servername=https://www.googleapis.com/doubleclicksearch/; clientId: 372185870949-v1f1h94mvnad4hkajr4tonlia51vtj2n.apps.googleusercontent.com; clientSecret: **********; authUri: https://accounts.google.com/o/oauth2/auth; tokenUri: https://oauth2.googleapis.com/token; redirectUri: http://localhost;AuthenticationMethod:OAuth2; scope: https://www.googleapis.com/auth/doubleclicksearch; refreshtoken: **********;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Greenplum	Database	Yes	Yes	Yes	No	• **userIdPassword**: `jdbc:weblogic:greenplum://servername:port;KeyspaceName=keyspace; AuthenticationMethod=userIdPassword; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **kerberos**: `jdbc:weblogic:greenplum://servername:port;DatabaseName=databaseName;AuthenticationMethod=kerberos;ServicePrincipal=servicePrincipal; [property=value[;...]]; User: [Ignored] Password: [Ignored]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Hortonworks Hive	Database	Yes	Yes	Yes	No	• **delegationToken**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=delegationToken;DelegationToken=delegationToken; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **kerberos**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=kerberos;ServicePrincipal=servicePrincipal; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **userIdPassword**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=userIdPassword; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **none**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=none; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
HubSpot	Application	Yes	Yes	Yes	No	**OAuth2**: `jdbc:weblogic:hubspot:AuthenticationMethod=OAuth2;ClientID=<CLIENT_ID>;ClientSecret=<CLIENT_SECRET>;AccessToken=<ACCESS_TOKEN>;`	
Hypersonic SQL	Database	Yes	Yes	Yes	Yes	`jdbc:hsqldb:.;ifexists=true[;<property>=<value>...]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
IBM BigInsights	Database	Yes	Yes	Yes	No	• **delegationToken**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=delegationToken;DelegationToken=delegationToken; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **kerberos**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=kerberos;ServicePrincipal=servicePrincipal; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **userIdPassword**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=userIdPassword; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **none**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=none; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Base	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
IBM DB2 Hosted	Database	Yes	Yes	Yes	Yes	• **cleartext**: `jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=cleartext;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **client**: `jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=client;[property=value[;...]]; User: [Ignored] Password: [Ignored]` • **encryptedPassword**: `jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=encryptedPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **encryptedPasswordAES**: `jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=encryptedPasswordAES;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **encryptedUIDPassword**: `jdbc:weblogic:db2://`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
						servername:port;DatabaseName=database;AuthenticationMethod=encryptedUIDPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **encryptedUIDPasswordAES**: jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=encryptedUIDPasswordAES;[property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **kerberos**: jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=kerberos;ServicePrincipal=servicePrincipal; [property=value[;...]]; User: [Ignored] Password: [Ignored] • **pluginSecurity**: jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=pluginSecurity;PluginName=pluginName; [property=value[;...]]; User: [Ignored] Password: [Ignored]	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
IBM DB2 UDB	Database	Yes	Yes	Yes	Yes	• **cleartext**: `jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=cleartext;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **client**: `jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=client;[property=value[;...]]; User: [Ignored] Password: [Ignored]` • **encryptedPassword**: `jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=encryptedPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **encryptedPasswordAES**: `jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=encryptedPasswordAES;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **encryptedUIDPassword**: `jdbc:weblogic:db2://`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
						servername:port;DatabaseName=database;AuthenticationMethod=encryptedUIDPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **encryptedUIDPasswordAES**: jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=encryptedUIDPasswordAES;[property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **kerberos**: jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=kerberos;ServicePrincipal=servicePrincipal; [property=value[;...]]; User: [Ignored] Password: [Ignored] • **pluginSecurity**: jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=pluginSecurity;PluginName=pluginName; [property=value[;...]]; User: [Ignored] Password: [Ignored]	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
IBM DB2 Warehouse	Database	Yes	Yes	Yes	Yes	• **cleartext**: `jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=cleartext;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **client**: `jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=client;[property=value[;...]]; User: [Ignored] Password: [Ignored]` • **encryptedPassword**: `jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=encryptedPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **encryptedPasswordAES**: `jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=encryptedPasswordAES;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **encryptedUIDPassword**: `jdbc:weblogic:db2://`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
						servername:port;DatabaseName=database;AuthenticationMethod=encryptedUIDPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **encryptedUIDPasswordAES**: jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=encryptedUIDPasswordAES;[property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **kerberos**: jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=kerberos;ServicePrincipal=servicePrincipal; [property=value[;...]]; User: [Ignored] Password: [Ignored] • **pluginSecurity**: jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=pluginSecurity;PluginName=pluginName; [property=value[;...]]; User: [Ignored] Password: [Ignored]	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Base	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
IBM DB2/400	Database	Yes	Yes	Yes	Yes	• **cleartext**: `jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=cleartext;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **client**: `jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=client;[property=value[;...]]; User: [Ignored] Password: [Ignored]` • **encryptedPassword**: `jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=encryptedPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **encryptedPasswordAES**: `jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=encryptedPasswordAES;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **encryptedUIDPassword**: `jdbc:weblogic:db2://`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
						servername:port;DatabaseName=database;AuthenticationMethod=encryptedUIDPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **encryptedUIDPasswordAES**: jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=encryptedUIDPasswordAES;[property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **kerberos**: jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=kerberos;ServicePrincipal=servicePrincipal; [property=value[;...]]; User: [Ignored] Password: [Ignored] • **pluginSecurity**: jdbc:weblogic:db2://servername:port;DatabaseName=database;AuthenticationMethod=pluginSecurity;PluginName=pluginName; [property=value[;...]]; User: [Ignored] Password: [Ignored]	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Base	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Informix	Database	Yes	Yes	Yes	No	• jdbc:weblogic:informix://servername:port;InformixServer=informixserver;Database=databasename;[property=value[;...]];User: [Mandatory] Password: [Mandatory] • jdbc:weblogic:informix://servername:port;InformixServer=informixserver;DatabaseName=databasename;AlternateServers=(alternateserver);ConnectionRetryCount=connectionretrycount;ConnectionRetryDelay=connectionretrydelay;[property=value[;...]];User: [Mandatory] Password: [Mandatory]	
Jira	Application	Yes	Yes	Yes	No	**Basic**: `jdbc:weblogic:autorest:servername=https://your-domain.atlassian.net;AuthenticationMethod=Basic;`	
Klaviyo	Application	Yes	Yes	Yes	No	**URLParameter**: `jdbc:weblogic:autorest:servername=https://a.klaviyo.com/api;AuthenticationMethod=URLParameter;securitytoken=*****;authparam=api_key;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Magento	Application	Yes	No	No	Depends on the driver	`jdbc:weblogic:autorest://servername; [property=value[;...]];`	Requires driver installation
Mailchimp	Application	Yes	Yes	Yes	No	**BearerToken**: `jdbc:weblogic:autorest:servername=us20.api.mailchimp.com;AuthenticationMethod=BearerToken;securitytoken=*******;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
MapR Hive	Database	Yes	Yes	Yes	No	• **delegationToken**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=delegationToken;DelegationToken=delegationToken;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **kerberos**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=kerberos;ServicePrincipal=servicePrincipal;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **userIdPassword**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=userIdPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **none**: `jdbc:weblogic:hive://servername:port;DatabaseName=database;AuthenticationMethod=none;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Marketo	Application	Yes	Yes	Yes	No	• **OAuth2**: `jdbc:weblogic:autorest:servername=https://{api_id}.mktorest.com;AuthenticationMethod=OAuth2;clientid=a1234bc5-67d8-9e01-f23g-4h567ijk89l0;clientsecret=**********;tokenuri=https://{api_id}.mktorest.com/identity/oauth/token;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Microsoft Dynamics 365	Application	Yes	Yes	Yes	Yes	• **OAuth2**: `jdbc:weblogic:dynamics365:ServiceURL=serviceurl; AuthenticationMethod=OAuth2;ClientID=client_id; ClientSecret=client_secret;TokenURI=token_uri;RefreshToken=refresh_token; [property=value[;...]]; User: [Ignored] Password: [Ignored]` • **NTLM**: `jdbc:weblogic:dynamics365:ServiceURL=serviceurl; AuthenticationMethod=NTLM;NTLMDomain=nTLMDomain;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Microsoft Share Point	Application	Yes	Yes	Yes	Yes	**Basic**: `jdbc:weblogic:sharepoint:ServiceURL=service url; AuthenticationMethod=Basic;ClientID=client_id; ClientSecret=client_secret;TokenURI=token_uri;RefreshToken=refresh_token; [property=value[;...]]; User: [Ignored] Password: [Ignored]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Microsoft SQL Server	Database	Yes	Yes	Yes	Yes	• **activeDirectoryPassword**: `jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=activeDirectoryPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **ActiveDirectoryServicePrincipal**: `jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=ActiveDirectoryServicePrincipal;ServicePrincipal=clientID;Secret=clientSecret;[property=value[;...]];User: [Ignored] Password: [Ignored]` • **auto**: `jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=auto;[property=value[;...]]; User: [Optional] Password: [Optional]` • **kerberos**: `jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database;AuthenticationMethod=kerberos;ServicePrincipalName=service_principal_name;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
						[property=value[;...]]; User: [Ignored] Password: [Ignored] • **ntlm**: jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database; AuthenticationMethod=ntlm; [property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **ntlmjava**: jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database; AuthenticationMethod=ntlmjava; [property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **ntlm2java**: jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database; AuthenticationMethod=ntlm2java; [property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **userIdPassword**: jdbc:weblogic:sqlserver://hostname:port;DatabaseName=database; AuthenticationMethod=userIdPassword;	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
						[property=value[;...]]; User: [Mandatory] Password: [Mandatory]	
MongoDB	Database	Yes	Yes	Yes	Yes	• **None**: jdbc:weblogic:mongodb://host:port;AuthenticationDatabase=auth_db; AuthenticationMethod=None; [property=value[;...]]; User: [Ignored] Password: [Ignored] • **UserIDPassword**: jdbc:weblogic:mongodb://host:port;AuthenticationDatabase=auth_db; AuthenticationMethod=UserIDPassword; [property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **Plain**: jdbc:weblogic:mongodb://host:port;AuthenticationDatabase=auth_db; AuthenticationMethod=Plain; [property=value[;...]]; User: [Mandatory] Password: [Mandatory] • **Kerberos**: jdbc:weblogic:mongodb://host:port;AuthenticationDatabase=auth_db; AuthenticationMethod=Kerberos; [property=value[;...]]; User: [Ignored] Password: [Ignored]	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
MySQL	Database	Yes	Yes	Yes	Yes	`jdbc:mysql://<host>[:<port>]/[<database>]?relaxAutoCommit=true&useCursorFetch=true[&<property>=<value>...]`	Make sure that the system variable property `sql_require_primary_key` is set to OFF. Otherwise, an ADW to MySQL mapping could fail with a “Table does not exist” error.

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
MySQL Heatwave	Database	Yes	Yes	Yes	Yes	`jdbc:mysql://<host>[:<port>]/[<database>]?relaxAutoCommit=true&useCursorFetch=true[&<property>=<value>...]`	If MySQL Heatwave database is created with high availability, then write operation is not supported. Make sure that the system variable property `sql_require_primary_key` is set to OFF. Otherwise, an ADW to MySQL Heatwave mapping could fail with a “Table does

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
							not exist" error.
Netezza	Database	Yes	No	No	Depends on the driver	-	Oracle Data Transforms uses the Netezza JDBC to connect to a NCR Netezza database. This driver must be installed in your Data Transforms `userlibs` directory. You can download the Netezza JDBC driver from the IBM website.

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Oracle	Database	Yes	Yes	Yes	Yes	• `jdbc:oracle:thin:@<host>:<port>:<sid>` • `jdbc:oracle:thin:@<host>:<port/ServiceName>` • `jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)(HOST=<host>)(PORT=<port>))(CONNECT_DATA=(SERVICE_NAME=<Service Name>)))`	For Data Integrator Web Edition, write operation is supported only on Oracle cloud database targets. For details refer to the Oracle terms of use before deploying the image from OCI marketplace.
Oracle Analytics Cloud	Application	Yes	Yes	Yes	No	-	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data base	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Oracle Business Intelligence Cloud (BICC) Connector	Application	Yes	Yes	Yes	No	-	For information about creating a connection using Oracle Business Intelligence Cloud (BICC) Connector, see Create an Oracle Business Intelligence Cloud Connector Connection.
Oracle EBS	Application	Yes	Yes	Yes	Yes	`jdbc:oracle:thin:@<host>:<port>:<sid>`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data base	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Oracle Financials Cloud	Application	Yes	Yes	Yes	No	-	For information about creating a connection using Oracle Financials Cloud, see Create an Oracle Financials Cloud Connection.
Oracle Fusion ERP	Application	Yes	Yes	Yes	No	-	
Oracle Fusion Sales	Application	Yes	Yes	Yes	No	-	
Oracle Fusion Service	Application	Yes	Yes	Yes	No	-	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data base	Supported in Data Transforms built into OCI GoldenGate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Oracle GoldenGate – OCI	Service	Yes	Yes	Yes	Yes	-	
Oracle Marketing Cloud	Application	Yes	Yes	Yes	Yes	`jdbc:weblogic:eloqua:Company=<company_id>;[;property=value[;...]]`	
Oracle NetSuite	Application	Yes	Yes	Yes	No	`jdbc:ns://{Server Host}:{Server Port};ServerDataSource={Server Data Source};`	For information about creating a connection using Oracle Netsuite, see Create and Use an Oracle NetSuite Connection.

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data base	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Oracle Object Storage	Database	Yes	Yes	Yes	Yes	-	For information about creating a connection using Oracle Object Storage, see Create an Oracle Object Storage Connection.
Oracle People Soft	Application	Yes	Yes	Yes	No	`jdbc:oracle:thin:@<host>:<port>:<sid>`	
Oracle Sales Cloud	Application	Yes	Yes	Yes	No	`jdbc:weblogic:oraclesalescloud://<base_url>;WSCompressData=none`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Oracle Service Cloud	Application	Yes	Yes	Yes	No	`jdbc:weblogic:oracleservicecloud:loginHost=host;[property=value[;...]]`	When using multiple JDBC connections for Oracle Service Cloud in Oracle Data Transforms, ensure that each connection is uniquely identified to avoid conflict. To do this, add the `DatabaseName` property to the JDBC URL. For example, `jdbc:weblogic:oracleservicecloud:loginHost=h`

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
							`ost;DatabaseName=<UniqueValue>;` where *<UniqueValue>* is a unique value for each environment. For example, `DevConnection` or `ProdConnection`. This allows separate configuration files for each environment.
Oracle SIEBEL	Application	Yes	Yes	Yes	No	`jdbc:oracle:thin:@<host>:<port>:<sid>`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
PayPal	Application	Yes	Yes	Yes	No	**Basic**: `jdbc:weblogic:autorest:ServerName=https://api-m.sandbox.paypal.com;AuthenticationMethod=Basic;`	
Pivotal HD	Database	Yes	Yes	Yes	No	-	
Pivotal HDB	Database	Yes	Yes	Yes	No	• **userIdPassword**: `jdbc:weblogic:greenplum://servername:port;KeyspaceName=keyspace;AuthenticationMethod=userIdPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **kerberos**: `jdbc:weblogic:greenplum://servername:port;DatabaseName=databaseName;AuthenticationMethod=kerberos;ServicePrincipal=servicePrincipal;[property=value[;...]];User: [Ignored] Password: [Ignored]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
PostgreSQL	Database	Yes	Yes	Yes	Yes	• **userIDPassword**: `jdbc:weblogic:postgresql://servername:port;AuthenticationMethod=userIDPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **kerberos**: `jdbc:weblogic:postgresql://servername:port;AuthenticationMethod=kerberos;ServicePrincipalName=servicePrincipalName; [property=value[;...]]; User: [Ignored] Password: [Ignored]` • **EntraIDPassword**: `jdbc:weblogic:postgresql://servername:port;AuthenticationMethod=EntraIDPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **EntraIDServicePrincipal**: `jdbc:weblogic:postgresql://servername:port;AuthenticationMethod=EntraIDServicePrincipal;ServicePrincipal=servicePrincipal;Secret=secret; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Qmetry	Application	Yes	Yes	Yes	No	**URLParameter**: `jdbc:weblogic:autorest:servername=https://qtmcloud.qmetry.com;AuthenticationMethod=URLParameter;securitytoken=*****;authparam=apiKey;`	
Quick Books Online	Application	Yes	Yes	Yes	No	**OAuth2**: `jdbc:weblogic:autorest:servername=https://sandbox-quickbooks.api.intuit.com/v3/company/+++++; clientId: AB8oHGUgVpXQnqKehrhBrnnIEQ7vNa7YI4UOr6LEX1Ja6dx0CL; clientSecret: **********; authUri: https://appcenter.intuit.com/connect/oauth2; tokenUri: https://oauth.platform.intuit.com/oauth2/v1/tokens/bearer; redirectUri: http://localhost; AuthenticationMethod: OAuth2; refreshtoken: **********;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Quick Books Payments	Application	Yes	Yes	Yes	No	**OAuth2**: `jdbc:weblogic:autorest:servername=https://sandbox.api.intuit.com; clientId=ABCd1efGH2Ijkl3mN4OPQrsTu5vWxyZa6bCdEFgHijk7LMnOpQ;clientSecret=**********; authUri=https://appcenter.intuit.com/connect/oauth2; tokenUri=https://oauth.platform.intuit.com/oauth2/v1/tokens/bearer;redirectUri=http://localhost; AuthenticationMethod=OAuth2;scope=com.intuit.quickbooks.payment;refreshtoken=**********;`	
Quora Ads	Application	Yes	Yes	Yes	No	**OAuth2**: `jdbc:weblogic:autorest:ServerName=api.quora.com/ads;AuthenticationMethod=OAuth2;ClientID=xxxxxxx;ClientSecret=xxxxxxxx;AccessToken=xxxxxxx;Scope=ads_read;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Sage	Application	Yes	Yes	Yes	No	**OAuth2**: jdbc:weblogic:autorest:servername=api.accounting.sage.com;clientId:**********;clientSecret:**********;authUri:https://www.sageone.com/oauth2/auth/central?filter=apiv3.1;tokenUri:https://oauth.accounting.sage.com/token,redirectUri: http://localhost;AuthenticationMethod: OAuth2; scope: full_access; refreshtoken: **********;	
Salesforce Chatter	Application	Yes	Yes	Yes	No	**OAuth2**: jdbc:weblogic:autorest:servername=progresscom4-dev-ed.my.salesforce.com/services/data/v52.0; AuthenticationMethod=OAuth2; clientid=3MVG9SOw8KERNN08IkuAFNRo0MZbipoL4zeEhQzdT4IrNnqW7_UBvm1K7vZ4Y_OzbNxFn0.nJSfVBdqQuh4IL; clientsecret=**********; refreshtoken=**********; scope=full refresh_token; authuri=https://login.salesforce.com/services/oauth2/authorize; tokenuri=https://login.salesforce.com/services/oauth2/token;	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Salesforce.com	Application	Yes	Yes	Yes	Yes	• **userIDPassword**: `jdbc:weblogic:sforce://servername;AuthenticationMethod=userIDPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **OAuth2**: `jdbc:weblogic:sforce://servername;AccessToken=your_access_token;AuthenticationMethod=userIDPassword;[property=value[;...]]; User: [Ignored] Password: [Ignored]`	
SAP BW/4HANA	Database	Yes	Yes	Yes	No	• **Basic**: `jdbc:weblogic:s4hana:ServerName=servername;AuthenticationMethod=Basic;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **HTTPHeader**: `jdbc:weblogic:s4hana:ServerName=servername;AccessToken=your_access_token;AuthenticationMethod=HTTPHeader;AuthHeader=api_key;SecurityToken=security_token;[property=value[;...]]; User: [Ignored] Password: [Ignored]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
SAP HANA	Application	Yes	Yes	Yes	No	• **Basic**: `jdbc:weblogic:s4hana:ServerName=servername;AuthenticationMethod=Basic; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **HTTPHeader**: `jdbc:weblogic:s4hana:ServerName=servername;AccessToken=your_access_token;AuthenticationMethod=HTTPHeader;AuthHeader=api_key;SecurityToken=security_token; [property=value[;...]]; User: [Ignored] Password: [Ignored]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
SAP NetWeaver	Database	Yes	Yes	Yes	No	• **Basic**: `jdbc:weblogic:s4hana:ServerName=servername;AuthenticationMethod=Basic; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **HTTPHeader**: `jdbc:weblogic:s4hana:ServerName=servername;AccessToken=your_access_token;AuthenticationMethod=HTTPHeader;AuthHeader=api_key;SecurityToken=security_token; [property=value[;...]]; User: [Ignored] Password: [Ignored]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
SAP S/4HANA Cloud	Application	Yes	Yes	Yes	No	• **Basic**: `jdbc:weblogic:s4hana:ServerName=servername;AuthenticationMethod=Basic; [property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **HTTPHeader**: `jdbc:weblogic:s4hana:ServerName=servername;AccessToken=your_access_token;AuthenticationMethod=HTTPHeader;AuthHeader=api_key;SecurityToken=security_token; [property=value[;...]]; User: [Ignored] Password: [Ignored]`	
Semrush	Application	Yes	Yes	Yes	No	`jdbc:weblogic:autorest://servername; [property=value[;...]];`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data base	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
ServiceNow	Service	Yes	Yes	Yes	No	**OAuth2**:jdbc:weblogic:autorest:servername=instance.service-now.com/api; clientId=123a4b567c8901234567d8901e234fg5;clientSecret=**********; authUri=https://instance.service-now.com/oauth_auth.do; tokenUri=https://instance.service-now.com/oauth_token.do;redirectUri=http://localhost; AuthenticationMethod=OAuth2;refreshtoken=**********;	
Shopify	Application	Yes	Yes	Yes	No	jdbc:weblogic:autorest://servername; [property=value[;...]];	
Snowflake	Database	Yes	Yes	Yes	Yes	**Basic**: jdbc:weblogic:snowflake:AccountName=account_name;DatabaseName=database_name;Schema=schema_name;Warehouse=warehouse_name; [property=value[;...]];	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Square	Application	Yes	Yes	Yes	No	**BearerToken**: `jdbc:weblogic:autorest:servername=sandbox.api.intuit.com;AuthenticationMethod:BearerToken;securitytoken=**********;`	
Stripe	Application	Yes	Yes	Yes	No	**BearerToken**: `jdbc:weblogic:autorest:servername=https://api.stripe.com;AuthenticationMethod=BearerToken;securitytoken=**********;`	
SAP Sybase ASE	Database	Yes	Yes	Yes	Yes	`jdbc:weblogic:sybase://hostname:port[;property=value[;...]]`	
SAP Sybase IQ	Database	Yes	Yes	Yes	Yes	`jdbc:weblogic:sybase://hostname:port[;property=value[;...]]`	
Sybase As Anywhere	Database	Yes	Yes	Yes	Yes	`jdbc:weblogic:sybase://hostname:port[;property=value[;...]]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
TeamCity	Application	Yes	Yes	Yes	No	**BearerToken**: `jdbc:weblogic:autorest:ServerName=https://{your_teamcity_servername}.com;AuthenticationMethod=BearerToken;SecurityToken=Ab12cdEfG3HIJkL.MNOpQRStuvWX4yZabcD56cD7eF8GHiJ90Klm.nOpqRst2u3VWxYZabC4DEfGhIjk;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Teradata	Database	Yes	No	No	Depends on the driver	`jdbc:weblogic:teradata://<host>:<port>/<database>[;property=value[;...]]`	Data Transforms uses the Teradata JDBC Driver to connect to a Teradata Database. To use Teradata as a data source the Teradata Gateway for JDBC must be running, and this driver must be installed in your Data Transforms `userlibs` directory. You can download the JDBC driver from the

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data base	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
							Teradata website.

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Teradata 17+	Database	Yes	No	No	Depends on the driver	-	Data Transforms uses the Teradata JDBC Driver to connect to a Teradata Database. To use Teradata as a data source the Teradata Gateway for JDBC must be running, and this driver must be installed in your Data Transforms `userlibs` directory. You can download the JDBC driver from the

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
							Teradata website..
Tumblr	Application	Yes	Yes	Yes	No	`jdbc:weblogic:autorest://servername; [property=value[;...]]`	
Veeva CRM	Application	Yes	Yes	Yes	Yes	• **userIDPassword**: `jdbc:weblogic:sforce://servername;AuthenticationMethod=userIDPassword;[property=value[;...]]; User: [Mandatory] Password: [Mandatory]` • **oauth2.0**: `jdbc:weblogic:sforce://servername;AccessToken=your_access_token;AuthenticationMethod=oauth2.0; [property=value[;...]]; User: [Ignored] Password: [Ignored]`	
Volusion	Application	Yes	Yes	Yes	No	`jdbc:weblogic:autorest://servername; [property=value[;...]]`	
Wistia	Application	Yes	Yes	Yes	No	**Basic**: `jdbc:weblogic:autorest:servername=https://api.wistia.com;authenticationmethod=Basic;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
WooCommerce	Application	Yes	No	No	Depends on the driver	`jdbc:weblogic:autorest://servername;[property=value[;...]];`	Requires driver installation

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
WordPress	Application	Yes	Yes	Yes	No	**OAuth2**: `jdbc:weblogic:autorest:servername=public-api.wordpress.com/rest/v1.1;AuthenticationMethod=OAuth2;clientid=*****;clientsecret=*********;accesstoken=*********;scope=global;oauthcode=**********;authuri=https://public-api.wordpress.com/oauth2/authorize;tokenuri=https://public-api.wordpress.com/oauth2/token;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Workday	Application	Yes	No	No	Depends on the driver	`jdbc:weblogic:autorest://servername; [property=value[;...]];`	Requires driver installation
X	Application	Yes	Yes	Yes	No	-	
Xero	Application	Yes	Yes	Yes	No	`jdbc:weblogic:autorest://servername; [property=value[;...]]`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Data Warehouse	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Yelp	Application	Yes	Yes	Yes	No	**BearerToken**: `jdbc:weblogic:autorest:servername=https://api.yelp.com/v3;AuthenticationMethod=BearerToken;securitytoken=**********;`	
Zendesk	Application	Yes	Yes	Yes	No	**OAuth2**: `jdbc:weblogic:autorest:servername=https://example.zendesk.com;AuthenticationMethod=OAuth2;clientid=client_id_zendesk; clientsecret=********; scope=read write;authuri=/oauth/authorizations/new; tokenuri=/oauth/tokens;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Zoho CRM	Application	Yes	Yes	Yes	No	**OAuth2**: `jdbc:weblogic:autorest:AuthenticationMethod=OAuth2;clientid=a1234bc567d89e01f23g; clientsecret: **********;scope=ZohoCRM.settings.ALL,ZohoCRM.modules.ALL,ZohoCRM.notifications.ALL,ZohoCRM.org.READ,ZohoCRM.users.READ,ZohoCRM.settings.roles.READ,ZohoCRM.settings.profiles.READ,ZohoCRM.modules.Contacts.READ, ZohoSearch.securesearch.READ; authuri=https://accounts.zoho.com/oauth/v2/auth;tokenuri=https://accounts.zoho.com/oauth/v2/token;redirecturi=http://localhost; refreshToken=**********;`	

Name	Type	Supported in Data Integrator: Web Edition	Supported in Data Transforms built into Autonomous Database	Supported in Data Transforms built into OCI Golden Gate	Supports Write Operation	Authentication Mode: Connection URL Example	Notes
Zoom	Application	Yes	Yes	Yes	No	**OAuth2**: `jdbc:weblogic:autorest:servername=https://api.zoom.us/v2;clientId=0aBcDeF_GhI2j_KlMnOpQr;clientSecret=**********;authUri=https://zoom.us/oauth/authorize;tokenUri=https://zoom.us/oauth/token;redirectUri=http://localhost;AuthenticationMethod=OAuth2;accesstoken:**********;`	

Create Custom Connectors

APPLIES TO: Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.

The Custom Connections page of the Administration tab of Oracle Data Transforms helps you to create custom connectors that point to any JDBC supported data sources.

The custom connectors will be listed in the Create Connection page where you can use them to connect data sources to Data Transforms. See Work with Connections for more information. To create a new connector:

1. In the left pane, click **Administration**.
   A warning message appears.

2. Click **Continue**.
3. In the left pane, click **Custom Connections**.
   Custom Connections screen appears.
4. Click **Create Connection Type**.
   The Create Connection Type page appears.
5. From the **Category** drop-down select the type of connection that you wish to create whether database, application, or service.
6. Enter a name for the connection.
7. Enter the name of the JDBC Driver of the source connection. For example, `oracle.jdbc.OracleDriver`.

> **Note:**
>
> For connectors that require driver installation, you need to copy the jar files to the `/u01/oracle/transforms_home/userlibs` directory before you add the connection.

8. Click **OK**.

The newly created custom connection appears in the list and are available in the Create Connection page.

Click the **Actions** icon (⋮) next to the selected connection to perform the following operations:

- Select **Edit**, to edit the provided connection details.
- Click **Export** to export the connection. See Export Objects.
- Select **Delete**, to delete the created connection.

> **Note:**
>
> You cannot delete custom connectors that have existing connections.

Create a Data Transforms Connection for Remote Data Load

You can connect to an existing Data Transforms instance and run a data load remotely.

To create this connection, you need to specify the URL of the Data Transforms instance along with the name of the ODI rest API from where you want to run the data load.

To define a Data Transforms connection:

1. From the left pane of the Home page, click the **Connections** tab.

   Connections page appears.
2. Click **Create Connection**.

   Create Connection page slides in.
3. For **Select Type**,
   - In the **Name** field, enter the name of the newly created connection

- Select **Services** as the type of connection that you wish to create.

4. In the **Endpoint URL** textbox, enter the URL of the ODI rest API from where you want to run the data load. Enter the URL in the format `http://<host-ip-address>:<port>/odi-rest`.
5. In the **User** text box enter SUPERVISOR as the user name.
6. In the **Password** text box enter the ODI Supervisor password.
7. After providing all the required connection details, click **Test Connection** to test the established connection.
8. Click **Create**.

   The new connection is created.

The newly created connections are displayed in the **Connections** page.

Click the **Actions** icon (⋮) next to the selected connection to perform the following operations:

- Select **Edit** to edit the provided connection details.
- Select **Test Connection** to test the created connection.
- Click **Export** to export the connection. See Export Objects.
- Select **Delete Schema** to delete schemas.
- Select **Delete Connection** to delete the created connection.

You can also search for the required connection to know its details based on the following filters:

- **Name** of the Connection.
- **Technology** associated with the created Connection.

Create a Delta Share Connection

Databricks Delta Share is an open protocol for secure data sharing. Oracle Data Transforms integrates with Delta Share to load data to Oracle Autonomous Database. You can use the Delta Share connection to load data from Databricks or Oracle Data Share.

To use Databricks as a source, you need to specify the URL of the Delta Sharing server along with the bearer token that lets you access the Delta Lake share server. To use Oracle Data Share as a source, you need to specify the URL for the token end point along with a client ID and the secret key.

This topic has the following sections:

- Creating the Delta Share Connection
- Creating and Running a Delta Share Data Load

Creating the Delta Share Connection

To define a Delta Share connection:

1. From the left pane of the Home page, click the **Connections** tab.
   Connections page appears.
2. Click **Create Connection**.
   Create Connection page slides in.

3. Do one of the following:
   - In the **Select Type** field, enter the name or part of the name of the connection type.
   - Select the **Databases** tab.
4. Select **Delta Share** as the connection type.
5. Click **Next**.
6. The **Connection Name** field is pre-populated with a default name. You can edit this value.
7. In the **Share Endpoint URL** textbox, enter the URL of the Delta Sharing server. Enter the value in the `<host>:<port>/<shareEndpoint>/` format.
8. In the Connection section, do one of the following:
   - Select **Oracle Data Share** and provide the **Token Endpoint URL**, **Client ID**, and **Client Secret** for accessing the share.
     You can get this information from the Delta Share Profile JSON document that you will need to download from supplied to you by the Share Provider. (This is also where they get the Share Endpoint URL from)

     You can get this information from the Delta Share Profile JSON document that you can download from the activation link that is provided by the Data Share provider to access their share.
   - Select **Databricks** and in the **Bearer Token** text box enter the token for connecting to the Delta Sharing server.
9. If you need to use a proxy to access the Delta Share Server or Delta Share Storage configure the following settings:
   - In the **Proxy Host** textbox, enter the host name of the proxy server to be used for the connection.
   - In the **Proxy Port** textbox, enter the port number of the proxy server.
   - Select the following checkboxes depending on where the proxy is required:
     - **Use Proxy to access Delta Share Server**
     - **Use Proxy to access Delta Share Storage**
10. Click **Test Connection**, to test the established connection.
11. After providing all the required connection details, click **Create**.
    The new connection is created.

The newly created connections are displayed in the **Connections** page.

Click the **Actions** icon (⋮) next to the selected connection to perform the following operations:

- Select **Edit**, to edit the provided connection details.
- Select **Test Connection**, to test the created connection.
- Click **Export** to export the connection. See Export Objects.
- Select **Delete Schema**, to delete schemas.
- Select **Delete Connection**, to delete the created connection.

You can also search for the required Connection to know its details based on the following filters:

- **Name** of the Connection.

- **Technology** associated with the created Connection.

Creating and Running a Delta Share Data Load

To load data from Delta Share into Oracle Autonomous Database, the Oracle connection user must be an Admin user. Admin privileges are required so that the Oracle user can create and insert data into tables in another schema.

When you run the data load, Data Transforms loads the data onto a corresponding table in the target schema. The data load runs incrementally. The very first time you run a data load, Data Transforms copies all the data into new tables. For every subsequent data load run, it only uploads the changes. Any additions or deletions in the records will reflect in the target tables. Note that if there is any metadata change in the table, for example a column is added, Data Transforms creates a new table to load the data on to the target server. You could create a workflow, add the data load as a step, create a schedule to run the workflows at a predefined time interval. See Create a New Workflow.

To create and run a Delta Share data load:

1. Do one of the following:
   - On the Home page, click **Load Data**. The Create Data Load wizard appears.
     In the Create Data Load tab, enter a name if you want to replace the default value, add a description, and select a project from the drop-down.
   - On the Home page, click **Projects**, and then the required project tile. In the left pane, click **Data Loads**, and then click **Create Data Load**. The Create Data Load wizard appears.
     In the Create Data Load tab, enter a name if you want to replace the default value and add a description.
2. Click **Next**.
3. In the Source Connection tab,
   a. From the **Connection Type** drop-down, select **Delta Share**.
   b. from the **Connection** drop-down, select the required connection from which you wish to add the data entities.
   c. Select the share that you want to load tables from the **Share** drop-down. The drop-down lists all the shares for the selected connection.
   d. Click **Next**.
4. In the Target Connection tab,
   a. From the **Connection Type** drop-down, select **Oracle** as the connection type.

   > **Note:**
   > This drop-down lists only JDBC type connections.

   b. From the **Connection** drop-down, select the required connection from to you wish to load the data entities.
   c. Enter a unique name in the **Schema** textbox.
   d. Click **Save**.

   The Data Load Detail page appears listing all the tables in the selected share with their schema names.

> **Note:**
>
> For Delta Share data loads the Data Load Detail page only includes the option. You cannot apply different actions - incremental merge, incremental append, recreate, truncate, append - on the data entities before loading it to the target schema. This is to make sure that the data is consistent between the Delta Sharing server and the target schema.

5. Click to run the data load.
   A confirmation prompt appears when the data load starts successfully.

To check the status of the data load, see the Status panel on the right below the Target Schema details. For details about the Status panel, see Monitor Status of Data Loads, Data Flows, and Workflows. This panel shows links to the jobs that execute to run this data load. Click the link to monitor the progress on the Job Details page. For more information about jobs, see Create and Manage Jobs.

All the loaded data entities along with their details are listed in the Data Entities page. To view the statistics of the data entities, click the **Actions** icon ( ) next to the data entity, click **Preview**, and then select the **Statistics** tab. See View Statistics of Data Entities for information.

Create an Oracle Business Intelligence Cloud Connector Connection

Oracle Business Intelligence Cloud Connector (BICC) allows you to extract business data from a data source and load it into Autonomous Database.

To create an Oracle BICC connection you need to first configure external storage using the OCI Object Storage Connection tab in the BICC Console. You need to specify these connection details when you define the connection in Oracle Data Transforms.

You can use the BICC connection to choose the offerings whose data stores you want to extract. Data Transforms uses an Oracle Object Storage Data Server used by Oracle BICC to stage the extracted files, which you can then use as a source for mapping. Note that you cannot use an Oracle BICC connection as a target for mapping.

To define an Oracle BICC connection,

1. From the left pane of the Home page, click the **Connections** tab.

   Connections page appears.
2. Click **Create Connection**.

   Create Connection page slides in.
3. Do one of the following:
   - In the **Select Type** field, enter the name or part of the name of the connection type.
   - Select the **Applications** tab.
4. Select **Oracle BI Cloud Connector** as the connection type.
5. Click **Next**.
6. The **Connection Name** field is pre-populated with a default name. You can edit this value.
7. Enter the URL in the **BI Cloud Connector Service URL** textbox.

8. In the **Connection** section, enter the following details:
   - In the **User** text box enter the user name configured in the Oracle BI Cloud Connector Console.
   - In the **Password** text box enter the password configured in the Oracle BI Cloud Connector Console.
9. In the **Storage** section, enter the following details:
   - In the **External Storage BICC Name** text box enter the name of the external storage as it appears in the Oracle BI Cloud Connector Console.
   - In the **External Storage Bucket** text box specify the bucket into which extracts are uploaded. Bucket names are obtained in the OCI Console.
   - In the **External Storage Name Space** text box specify the namespace. Namespace is obtained in the OCI Console.
   - In the **External Storage Region** text box enter the OCI Object Storage region.
   - In the **External Storage User** text box enter your Oracle Cloud Infrastructure username.
   - In the **External Storage Token** text box enter the auth token.
10. Click **Test Connection** to test the established connection.
11. Click **Create**.

    The new connection is created.

The newly created connections are displayed in the **Connections** page.

Click the **Actions** icon (⋮) next to the selected connection to perform the following operations:

- Select **Edit**, to edit the provided connection details.
- Select **Test Connection**, to test the created connection.
- Click **Export** to export the connection. See Export Objects.
- Select **Delete Schema**, to delete schemas.
- Select **Delete Connection**, to delete the created connection.

You can also search for the required Connection to know its details based on the following filters:

- **Name** of the Connection.
- **Technology** associated with the created Connection.

Create and use an Oracle Cloud Infrastructure Generative AI Connection

Oracle Cloud Infrastructure (OCI) Generative AI enables organizations to automate text summarization and dynamic content generation. Data Transforms integrates with OCI Generative AI to support the use of embedding vectors in a data flow.

Before you create an OCI Generative AI connection you create an Oracle Database 23ai connection. Data Transforms will use this connection to test the OCI Generative AI connection. To create the OCI Generative AI connection you need to specify details such as the OCI URL, User OCID, Tenancy OCID, Compartment OCID, Private Key, and Fingerprint information.

See Use Text Embedding Vector in a Data Flow for information on how you will use this connection to add vector embedding in a data flow.

To define Oracle Cloud Infrastructure Generative AI Connection:

1. From the left pane of the Home page, click the **Connections** tab.

   Connections page appears.
2. Click **Create Connection**.

   Create Connection page slides in.
3. For **Select Type**,
   - In the **Name** field, enter the name of the newly created connection.
   - Select **Services** as the type of connection.
4. Select **OCI Generative AI**, and **Next**.
5. For **Connection Details**, provide the following details:
   - **OCI URL** - The endpoint URL of the OCI Generative AI service.
   - **User OCID** - The user OCID from the Oracle Cloud Infrastructure Console.
   - **Tenancy OCID** - The tenancy OCID from the Oracle Cloud Infrastructure Console.
   - **Compartment OCID** - The compartment OCID from the Oracle Cloud Infrastructure Console.
   - **Private Key** - The private key in the PEM format. Specify the path to your downloaded private key file.
   - **Fingerprint** - The fingerprint of the key that was just added.
6. After proving all the required connection details, click **Create**.

   The new connection is created.
7. Click **Test Connection**, to test the established connection.

   A pop-up appears listing the Oracle Database 23ai connections that you have configured. Select the option you want to use to test this connection.

The newly created connection is displayed in the Connections page.

Create an Oracle Financials Cloud Connection

You can fetch real time transactional data from Oracle Financials Cloud REST endpoints, import the data entities into Data Transforms, and use them as a source in a data flow.

To create an Oracle Financials Cloud connection you need to choose a temporary schema where Data Transforms can create data entities after the reverse-engineering operation.

To define an Oracle Financials Cloud connection,

1. From the left pane of the Home page, click the **Connections** tab.

   Connections page appears.
2. Click **Create Connection**.

   Create Connection page slides in.
3. Do one of the following:
   - In the **Select Type** field, enter the name or part of the name of the connection type.
   - Select the **Applications** tab.

4. Select **Oracle Financials Cloud** as the connection type.
5. Click **Next**.
6. The **Connection Name** field is pre-populated with a default name. You can edit this value.
7. In the **REST Service URL** textbox, enter the URL of the endpoint that services the REST resources.
8. In the **Proxy Host** textbox, enter the host name of the proxy server to be used for the connection.
9. In the **Proxy Port** textbox, enter the port number of the proxy server.
10. In the **User** text box enter the user name for connecting to the REST endpoint.
11. In the **Password** text box enter the password for connecting to the REST endpoint.
12. Choose a connection from the **Staging Connection** drop-down list. The list displays only existing Autonomous Database connections. To use a different connection, create the connection before you reach this page.
13. After providing all the required connection details, click **Create**.

    The new connection is created.
14. Click **Test Connection**, to test the established connection.

The newly created connections are displayed in the **Connections** page.

Click the **Actions** icon (⋮) next to the selected connection to perform the following operations:

- Select **Edit**, to edit the provided connection details.
- Select **Test Connection**, to test the created connection.
- Click **Export** to export the connection. See Export Objects.
- Select **Delete Schema**, to delete schemas.
- Select **Delete Connection**, to delete the created connection.

You can also search for the required Connection to know its details based on the following filters:

- **Name** of the Connection.
- **Technology** associated with the created Connection.

Create and Use an Oracle NetSuite Connection

You can use the Oracle NetSuite JDBC Driver or OAuth 2.0 authentication to connect to the Oracle NetSuite application. For Oracle NetSuite connections, Data Transforms allows you to load pre-built dataflows and workflows that you can run to transfer data from NetSuite to your target schema.

This topic has the following sections:

- Configuring Access Permissions Required for Building Data Warehouse
- Creating the Oracle NetSuite Connection
- Using the Build Data Warehouse Wizard
- Running the Pre-Built Workflows to Load Data into the Target Schema

Configuring Access Permissions Required for Building Data Warehouse

Before you create a NetSuite connection or use the Build Data Warehouse Wizard in Data Transforms, you need to login to NetSuite as an administrator, enable SuiteAnalytics Connect, create a custom role, and set the access permissions that are required to build the NetSuite Data Warehouse.

To configure the access permissions:

1. Log in to NetSuite as an administrator using the following URL:
   https://system.netsuite.com/pages/customerlogin.jsp
2. Check whether the SUITEANALYTICS CONNECT feature is enabled.
   Go to *Setup* > *Company* > *Enable Features*. Click the **Analytics** tab and verify that the **SuiteAnalytics Connect** box is checked.
3. Add a custom role.
   Go to *Setup* > *Users/Roles* > *User Management* > *Manage Roles*. Click **New Role**, add the required details, and click **Save** to create a custom role.
4. Assign the required permissions to the custom role.
   Go to *Setup* > *Users/Roles* > *User Management* > *Manage Roles*. Click **Customize** next to the name of the custom role for which you would like to add the SuiteAnalytics Connect permission.

   Click the **Transactions** tab under the **Permissions** tab and assign the following permissions:

   - Account Detail
   - Accounting Lists
   - Accounting Management
   - Accounts
   - Accounts Payable
   - Accounts Payable Graphing
   - Accounts Payable Register
   - Accounts Receivable
   - Accounts Receivable Graphing
   - Accounts Receivable Register
   - Adjust Inventory
   - Adjust Inventory Worksheet
   - Amortization Reports
   - Amortization Schedules
   - Audit Trail
   - Balance Sheet
   - Bank Account Registers
   - Bill Purchase Orders
   - Billing Schedules
   - Bills
   - Bin Putaway Worksheet

- Bin Transfer
- Blanket Purchase Order
- Build Assemblies
- CRM Groups
- Calendar
- Cash Sale
- Cash Sale Refund
- Charge
- Charge - Run Rules
- Charge Rule
- Check
- Classes
- Commission Feature Setup
- Commission Reports
- Commit Orders
- Commit Payroll
- Competitors
- Component Where Used
- Contacts
- Count Inventory
- Create Allocation Schedules
- Credit Card
- Credit Card Refund
- Credit Card Registers
- Credit Memo
- Credit Returns
- Currency
- Currency Revaluation
- Custom Recognition Event Type
- Customer Deposit
- Customer Payment
- Customer Refund
- Customers
- Deferred Expense Reports
- Deleted Records
- Departments
- Deposit
- Deposit Application

- Documents and Files
- Edit Forecast
- Edit Manager Forecast
- Email Template
- Employee Commission Transaction
- Employee Commission Transaction Approval
- Employee Record
- Employee Reminders
- Employees
- Enter Opening Balances
- Enter Vendor Credits
- Equity Registers
- Establish Quotas
- Estimate
- Events
- Expense Report
- Expenses
- Export Lists
- Fair Value Dimension
- Fair Value Formula
- Fair Value Price
- Financial Statements
- Find Transaction
- Fixed Asset Registers
- Fulfill Orders
- Fulfillment Request
- General Ledger
- Generate Price Lists
- Generate Statements
- Imported Employee Expenses
- Inbound Shipment
- Income
- Income Statement
- Inventory
- Inventory Status Change
- Invoice
- Invoice Approval
- Invoice Sales Orders

- Item Fulfillment
- Item Receipt
- Item Revenue Category
- Items
- Lead Snapshot/Reminders
- Locations
- Long Term Liability Registers
- Make Journal Entry
- Manage Accounting Periods
- Manage Payroll
- Mass Updates
- Memorized Transactions
- Mobile Device Access
- Net Worth
- Non Posting Registers
- Notes Tab
- Opportunity
- Other Asset Registers
- Other Current Asset Registers
- Other Current Liability Registers
- Other Lists
- Other Names
- Ownership Transfer
- Pay Bills
- Pay Sales Tax
- Pay Tax Liability
- Paycheck Journal
- Payroll Items
- Perform Search
- Phone Calls
- Post Vendor Bill Variances
- Posting Period on Transactions
- Price Books
- Price Plans
- Process Payroll
- Project Revenue Rules
- Purchase Contract
- Purchase Order

- Purchase Order Reports
- Purchases
- Reconcile
- Reconcile Reporting
- Refund Returns
- Report Customization
- Report Scheduling
- Request For Quote
- Requisition
- Resource
- Return Authorization
- Revalue Inventory Cost
- Revenue Arrangement
- Revenue Arrangement Approval
- Revenue Element
- Revenue Recognition Field Mapping
- Revenue Recognition Plan
- Revenue Recognition Reports
- Revenue Recognition Rule
- Revenue Recognition Schedules
- SOAP Web Services
- Sales
- Sales By Partner
- Sales By Promotion
- Sales Order
- Sales Order Approval
- Sales Order Fulfillment Reports
- Sales Order Reports
- Sales Order Transaction Report
- Set Up Budgets
- Set Up SOAP Web Services
- Statement Charge
- Statistical Account Registers
- Store Pickup Fulfillment
- Subscription Change Orders
- Subscription Plan
- Subscriptions
- Subsidiaries

- SuiteAnalytics Connect
- SuiteAnalytics Workbook
- Tasks
- Tax
- Track Messages
- Transaction Detail
- Transfer Funds
- Transfer Inventory
- Transfer Order
- Transfer Order Approval
- Trial Balance
- Unbilled Receivable Registers
- Unbuild Assemblies
- Units
- Vendor Bill Approval
- Vendor Payment Approval
- Vendor Request For Quote
- Vendor Return Auth. Approval
- Vendor Return Authorization
- Vendor Returns
- Vendors
- Work Calendar
- Work Order
- Work Order Close
- Work Order Completion
- Work Order Issue

5. Add the **SuiteAnalytics Connect – Read All** permission.
   Click the **Setup** tab under the **Permissions** tab, select **SuiteAnalytics Connect – Read All** from the drop-down, and click **Add**.
6. Click **Save** to apply these permissions to the custom role.
7. Assign the custom role to a user. When you create a connection to NetSuite from Data Transforms, you will need to enter the credentials of this user to connect to the data server. See Creating the Oracle NetSuite Connection for information about creating the connection.
   To assign the custom role to the user, go to *Setup* > *Users / Roles* > *Manage Users*. Click **Edit** next to the name of the user, assign the custom role, and click **Save**.
8. To verify the access permissions, log in as the user that has the custom role assigned. Go to *Analytics* > *Datasets*. Click **New Dataset**. This page will list all the tables and record types that the user has access to. Search for “transaction” table, for example, to verify whether the user has access to the transaction table.

Creating the Oracle NetSuite Connection

You can create an Oracle NetSuite connection using JDBC connectivity or OAuth 2.0 authentication.

To define an Oracle NetSuite connection:

1. From the left pane of the Home page, click the **Connections** tab.
   Connections page appears.
2. Click **Create Connection**.
   Create Connection page slides in.
3. Do one of the following:
   - In the **Select Type** field, enter the name or part of the name of the connection type.
   - Select the **Applications** tab.
4. Select **Oracle NetSuite** as the connection type.
5. Click **Next**.
6. The **Connection Name** field is pre-populated with a default name. You can edit this value.
7. To specify the connection details, do one of the following:
   - To use JDBC connectivity, specify the following details:
     - **JDBC URL** - Enter the URL of the SuiteAnalytics Connect server to be used for the connection.
     - **User** - Enter the user name for connecting to the data server.
     - In the **Password** textbox enter the password for connecting to the data server.
     - In the **Account ID** textbox, enter the account ID for connecting to the data server.
     - In the **Role ID** textbox, enter the role ID for connecting to the data server.
   - To use OAuth 2.0 authentication, click the **OAuth 2.0** switch and then specify the following details:
     - **Username** - Enter the name of the user who has role access to login to NetSuite using OAuth 2.0 connection. This is the user you have assigned the custom role to in Configuring Access Permissions Required for Building Data Warehouse.
     - **Account ID** - Enter the account ID for connecting to the data server. You can get this information by logging into the NetSuite account and viewing the SuiteAnalytics connect information.
     - **Role ID** - Enter the role ID for connecting to the data server. You can get this information by logging into the NetSuite account and viewing the SuiteAnalytics connect information.
     - **Client ID** - Enter the client ID for connecting to the data server.
       To obtain the client ID, create an Integration record in NetSuite by enabling OAuth 2.0 Client Credentials Flow. Copy and save the Client ID that is displayed when the Integration Record is successfully created.

> **Note:**
>
> NetSuite no longer supports the RSA PKCSv1.5 scheme for token signing for NetSuite OAuth 2.0 client credentials flow. Any integrations that rely on the RSA PKCSv1.5 scheme will need to be updated to use the RSA-PSS scheme. Refer to the Oracle NetSuite documentation for more information.

- **Public Certificate** and **Private Key** - Use the OpenSSL commands to generate the key pair in the required PEM format. For example,
  ```
  openssl req -new -x509 -newkey rsa:4096 -keyout private.pem -sigopt
  rsa_padding_mode:pss -sha256 -sigopt rsa_pss_saltlen:64 -out public.pem
  -nodes -days 365
  ```
  Paste the contents of `public.pem` in the Public Certificate field. Paste the contents of `private.pem` in the Private Key field.
- **Certificate ID** - Enter the Certificate ID for connecting to the data server.
  To get the certificate ID, use the NetSuite OAuth 2.0 Client Credentials (M2M) Setup to add the public certificate file (`auth-cert.pem`) to the certificate key list and copy the generated Certificate ID.

8. If the source that you want to use for mapping is a saved search, you need to also specify the following details in **Saved Search Extraction**:
   - **Application ID**: Enter the NetSuite Application ID for Data Transforms.
   - **Version**: Enter the NetSuite version number.
9. Select the checkbox in **Build Data Model** to install pre-built dataflows and workflows that you can run to extract data from NetSuite and move it to your Oracle target schema using the Build Data Warehouse wizard.
10. Click **Test Connection**, to test the established connection.
11. After providing all the required connection details, click **Create**.
    The new connection is created.

The newly created connections are displayed in the **Connections** page.

Click the **Actions** icon (⋮) next to the selected connection to perform the following operations:

- Select **Edit**, to edit the provided connection details.
- Select **Test Connection**, to test the created connection.
- Select **Build Data Warehouse**, to select the functional areas and create the NetSuite Data Warehouse in the target schema. See Using the Build Data Warehouse Wizard for more information.
- Click **Export** to export the connection. See Export Objects.
- Select **Delete Schema**, to delete schema.
- Select **Delete Connection**, to delete the created connection.

You can also search for the required Connection to know its details based on the following filters:

- **Name** of the Connection.
- **Technology** associated with the created Connection.

Using the Build Data Warehouse Wizard

Data in your NetSuite account is grouped into business or subject areas in the Analytics Warehouse. The Build Data Warehouse wizard allows you to select the areas that you want to include in the newly created Data Warehouse.

To use the Build Data Warehouse Wizard:

1. On the Home page, click the **Connections** tab. The **Connections** page appears.
2. Click the Actions icon (⋮) next to the Oracle NetSuite connection that you want to use to build the data warehouse and click **Build Data Warehouse**.
   The Build Data Warehouse wizard opens.
3. From the **Connection** drop-down list, choose the Autonomous Database connection where your target schema resides.
4. From the **Staging Schema** drop-down, all schema corresponding to the selected connection are listed in two groups:
   - Existing Schema (ones that you've imported into Oracle Data Transforms) and
   - New Database Schema (ones that you've not yet imported).

   Select the schema that you want to use from the drop-down.
5. Similarly select the **Target Schema**.
6. Click **Next**.
7. Select the **NetSuite Business Areas** that you want to use to transfer data from the NetSuite Data Warehouse to the target schema.
8. Click **Save**.
   Data Transforms starts the process to build the data warehouse. Click **Jobs** on the left pane of the Home page to monitor the progress of the process. When the job completes successfully, Data Transforms creates a Project folder that includes all the pre-built workflows and dataflows, which you can run to transfer data from the NetSuite connection to your target schema. See Running the Pre-Built Workflows to Load Data into the Target Schema for more information.

Running the Pre-Built Workflows to Load Data into the Target Schema

When the Build Data Warehouse wizard completes successfully, Data Transforms creates a project that includes all the pre-built data flows and workflows that you can run to extract data from a NetSuite connection and load it into your target schema.

To view and run the pre-built workflows:

1. Click **Projects** on the left pane of the Home page and select the newly created NetSuite project.
2. Click **Workflows** in the left pane. The following pre-built workflows are listed in the Project Details page:
   - Stage NetSuite Source to SDS
   - Extract Transaction Primary Keys
   - Load SDS to Warehouse
   - Apply Deletes
   - All Workflows

3. Click the Actions icon (⋮) next to the workflow you want to run and click **Start**. Oracle recommends that you run **All Workflows** to execute all the pre-built workflows.

   To see the status of the workflow, click **Jobs** from the left pane in the current project. When the job completes successfully, all the data from the NetSuite connection is loaded into the target schema.

Create an Oracle Object Storage Connection

You can use Data Transforms to upload data from Oracle Object Storage to Autonomous Database.

The OCI Object Storage dedicated endpoints feature allows OCI customers to securely access the storage buckets. See Object Storage Dedicated Endpoints for more information. You need to use the new URL format when you create Object Storage connections in Data Transforms. For users that already have an Object Storage connection, the existing URL is automatically updated to the new URL format.

To create an Oracle Object Storage connection you need to have an Oracle Cloud Infrastructure username and an auth token. See Getting an Auth Token for information about how to generate the auth token. You need to specify these details when you define the connection in Oracle Data Transforms.

Note the following:

- To use an Oracle Object Storage connection to import data into Data Transforms, you must use a public IP address to access the compute node. If you want to use a private IP address to access the Object Storage service, make sure that you have access to the Internet.
- The supported file format for loading data from Oracle Object Storage to Autonomous Database and vice versa is CSV.
- The supported data types are Numeric, Double, String, and Date.
- Data load is not supported for Oracle Object Storage connections.
- To create a mapping to Object Storage, the source technology must support the `DBMS_CLOUD` package. Otherwise, the mapping will fail.
- Data Transforms does not support cross-realm connectivity for Object Storage. For example, you cannot connect from Data Transforms configured in a Government Cloud region to Object Storage that is in a commercial realm.

To define an Oracle Object Storage connection,

1. From the left pane of the Home page, click the **Connections** tab.

   Connections page appears.

2. Click **Create Connection**.

   Create Connection page slides in.

3. Do one of the following:
   - In the **Select Type** field, enter the name or part of the name of the connection type.
   - Select the **Databases** tab.
4. Select **Oracle Object Storage** as the connection type.
5. Click **Next**.
6. The **Connection Name** field is pre-populated with a default name. You can edit this value.

7. Enter the URL in the **Object Storage URL** textbox. You can enter the URL in either of the following formats:

   - URL with fully qualified domain name.
     For example,

     ```
     https://<namespace>.swiftobjectstorage.<your-region>.oci.customer-oci.com/v1/<your-namespace>/<your-bucket>
     ```

     ```
     https://<namespace>.objectstorage.<your-region>.oci.customer-oci.com/n/<your-namespace>/b/<your-bucket>/o
     ```

   - If you want to use the URL provided by the OCI Console, specify the URL only till the name of the bucket.
     For example,

     ```
     https://<namespace>.swiftobjectstorage.<your-region>.oci.customer-oci.com/v1/<your-namespace>/<your-bucket>
     ```

     ```
     https://<namespace>.objectstorage.<your-region>.oci.customer-oci.com/n/<your-namespace>/b/<your-bucket>/o
     ```

   - If you choose **Credential** as the **Connection Mode** (see step 6), specify the URL till *bucketname*/o
     For example,

     ```
     https://<namespace>.objectstorage.<your-region>.oci.customer-oci.com/n/<your-namespace>/b/<your-bucket>/o/
     ```

     > **Note:**
     >
     > Credential mode is available only to Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.

   The values for Region, Namespace and Bucket are auto-populated based on the URL provided.

8. To select the **Connection Mode** do one of the following:

   - Select **Swift Connectivity**, and provide the following details:
     - In the **User Name** text box enter your Oracle Cloud Infrastructure username. For tenancies that support identity domains, specify the domain name along with the username. For example, *<identity-domain-name>*/*<username>*.
     - In the **Token** text box enter the auth token.
   - (This is applicable only to Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.) Select **Credential** and provide the ODI credential in the **Enter Credential** text box.
     You must create the credential in the repository and in the Autonomous Database that you created during instance creation. When you create a data flow to map data from Object Storage to Autonomous Database you need to create the ODI credential in the target schema as well. Before you run the mapping, make sure that you select the step

and in the Properties panel, set the **Create credential** KM option to `false`. Otherwise, the credential-based connection will fail.

To create the credential, execute the following script:

```
begin DBMS_CLOUD.create_credential( credential_name => '<Credential
name>',
      username => '<oci username>', password => '<auth token>' ); end;
```

9. Click **Create**.

   The new connection is created.

10. Click **Test Connection**, to test the established connection.

The newly created connections are displayed in the **Connections** page.

Click the **Actions** icon (⋮) next to the selected connection to perform the following operations:

- Select **Edit**, to edit the provided connection details.
- Select **Test Connection**, to test the created connection.
- Click **Export** to export the connection. See Export Objects.
- Select **Delete Schema**, to delete schemas.
- Select **Delete Connection**, to delete the created connection.

You can also search for the required Connection to know its details based on the following filters:

- **Name** of the Connection.
- **Technology** associated with the created Connection.

Create a REST Server Connection

You can connect to any REST service endpoint, import the data entities into Data Transforms, and use them as source in a data flow.

To create a generic REST connector, you need to provide the JDBC URL, username, and password to connect to the endpoint. You can also create and upload a config file that contains information such as the authentication methods, endpoints, and tables that you want to import data entities from.

The **Application** tab on the **Create Connection** page includes two connection options to create a generic REST connection - Generic REST and Generic REST Config. This topic has the following sections:

- Creating a Generic REST Connection
- Creating a Generic Rest Connection Using a Config File

The newly created connections are displayed in the **Connections** page.

Click the **Actions** icon (⋮) next to the selected connection to perform the following operations:

- Select **Edit** to edit the provided connection details.
- Select **Test Connection** to test the created connection.

- Click **Export** to export the connection. See Export Objects.
- Select **Delete Schema** to delete schemas.
- Select **Delete Connection** to delete the created connection.

You can also search for the required connection to know its details based on the following filters:

- **Name** of the connection.
- **Technology** associated with the created connection.

Creating a Generic REST Connection

To create this connection you need to specify the REST service URL and choose a temporary schema where Data Transforms can create data entities after the reverse-engineering operation.

To define a REST server connection:

1. From the left pane of the Home page, click the **Connections** tab.
   Connections page appears.
2. Click **Create Connection**.
   Create Connection page slides in.
3. Do one of the following:
   - In the **Select Type** field, enter the name or part of the name of the connection type.
   - Select the **Applications** tab.
4. Select **Generic Rest** as the connection type.
5. Click **Next**.
6. The **Connection Name** field is pre-populated with a default name. You can edit this value.
7. In the **REST Service URL** textbox, enter the URL of the endpoint that services the REST resources.
8. In the **Proxy Host** textbox, enter the host name of the proxy server to be used for the connection.
9. In the **Proxy Port** textbox, enter the port number of the proxy server.
10. In the **User** text box enter the user name for connecting to the REST endpoint.
11. In the **Password** text box enter the password for connecting to the REST endpoint.
12. Choose a connection from the **Staging Connection** drop-down list. The list displays only existing Autonomous Database connections. To use a different connection, create the connection before you reach this page.
13. After providing all the required connection details, click **Test Connection** to test the connection.
14. Click **Create**.
    The new connection is created.

Creating a Generic Rest Connection Using a Config File

APPLIES TO: Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.

To create a generic REST connector, you need the JDBC URL, username, password, and a config file. The config file is a model file with the *file_name.rest* naming convention that you

need to upload when you create a REST Server connection. You need to specify the endpoints, table mappings, and the authentication methods to create the config file. You can create the config file using any text editor.

To define a REST server connection using a config file:

1. From the left pane of the Home page, click the **Connections** tab.
   Connections page appears.
2. Click **Create Connection**.
   Create Connection page slides in.
3. Do one of the following:
   - In the **Select Type** field, enter the name or part of the name of the connection type.
   - Select the **Applications** tab.
4. Select **Generic Rest Config** as the connection type.
5. Click **Next**.
6. The **Connection Name** field is pre-populated with a default name. You can edit this value.
7. Use the **Config File** text box to upload the config file that you want to use.
8. In the **JDBC URL** textbox, enter the URL to connect to the server.
9. In the **User** and **Password** text boxes enter the user name and password for connecting to the REST endpoint. You may leave these fields blank if these values are not applicable or are already mentioned in the JDBC URL.
10. After providing all the required connection details, click **Test Connection** to test the connection.
11. Click **Create**.
    The new connection is created.

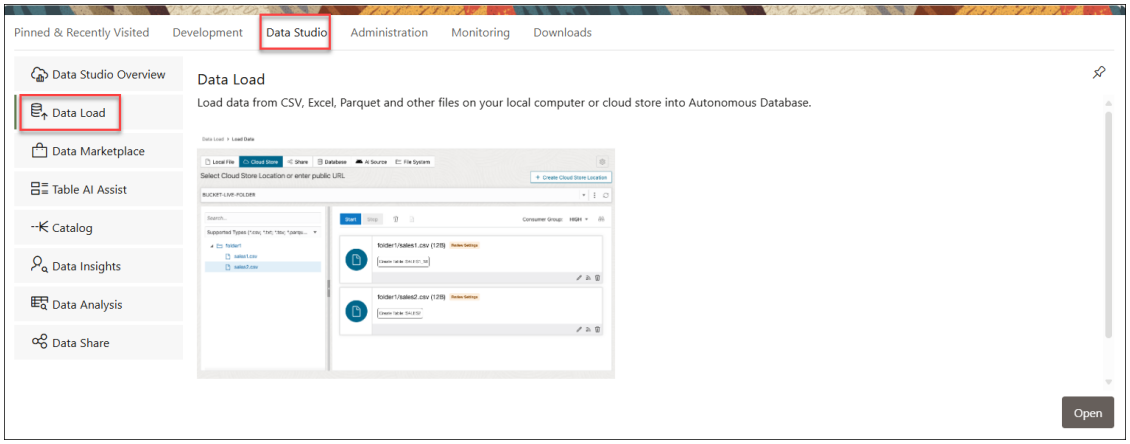
The Data Load Page

Use the Data Load dashboard page to make more data available to your Oracle Autonomous Database. You can load data from files or databases, from links to external databases or cloud storage files, or from a live feed of data from cloud storage.

From the Data Load page, you can also explore the data in your autonomous database and manage your cloud storage locations.

> **Note:**
>
> If you do not see the Data Load card then your database user is missing the required DWROLE role.

To reach the Data Load page, click the **Data Studio** tab in the Database Actions page and select the **Data Load** pane.



You can alternatively click the Selector icon and select **Data Load** from the Data Studio menu in the navigation pane.

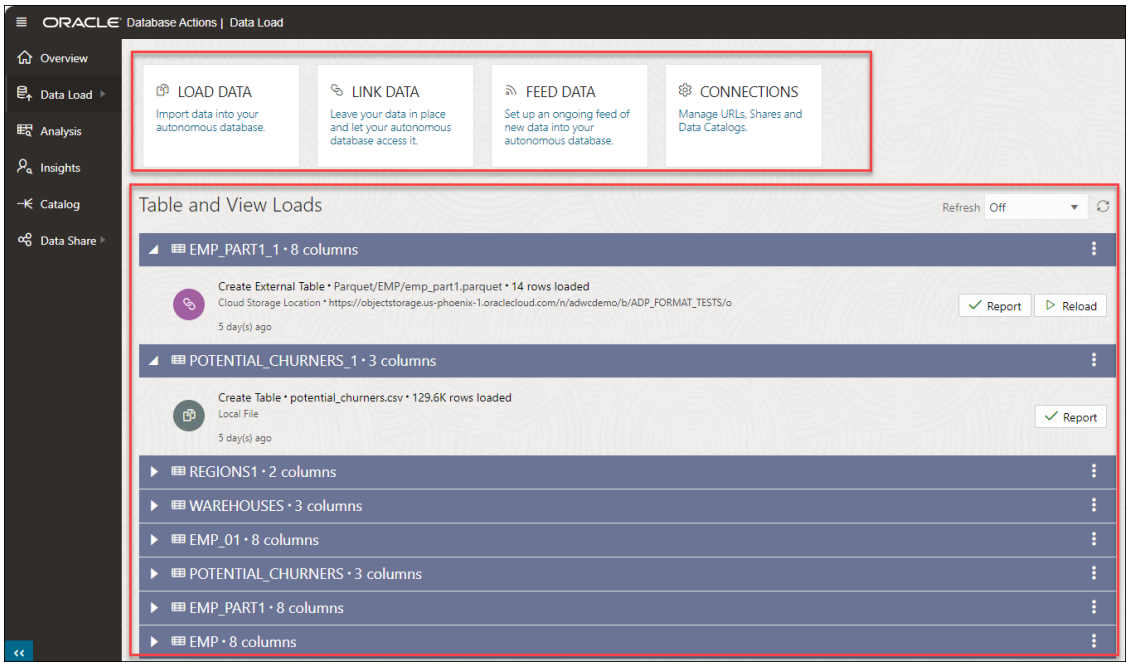
The Data Load Dashboard provides a central location for all data load and live feed tasks. The various icons on the dashboard, provide access to loading data, linking data, setting up an ongoing feed, checking data load jobs, and managing connections.

You can view the following tiles on the top section of the dashboard:

- Load Data
- Link Data
- Feed Data
- Connections

The bottom section of the dashboard displays the tables, views, and life feed you load into the tool.

The following image illustrates the **Data Load Dashboard**.



After drilling down onto a table, in the **Table and View Loads** section, you can view the details of the table load such as the name of the table, the number of rows and columns in the table, and the time you load the table.

Tables which are not active for last 7 days are collapsed.

You can also view the source location of the data you load. Click the location to view objects in the location.

Click the **Actions** icon beside the table name and select **Table** to view the following properties of the table:

- View Details
- Gather Statistics
- Register to Cloud Link
- Publish Listings to Marketplace
- Create Analytic View
- The Table AI Assist Tool
- Export Data to Cloud
- The Data Analysis Tool
- Edit
- **Drop** : Select **Drop** to delete the table.

Select **Report** to view a report of total rows processed successfully and failed for the selected table.

Select **Reload** to reload the source of the data load job. Clicking Reload takes you to the **Link Data** page where you can link your data source to the Autonomous Database.

> **Note:**
>
> For tables you load via Cloud Store, you can view **Create Live Feed** icon. See Create Live Feed from Data Load. After the creation of live feed is successful, you will view a Live Feed badge besides the table load header.

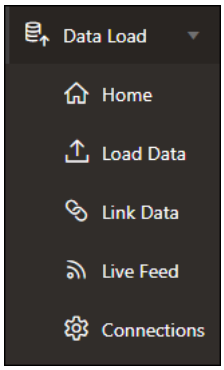
In case of errors, you can reload data with suggested fixes.

To load or create links to data or create live table feeds, on the Data Load dashboard, select a combination of an operation and a data source location. The following table lists the operations and the source locations that support those operations.

Operation	Source Location	Description
Load data	Local file Database Cloud storage Directory Share AI Sources	Load data from files on your local device, from remote databases, or from cloud storage into tables in your Oracle Autonomous Database.

Operation	Source Location	Description
Link data	Database Objects in Oracle Cloud Infrastructure (OCI) Catalog	Create external tables or views in your Oracle Autonomous Database that link to data in cloud storage or remote databases. Changes to the source data automatically appear in the target objects.
Feed data	Cloud storage	Set up a feed of data from a cloud storage bucket into a table. Changes to the source data load into the target table as scheduled or on demand.
Connections	Cloud storage Data Catalog links Share Providers	Create a cloud storage location. Subscribe to Share Provider and Register Data Catalog. You can manage cloud storage locations, Data Catalog and Shares.

The left navigation pane of the Data Load Dashboard provides links to the Data Load features and pages. When you select a link from the left navigation pane, the corresponding page appears in the target content area.



Topics:

- Data Studio Settings
  This topic explains the Data Studio settings on the Data Studio tool.
- Managing Connections
  Connections that are established from the Data Studio to the cloud, catalogs and shares are listed on this page.
- Loading Data
  You can load data from files on your local device, from remote databases, or from cloud storage buckets from directories and share providers. The file formats that you can load are CSV, XLS, XLSX, TSV, TXT, XML, JSON, GEOJSON, AVRO, PARQUET, GZ, GZIP, ZIP, PDF, PNG, JPG, JPEG and TIFF.
- Linking Data
  You can link to data in remote databases or in cloud storage buckets.
- Feeding Data
  You can run a live table feed on demand, on a schedule, or as the result of a notification.

- DBMS_LIVE_FEED Package Reference
  This chapter provides information about the packages you use with the Data Load Tool in Data Studio. The `DBMS_LIVE_FEED` topic also covers the procedures included in the `DBMS_LIVE_FEED` package.

Create and Run Data Loads

A data load allows you to load multiple data entities from a source connection to a target connection.

Topics

- Create a Data Load
  To create a data load you need to define the source connection and the target connection.
- Run a Data Load
  After you create the data load, you are taken to the Data Load Detail page that displays the details that you need to run a data load.

Create a Data Load

To create a data load you need to define the source connection and the target connection.

You can create a data load in either of the following ways:

- Using the Create Data Load wizard.
- Using the Projects page.

> **Note:**
>
> Data load is not supported for Oracle Object Storage connections.

The newly created data load is displayed in the Data Loads page of the associated project.

Click the Actions icon (⋮) next to the selected data load to edit, rename, start, export, or delete it.

Creating a Data Load from the Home Page

This section describes the generic steps to create a data load. If you plan to load and transform data using OCI GoldenGate, create the data load using the Projects page. See Creating a Data Load from the Projects Page.

To create a data load from the Home page:

1. From the left pane, click the Home tab. Click **Load Data**.
   The Create Data Load wizard appears.
2. In the Name field, enter a name for the data load. The field is pre-populated with a default name. You can edit this value.
3. Add a description. This is optional.
4. Select a project name from the drop-down. If this your first time here, click the + icon to create a project. If you have logged in as SUPERVISOR, the default project name is `Home`. For other users, the default project name is in the format `<username>_Home`. You can edit the default value. See Work with Projects for more information about projects.
5. Click **Next**.

6. To define your source connection, from the **Connection** drop-down, select the required connection from which you wish to add the data entities. Alternatively, click the + icon to create a new connection. See Work with Connections for more details about connections.
7. In the **Schema** drop-down, all schema corresponding to the selected connection are listed in two groups:
   - Existing Schema (ones that you've imported into Oracle Data Transforms) and
   - New Database Schema (ones that you've not yet imported).

   Select the schema that you want to use from the drop down.

> **Note:**
>
> If there is missing information such as user name or password not specified, wallet missing, and so on, the list may fail to populate with a “This connection has missing information.” error. Click the Edit icon ( ) to open the Update Connection page where you can fill in the missing details.

8. Click **Next**.
9. Similarly, define the target connection.
10. Click **Save**.
    The Data Load Detail page appears listing all the loaded data entities.

Creating a Data Load from the Projects Page

To create a data load from the Projects page,

1. On the Home page, click **Projects**, and then the required project title. In the left pane, click **Data Loads**, and then click **Create Data Load**.
   The Create Data Load wizard appears.
2. In the Name field, enter a name for the data load.
3. Add a description. This is optional.
4. Select the source and target schemas.

> **Note:**
>
> Make sure that you have created connections before you plan to create a data load using the Projects page. See Work with Connections for more details about connections.

   - To use the OCI GoldenGate Deployment Console to load data entities:
     a. Select the **Use GoldenGate** checkbox.
        The Create Data Load page now shows fields that are specific to OCI GoldenGate.
     b. Select the GoldenGate connection from the drop-down.
     c. To select the source connection, select the Registered Database and the Schema.
     d. Similarly, define the target connection.
   - To use all other connection types to load entities:

a. To define your source connection, select the connection from which you wish to add the data entities from the **Connection** drop-down.

b. From the **Schema** drop-down select the schema that you want to use. All schema corresponding to the selected connection are listed in two groups

   – Existing Schema (ones that you've imported into Oracle Data Transforms) and

   – New Database Schema (ones that you've not yet imported).

5. Click **Create**.
   The Data Load Detail page appears listing all the loaded data entities.

Run a Data Load

After you create the data load, you are taken to the Data Load Detail page that displays the details that you need to run a data load.

It includes the details of the source schema, the data entities that are loaded from the source schema, and the details of the target schema. You can choose the action that you want to apply on each data entity – recreate, truncate, append - and load tables in bulk to the target schema. When you run a data load, multiple jobs run in the background to complete the request.

You can run a data load in either of the following ways:

- Using the Data Load Detail Page.
- Running a Data Load from the Workflow Details Page.

> **Note:**
>
> Data load is not supported for Oracle Object Storage connections.

Running a Data Load from the Data Load Detail Page

The Data Load Detail page displays the information that you need to run a data load. You can apply different actions - incremental merge, incremental append, recreate, truncate, append - on the data entities before loading it to the target schema.

> **Note:**
>
> APPLIES TO: Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.
> If the data load is huge, you might want to increase the memory of the ODI Agent to avoid any issues. Follow the instructions in Increase the Memory of ODI Agent before you start to run the data load.

To run a data load from the Data Load Detail Page:

1. In the Data Load Detail page, select the data entities that you want to move to the target schema.
   To filter the list, you can do one of the following:
   - Enter the name or part of the name in the **Name** text box. This search returns data entities that include the specified string in the name. For example, the search string AD returns results such as ADMIN, ADDRESS, BADGE, UPGRADE, WORKLOAD, and so on.

- Turn on the **Use Regular Expression** toggle to use pattern matching to search for particular strings of characters. For example, the search string CO.* returns results such as CONTACT, COUNT, COUNTRY and so on.

You can select multiple data entities and load them to the target schema. You can also sort the displayed list using the following options:

- All - Displays all the data entities that match the search criteria.
- Selected - Displays all the rows that you selected.
- Unselected - Displays all the unselected rows.
- Invalid – Displays all the invalid rows.

> **Note:**
>
> These options display the list of data entities based on the search criteria. To view the list of all data entities, clear any applied filters.

2. Click on the required icon to choose any of the following actions:

- Incremental Merge - Updates the data in the selected column by comparing the source table with the target table based on the specified merge key. To use this option, select the column that you want to merge and then select the merge key. To use this option, select the column that you want to merge and then select the merge key. Click the Validate icon ( ) to validate the selected values.

  > **Note:**
  >
  > – The last update values for the Incremental Column are stored in the Data Load. If you create a new data load with the same table, this option acts like Append, Truncate, and Recreate the first time you run it.
  >
  > – If the target table is not truncated, do not start the Data Load using the Append, Truncate, or Recreate mode and then switch the action to Incremental Merge. This will cause every existing row in the target table to be updated. Oracle recommends that you start the Data Load using the Incremental Merge action.

- Incremental Append - Updates data in the selected column in the target schema. To use this option, select the column that you want to update and click the Validate icon ( ) to validate the selection.
  The first time you run the Data Load, this option works in the same way as the Append, Truncate, and Recreate options, meaning all rows are inserted into the target. On all subsequent Data Load runs, the Incremental Column is used to filter only new rows to append data to avoid the same rows from being loaded more than once.

> **Note:**
>
> – The last update values for the Incremental Column are stored in the Data Load. If you create a new data load with the same table, this option acts like Append, Truncate, and Recreate the first time you run it.
>
> – If the target table is not truncated, do not start the Data Load using the Append, Truncate, or Recreate mode and then switch the action to Incremental Append. Oracle recommends that you start the Data Load using the Incremental Append action.

- Recreate – If the table is already present in the target schema, drops the existing table and recreates it.

> **Note:**
>
> This option is not available for data entities that are loaded using OCI GoldenGate.

- Truncate – If the table is already present in the target schema, deletes all the data from the selected table. Nothing is dropped.

> **Note:**
>
> For Delta Share data loads the Data Load Detail page only includes the option. You cannot apply different actions - incremental merge, incremental append, recreate, truncate, append - on the data entities before loading it to the target schema. This is to make sure that the data is consistent between the Delta Sharing server and the target schema.

- Append – Inserts all rows in the Dataset into the target. If the Dataset contains records that already exist in the target, there will be duplicate records. Use Append if the source dataset is expected to have only new records or duplicate rows in the target are not a concern.
- Do Not Load – Skips the selected data entity from the data load job. After you click **Save**, these data entities are no longer available for future data load jobs.

You can select multiple data entities and apply different actions. The unsaved rows are highlighted in bold.

> **Note:**
>
> These options are not available for Delta Share connections.

3. To specify how you want to store the source column names in the target tables, click **Settings**, which is on the right side of the Data Load Detail page. Choose one of the following:
   - **Retain original names by enclosing all names with delimiters** - Creates column names with the same names as is from the source tables in the target table.

- **Use no delimiters** - This is the default selection. Converts all the column names to upper case and replaces spaces and special characters with underscores.
  The following options are applicable to reserved words such as `Date`, `Timestamp`, `Start`, and so on.
  - **Enclose with delimiters** - This is the default selection. Encloses column names that are reserved words with delimiters (not all column names).
  - **Use a prefix** - Adds the specified prefix to column names that are reserved words (not all column names).

  For column names that have the same name after conversion, the names are suffixed with a numeric value to maintain uniqueness. For example column names

  `Date, date, DATE, Item_@Code, Item$$Code, Item%%Code`

  are created in the target table as

  `DATE, DATE_0, DATE_1, ITEM__CODE, ITEM__CODE_0, ITEM__CODE_1.`

> **Note:**
>
> Once the data load is run, the selected options are applied and retained for all subsequent runs. You cannot change the configuration.

4. Click  to save the changes. A green checkmark () in the row indicates that the changes are saved.
5. To start the data load,
   - Click  .
   - For GoldenGate data loads, click  .

   A confirmation prompt appears when the data load starts successfully.

To check the status of the data load, see the Status panel on the right below the Target Schema details. For details about the Status panel, see Monitor Status of Data Loads, Data Flows, and Workflows. This panel shows links to the jobs that execute to run this data load. Click the link to monitor the progress on the Job Details page. For more information about jobs, see Create and Manage Jobs.

Running a Data Load from the Workflow Details Page

You can add multiple data loads to a workflow along with data flows or workflows and run them as separate steps. The left panel of the Workflow Details page lists the data flows, workflows, and data loads that are available for use.

For data loads, the left panel lists the following two folders:

- **Data Loads** - This folder lists all the data loads that you have created in the local Data Transforms instance. When you select a step in the workflow that is a local data load, the Properties Panel available on the right side of the design canvas displays the Type as `Data Load` and the Linked Object as Home>`<nameofDataLoad>`.
- **Remote Data Loads** - This folder lists all the data loads that you have created in a remote Data Transforms instance. See Create a Data Transforms Connection for Remote Data Load. When you select a step in the workflow that is a remote data load, the Properties Panel available on the right side of the design canvas displays the Type as `Remote Data`

`Load` and the Linked Object as `Name of the Data Transforms connection` > `<nameofDataLoad>`.

After you have added the data loads to the workflow, click  to execute them.

Work with Data Entities

A Data Entity is a tabular representation of a data structure.

It includes Database Tables or Views that can be used in a mapping as a source or target. They are simply the metadata for sources and targets. They are used in creating data flows.

You can add Data Entities to your newly created project in one of the following two ways:

- Importing Data Entities
- Creating Data Entities

All the newly created or imported Data Entities along with their details are displayed in the **Data Entities** page. The details include:

- **Name** of the Data Entity
- **Connection** for which the Data Entity was created
- **Schema** to which the Data Entity is associated
- Click the **Actions** icon (⋮) next to the selected Data Entity to perform the following operations:
  - Select **Edit**, to edit the existing details.
  - Select **Preview**, to preview the selected Data Entity. If the data entity belongs to an Oracle database, you can also view statistics of the table. See View Statistics of Data Entities for more details.
  - Select **Delete**, to delete the selected Data Entity.
- To delete the Data Entities in bulk, in the Data Entities page, select the check boxes of the respective Data Entities and click **Delete**.
- You can also search for the required Data Entity to know its details based on the following filters:
  - **Name** of the Data Entity
  - **Connection** for which the Data Entity was created
  - **Schema** to which the Data Entity is associated
  - **Tag** that is associated with the Data Entity

Topics

- Import Data Entities
  The easiest and most common way to create a Data Entity is by importing its structure from the connection type (for example, Oracle database).
- Create Data Entities
  You can manually create or update the Data Entities through the Oracle Data Transforms interface.
- Create Data Entities within the Data Flow editor
  If you have already created or imported your target data entity, then you would drag the data entity onto the Design Canvas and complete the column mappings and options.

- View Statistics of Data Entities
  The Preview tab displays detailed statistics of each data entity.

Import Data Entities

The easiest and most common way to create a Data Entity is by importing its structure from the connection type (for example, Oracle database).

To import existing Data Entities:

1. From the left pane of the Home page, click the **Data Entities** tab.
   Data Entities page appears.
2. Click **Import Data Entities**, to import the existing data entities.
   Import Data Entities page slides-in.
3. Select the **Connection** followed by **Schema** and then finally select the **Type of Objects** you want to import.
   For Oracle Object Storage connections, the Schema drop-down lists the name of the bucket that you specified in the URL when you created the connection.
4. [For Oracle Business Intelligence Cloud Connector (BICC) connections only] From the **Offerings to import for collection**, choose the offerings whose data stores you want to import. You must select at least one offering to import the BICC public view objects (PVO).

   > **Note:**
   >
   > The import of BICC PVOs can take a long time depending on the number of selected objects. To improve performance, Oracle recommends that you use a mask to filter and limit the number of PVOs that you want to import.

5. Choose a **Mask**/filter if you don't want to import every object in the schema. Depending on the Connection Type, you will be presented with further options for importing.

   > **Note:**
   >
   > For Oracle Object Storage connections, this value is case-sensitive. If **Batch similar files** is set to `True`, all the files that match the mask and have the same structure are grouped together into a single data entity.

6. [For Oracle Financials Cloud connections only] From the list in the **Resources** section, select the items that you want to import. When the import process completes, a table is created for each selected resource.
7. [For REST server connections only] In the **Resources** section, do the following:
   - In the **Resource URI** field, enter the URL of the REST service you want to import resources from.
   - Click the + icon.
   - In the **Name** column enter an identifier for the resource.
   - In the **Operation URI** column enter the URI of the resource.
   - Click **Test Resource** to check whether the entries are valid.
8. Click **Start**.

A Job is created and the corresponding Job ID is displayed for you to track the session. Click the Job ID to view the details of the job.

Upon successful execution of the job, all the selected Data Entities are imported. Click the Refresh icon present at the right corner of the Data Entities page, to see the new imported Data Entities.

Create Data Entities

You can manually create or update the Data Entities through the Oracle Data Transforms interface.

Data entities should possess the corresponding objects in the source connection to be used as a source in a data flow. Usually the import process makes sure that these objects are in coordination. However, whenever you manually create or update Data Entities always make sure to check if both the definitions are in coordination with each other.

When you use a Data Entity as a target then it doesn't have to exist previously in the target connection and can be created as a part of Data Flow execution.

To create a new Data Entity:

1. From the left pane of the Home page, click the **Data Entities** tab.
   Data Entities page appears.
2. Click **Create Data Entity**, to create a new data entity.
   Create Data Entity page appears.
3. In the Name text box, enter the name of the new Data Entity that you are creating.
4. From the **Connection** drop-down, select the required connection from which you wish to add the newly created Data Entity.

> **Note:**
>
> Oracle Financials Cloud connections are not listed here because you cannot manually create data entities for such connections. You can only import data entities from Oracle Financials Cloud REST endpoints using the Import Data Entities page. See Import Data Entities.

5. In the **Schema** drop-down, all schema corresponding to the selected connection are listed in two groups.
   - New Database Schema (ones that you've not imported from before) and
   - Existing Database Schema (ones that you've imported from before and are potentially replacing data entities).

   From the **Schema** drop-down, select the required schema.

> **Note:**
>
> For Oracle Object Storage connections, the Schema drop-down lists the name of the bucket that you specified in the URL when you created the connection.

6. From the **Type** drop-down, select the data entity type.

- **Table**: To define the table structure for the newly created Data Entity, click the + icon to add columns. For each column, you can specify parameters such as Name, Data Type, Length, Scale, Not Null. Double click on the cell to configure the value.
  Click the 'x' icon, to delete a row. Click the Up and Down arrows to sort the table rows.
- **Inline View**: To create the data entity using inline code, enter the Select statement in the **Query** tab. For example, `SELECT * FROM CUSTOMER`. Click **Validate**.
  The **Columns** tab displays a read-only list of the columns that the query returns. Click the **Preview** tab to see the column data.

7. In the **Tags** text box, enter a tag of your choice. You can use tags to filter the Data Entities displayed in the Data Entity Page.
8. For Oracle Object Storage connections, this page displays the following options:
   - **Contents** – Select the CSV file that contains the data you want to import. The metadata displayed in the preview table, such as the data type and length of columns, is based on the first row of the CSV file. Make sure that the CSV file has a header line. The header should contain only alphanumeric characters and no special characters.
   - **Group Files** – Select this check box if you want to group data from multiple CSV files into one data entity. For example, say you want to merge data from Employee_Data1.csv, Employee_Data2.csv, and Employee_Data3.csv into a single data entity.
     If you want to upload the CSV files using Windows, make sure that you first convert the files to Unix format to avoid character encoding issues during grouping.
   - **Resource Name** – Use this option along with **Group Files**. Enter the value you want to use to identify the files. The resource name should be a regular expression. You can use only an asterisk (*) as a wildcard character in the resource name. For example, `Employee_Data*.csv`.
9. For Oracle database connections, you can mark the data entities as a feature group. Expand **Advanced Options** and click the **Treat as Feature Group** checkbox.
10. Click **Save**.
    The new Data Entity is created.

Create Data Entities within the Data Flow editor

If you have already created or imported your target data entity, then you would drag the data entity onto the Design Canvas and complete the column mappings and options.

To create the definition of an entity while in the Data Flow editor,

1. Drag the data entity onto the Design Canvas.
2. Select the component and click the Add Data Entity icon  present on the top right corner of the target component.
3. **Add Data Entity** page appears allowing you to configure the following details of the target component:
   General tab
   - In the **Name** text box, enter the name of the newly created Data Entity.
   - From the **Entity Type** drop-down, select the data entity type.
     - **Table**
     - **ML Model**
       When you select this entity type the user interface changes as follows:

    * The **Connection Type** drop-down only lists **Oracle** as the option. Consequently, the **Connection** drop down only lists the Oracle connections that you have created.
    * The Add Data Entity wizard displays the **Properties** tab where you can select the Type of Learning, Function, Algorithm, and configure parameters to define the ML Model. See Use Machine Learning (ML) Models for detailed information about creating and using an ML Model data entity.

- From the **Connection Type** drop-down, select the required connection from which you wish to add the newly created Data Entity. The Connection drop-down is populated with the connections you have created with the associated connection type.
- From the **Connection** drop-down, select the server name from which you wish to add the newly created Data Entity.
- In the **Schema** drop-down, all schema corresponding to the selected connection are listed in two groups.
  – New Database Schema (ones that you've not imported from before) and
  – Existing Database Schema (ones that you've imported from before and are potentially replacing data entities).

  From the **Schema** drop-down, select the required schema.

  > **Note:**
  >
  > For Oracle Object Storage connections, the Schema drop-down lists the name of the bucket that you specified in the URL when you created the connection.

- In the **Tags** text box, enter a tag of your choice. You can use tags to filter the Data Entities displayed in the Data Entity Page.
- For Oracle database connections, you can mark the data entities as a feature group. Expand **Advanced Options** and click the **Treat as Feature Group** checkbox.
- Click **Next**.

Columns tab

- Click the  Add Columns icon, to add new columns to the newly created Data Entity.
  A new column is added to the displayed table.
- The table displays the following columns:
  – Name
  – Data Type - Click the cell to configure the required Data Type.
  – Scale
  – Length
  – Actions - Click the cross icon to delete the created column.
- To delete the columns in bulk, select the columns and click the  Delete icon.

- To search for the required column details, in the Search text box enter the required column name and click enter. The details of the required column are displayed.
- Click **Next**.

Preview Data Entity tab

It displays a preview of all the created columns and their configured details. If the data entity belongs to an Oracle database, you can also view statistics of the table. See View Statistics of Data Entities for more information.

4. Click **Save**.
The new target Data Entity is created.
5. Expand the **Properties Panel** in the right pane to view the following settings of the created components:
   - General - Displays the Name of the component along with its Connection and Schema details.
   - Attributes - Displays the details of all the attributes associated to the component.
   - Column Mapping - Click Auto Map to map all the columns automatically.
   - Preview - Click to have a preview of the component.
   - Options - Change the options as appropriate.

View Statistics of Data Entities

The Preview tab displays detailed statistics of each data entity.

> **Note:**
>
> This feature is available for Oracle database tables only.

You can view the statistics of a selected data entity in one of following ways:

- In the Data Entities list, click the **Actions** icon (⋮) next to the Data Entity and click **Preview**. Select the **Statistics** tab to view the statistics of the selected data entity.
- On any data flow click on any source or target data entity, and expand the properties panel in the right pane. Click **Preview**.

The statistical data is presented as follows:

- The total number of rows and columns in the data entity is displayed at the top.
- The statistics panel displays the thumbnail graphs for each column with information about the Min, Max, Distinct, and Null values.
- Two types of thumbnail representations are displayed based on the histogram:
  - A bar chart represents data for frequency and top-frequency histograms. The bar chart show the first top 10 values for the number of rows in the table.
  - A table lists data for Hybrid and Height-Balanced histograms. The table displays the entire data and is scrollable. The table displays the range for the values and the percentage of rows in each range.
- You can click each thumbnail to view the statistics of the column in a new browser tab.
- The detailed view of each chart also shows the type of histogram.

Work with Projects

A project is the top-level container, which can include multiple folders to organize your data flows or work flows into logical groups.

The Projects page displays all the existing projects, which includes projects created by all users not just the logged in user.

You can perform the following operations on a Project folder:

- View Details
- Edit the name
- Export the Project. See Export Objects.
- Delete

Creating a Project

To create a new project:

1. Click **Create Project** on the Projects page.
2. On the **Create Project** page, provide a project name.
3. Click **Create**.
   After creating the project, you're redirected to the **Projects** page.

Viewing Project Details

On the Projects page, select a project and click **View Details** from the project's **Actions** icon (⋮) or select a project to open its details page. The left panel displays all the resources associated with the selected project such as data loads, data flows, workflows, variables, schedules, and jobs.

Deleting a Project

When you delete a project, you also delete the resources it contains. Once you delete a project, it cannot be restored.

> **Note:**
>
> Be sure to review all contained resources before you delete the project.

To delete a project, on the Projects page, select Delete from the **Actions** icon (⋮) for the project you want to delete or select a project to open its details page, and then click **Delete**.

In the **Delete Project** dialog, click **Delete** to confirm the delete operation.

About Data Flows

A data flow defines how the data is moved and transformed between different systems.

A data flow in Data Transforms connects sources to targets through a flow of components such as Join, Filter, Aggregate, Set, Split, and so on. See Supported Database Functions for more information.
When you run a data flow, Data Transforms uses the joins, filters, mappings, and constraints to transform source data and load it to target tables. Note that you can run only one execution flow at a time. You cannot put multiple flows on a Data Flow and a flow cannot diverge into multiple flows.

- Create a Data Flow
  Create data flows to load data from a source connection, run transformations, and move the data to a target database.
- About Data Flow Editor
  The Data flow editor is divided into five parts, the Data Entity Panel, the Database Functions Toolbar, the Design Canvas, the Properties Panel, and the Status Panel.
- Supported Database Functions
  Oracle Data Transforms supports various database functions that you can drag and drop on the Design Canvas to connect components within a data flow.
- Add Components
  Add the data entities and database functions to the Design Canvas, and connect them in a logical order to complete your data flows.
- Use Text Embedding Vector in a Data Flow
  Data Transforms supports the use of vector datatype and embedding vectors in a data flow. Currently, Data Transforms integrates with OCI Generative AI service to convert input text into vector embeddings that you can use for data analysis and searches.
- Component Properties
  The Properties Panel displays various settings for components selected in the Design Canvas.
- Map Data Columns
  When you connect the source data entity with the target data entity, the column names are automatically mapped by the column names. You have a choice to map the columns by Position or by Name or map the columns manually using the Expression Editor.
- Validate and Execute a Data Flow
  After your mappings are ready, you can proceed to validate and execute the data flow.

Create a Data Flow

Create data flows to load data from a source connection, run transformations, and move the data to a target database.

You can create data flows in any of the following ways:

- From the Projects page
- From the Data Flows page within a project
- From the Home page

The newly created data flow is displayed in the Data Flows page of the associated project.

Click the Actions icon ( ) next to the selected data flow to edit, rename, copy, change folder, start, export, or delete it.

From the Projects page

To create a data flow from the Projects page,

1. On the Projects page, click **Create Data Flow**.
   Create Data Flow page appears:
2. In the **Name** field, enter a name for the new data flow.
3. Select **Create New Project**, if you wish to create new project folder for the newly created data flow.
4. Else, click **Add to Existing Projects**, if you wish to add the newly created data flow to an existing project folder.
5. If you have selected **Create New Project** for the previous option, in the **Project Name** field, enter the name of the newly created project.
6. Else, if you have selected **Add to Existing Projects** for the previous option, select the required project from the **Project Name** drop-down arrow.
7. In the **Description** field, enter a description for the newly created data flow.
8. Click **Create**.

From the Data Flows page within a project

To create a data flow from the Data Flows page within a project,

1. On the Projects page click the project tile you wish to create a new data flow for. The **Project Details** page appears.
2. In the **Data Flows** page, click **Create Data Flow**.
3. Provide the **Name** and **Description** of the new data flow.
4. Click **Next**.
5. To define your source connection, from the **Connection** drop-down, select the required connection from which you wish to add the data entities.
6. In the Schema drop-down, all schema corresponding to the selected connection are listed in two groups:
   - Existing Schema (ones that you've imported into Oracle Data Transforms) and
   - New Database Schema (ones that you've not yet imported).

   Select the schema that you want to use from the drop down. For Oracle Object Storage connections, the Schema drop-down lists the name of the bucket that you specified in the URL when you created the connection.
7. Click Save.
   The Data Flow Editor appears that allows you to create a new data flow.

From the Home page

To create a data flow from the Home page,

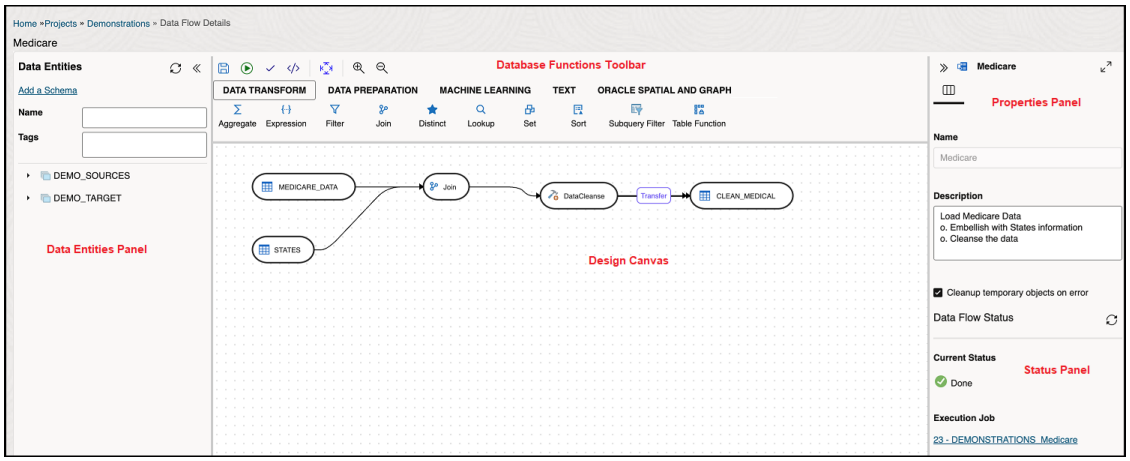
1. On the Home page, click **Transform Data**. The Create Data Flow page appears.
2. Provide the **Name** and **Description** of the new data flow.

3. Select a project name from the drop-down. Alternatively, click the + icon to create a project.
4. Click **Next**.
5. From the **Connection** drop-down, select the required connection from which you wish to add the data entities. Alternatively, click the + icon to create a new connection.
6. In the Schema drop-down, all schema corresponding to the selected connection are listed in two groups:
   - Existing Schema (ones that you've imported into Oracle Data Transforms) and
   - New Database Schema (ones that you've not yet imported).

   Select the schema that you want to use from the drop down.
7. Click **Save**.

About Data Flow Editor

The Data flow editor is divided into five parts, the Data Entity Panel, the Database Functions Toolbar, the Design Canvas, the Properties Panel, and the Status Panel.



- **Data Entities Panel**: The data entity panel displays the Data Entities that are available to use in your Data flows. The displayed list can be filtered using the Name and Tags fields. The panel includes options that let you add schemas, import data entities, remove any of the schemas that are associated with the data flow, and refresh data entities. See Add Components for information about how to use these options.
- **Database Functions Toolbar**: The Database Functions toolbar display the database functions that can be used in your data flows. Just like Data Entities, you can drag and drop the Database tools you want to use on the design canvas. See Supported Database Functions for more information.
- **Design Canvas**: The design canvas is where you build your transformation logic. After adding the Data Entities and Database Functions to the design canvas, you can connect them in a logical order to complete your data flows.
- **Properties Panel**: The properties panel displays the properties of the selected object on the design canvas. The Properties Panel is grouped into four Tabs. **General**, **Attributes**, **Preview Data**, **Column Mapping**, and **Options**. Not all tabs are available as they vary based on the selected object. See Component Properties to know more about these options.

- **Status Panel**: When you run a data flow, the Status Panel shows the status of the job that is running in the background to complete the request. You can see the status of the job that is currently running or the status of the last job. For more information about the Status panel, see Monitor Status of Data Loads, Data Flows, and Workflows.

After designing the required data flow,

- Click , to save the created/designed data flow.
- Click , to align the nodes of the designed data flow.
- Click , to execute the created data flow.
- Click , to validate the created data flow.
- Click , to maximize or minimize the created data flow diagram in the design canvas.

Supported Database Functions

Oracle Data Transforms supports various database functions that you can drag and drop on the Design Canvas to connect components within a data flow.

The Database Functions toolbar in the Data Flow editor includes the following database functions that can be used in your data flows. See Oracle Database SQL Language Reference for information about the database functions.

1. **Data Transformation**
   It contains the following components:
   - Aggregate
   - Expression
   - Filter
   - Join
   - Distinct
   - Lookup
   - Set
   - Sort
   - Subquery Filter
   - Table Function
2. **Data Preparation**
   It contains the following components:
   - Data Cleanse
   - Substitution
   - Equi_Width Binning
   - Quantile Binning
   - Lead
   - Lag
   - Replace

3. **Machine Learning**
   It contains the following components:
   - Prediction
   - Predition Model
   - Outlier Detection
   - Text Embedding Vector
4. **Text**
   It contains the following components:
   - REGEXP COUNT
   - REGEXP INSTR
   - REGEXP SUBSTR
   - REGEXP REPLACE
   - Edit Distance Similarity
   - Contains
5. **Oracle Spatial and Graph**
   It contains the following components:
   - Buffer Dim
   - Buffer Tol
   - Distance Dim
   - Distance Tol
   - Nearest
   - Simplify
   - Point
   - Geocode Tools:

     > **Note:**
     >
     > The following Geocode Tools work only in non-Autonomous Database environment.

     - Geocode As Geometry
     - Geocode
     - Geocode Addr
     - Geocode All
     - Geocode Addr All
     - Reverse Geocode

> **Note:**
>
> The following Geocode Tool works only in an Autonomous Database environment.

  - Geocode Cloud
- Spatial Join

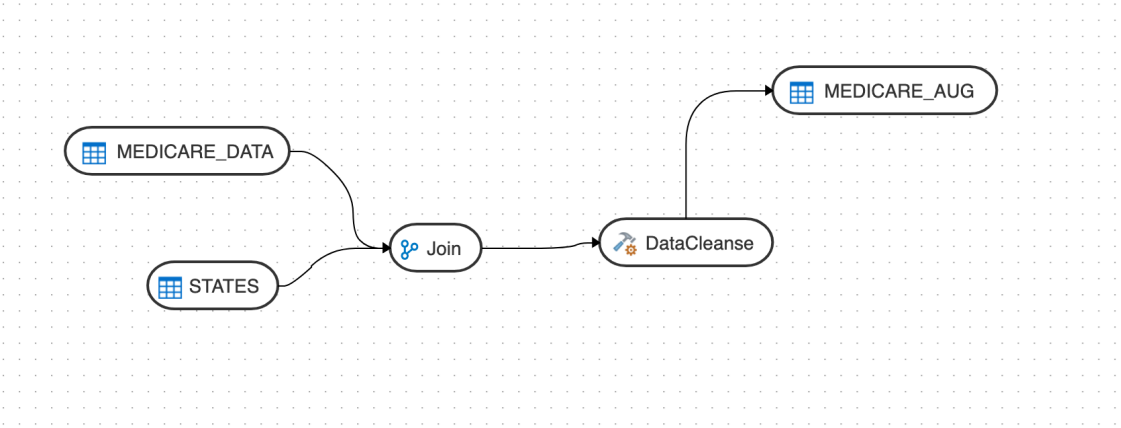
Add Components

Add the data entities and database functions to the Design Canvas, and connect them in a logical order to complete your data flows.

To add components to your data flow:

1. In the Data Entities panel, click **Add a Schema** to add schemas that contain the data entities that you want to use in the data flow.
2. In the Add a Schema page, select the connection and schema name.
3. Click **Import**.
4. In the Import Data Entities page, select the Type of Objects you want to import. Choose a Mask/filter if you don't want to import every object in the schema and click **Start**.
5. The Data Entities panel lists the imported data entities. The panel includes various options that let you do the following:
   - Refresh Data Entities – Click the Refresh icon to refresh the displayed list.
   - **Name** - Search for data entities by name.
   - **Tags** - Filter the data entities by the name of the tag used.
   - **Import Data Entities** - Right-click the schema to see this option. Use this option to import the data entities.
   - **Remove Schema** - Right-click the data entity to see this option. Use this option to remove the schema from the list. Note that this option does not delete the schema, it only removes the association of the schema with this data flow.
6. Similarly add more schemas to the Data Flow, if required.
7. Drag the required Data Entities that you want to use in the data flow and drop them on the design canvas.
8. From the **Database Functions** toolbar, drag the transformation component that you want to use in the data flow and drop them on the design canvas. You can use variables in the data flow. See Use Variables in a Data Flow for more information.
9. Select an object on the design canvas, and drag the **Connector** icon ( ) next to it to connect the components.
10. After saving the data flow, there may be a Transfer icon overlaid on one or more of the component connections. This indicates that ODI has detected an additional step and it is required to move the data between data servers. You can click on this Icon to view properties associated with this step.

For example:



Use Text Embedding Vector in a Data Flow

Data Transforms supports the use of vector datatype and embedding vectors in a data flow. Currently, Data Transforms integrates with OCI Generative AI service to convert input text into vector embeddings that you can use for data analysis and searches.

Before you use embedding vectors in a data flow, you need to do the following:

- Create an Oracle Database 23ai connection. See Work with Connections for generic instructions on how to create a connection in Data Transforms.
- Create an Oracle Cloud Infrastructure (OCI) Generative AI connection. See Create and use an Oracle Cloud Infrastructure Generative AI Connection.

To use vector embeddings in a data flow:

1. Follow the instructions in Create a Data Flow to create a new data flow.
2. In the Data Flow Editor click **Add a Schema** to define your source connection. From the **Connection** drop-down, select the Oracle Database 23ai connection and the schema that you want to use from the drop down. Click **OK**.
3. Drag the tables that you want to use as a source in the data flow and drop them on the design canvas.
4. From the Database Functions toolbar, click **Machine Learning** and drag the **Text Embedding Vector** transformation component drop it on the design canvas.
5. Click the **Text Embedding Vector** transformation component to view its properties.
6. In the **General** tab, specify the following:
   - **AI Service** - Select **OCI Generative AI** from the drop-down.
   - **Connection** - The drop-down lists all the available connections for the selected AI Service. Select the Oracle Database 23 ai connection that you want to use.
   - **AI Model** - The drop-down lists all the available models for the selected AI Service and Connection. The following models are listed:
     - "cohere.embed-english-light-v2.0"
     - "cohere.embed-english-light-v3.0"
     - "cohere.embed-english-v3.0"
     - "cohere.embed-multilingual-light-v3.0"
     - "cohere.embed-multilingual-v3.0"

You can also type in the model name.

7. In the **Column Mapping** tab, map the source column that you want to embed to the INPUT attribute of the operator. The only column available in the column mappings is `input_text`. Drag a text column from the available columns to the Expression column. This is the data that the vectors will be built on.
8. Drag the table that you want to use as a target in the data flow and drop it on the design canvas.
9. Save and execute the data flow.
   Data Transforms will build vectors for each of the rows in the source table and write that to the target table.

Component Properties

The Properties Panel displays various settings for components selected in the Design Canvas.

Depending on the component selected, you may see any of the following icons:

- General ( ) - Displays the name of the component along with its connection and schema details. You can edit some of these properties.
- Attributes ( ) - Displays the details of all the attributes associated with the component.
- Column Mapping ( ) - Allows you to map all the columns automatically. See Map Data Columns for more information.
- Preview ( ) - Displays a preview of the component. For Oracle tables, you can also view the statistics of the selected data entity. See View Statistics of Data Entities for details about the statistical information available.
- Options ( ) - Displays options such as
  - Truncate Table - Replaces any existing target table content with new data.
  - Append - Inserts records from the flow into the target. Existing records are not updated.
  - Incremental - Integrates data in the target table by comparing the records of the flow with existing records and updating the records when their associated data is not the same. Those that don't yet exist in the target are inserted.
    The option includes an Auto compression feature that is set to `True` by default. For data flow jobs that use the Incremental Update mode to load data onto a compressed Oracle target partition, the Auto compression feature recompresses the modified target partitions after the load completes successfully. For table partitions that are not originally compressed, the compression is skipped irrespective of whether Auto compression is set to true.

> **Note:**
>
> The Auto compression option is available to the ADMIN user or to a user with the DWROLE role. For data flows that have schema users other than ADMIN you need to either assign the DWROLE to the user or disable Auto compression to avoid execution errors.

Map Data Columns

When you connect the source data entity with the target data entity, the column names are automatically mapped by the column names. You have a choice to map the columns by Position or by Name or map the columns manually using the Expression Editor.

To map columns by Position or by Name:

1. Select the target Data Entity.
2. Click the arrow icon present on the top right corner to expand the Properties Panel. This will give you more space to work with.
3. In the **Properties Panel**, click the **Column Mapping** icon ( ).
4. To map the columns by Position or by Name, from the **Auto Map** drop-down menu, select **By Position** or **By Name**.

To map the columns manually:

1. From the **Auto Map** drop-down menu, select **Clear** to clear the existing mappings.
2. Drag and drop the attributes from the tree on the left to map with the Expression column.
3. To edit an expression, click the **Edit** icon of the respective column. The **Expression Editor** appears allowing you to perform the required changes (for example, you can just add an expression-"UPPER" or open the Expression Editor to edit the expression).

> **Note:**
>
> Use the expression editor only if you have complex expressions for a particular column.

4. Click **OK**.

Validate and Execute a Data Flow

After your mappings are ready, you can proceed to validate and execute the data flow.

Do the following:

1. Click Save.
   After saving, if data needs to be staged before transforming, Transfer button is added to one or more links. You can click these buttons to set more options, if available.
2. Click the **Code Simulation** icon ( ) if you want to check the code that will run to complete the tasks that are performed when you execute the data flow job. The source and target details are displayed in different colors for ease of reference. This is handy if you want to check if the mapping is correct before you run the job or if the job fails. Note that the code cannot be used for debugging. For detailed information about the job, see the Job Details page.
3. Click the **Validate** icon ( ) in the toolbar above the design canvas to validate the data flow.
4. After a successful validation, click the **Execute** icon ( ) next to the **Validate** icon to execute the data flow.

If you have added variables to the data flow, the **Variable values** page appears that displays the list of variables that you have added to the data flow. You can choose to use the current value, the default value, or set a custom value for each variable. Note that the custom value is applied only to the current run of the data flow. The custom value does not persist for any subsequent sessions.

A message appears that displays the execution Job ID and name. To check the status of the data flow, see the Status panel on the right below the Properties Panel. For details about the Status panel, see Monitor Status of Data Loads, Data Flows, and Workflows. This panel also shows the link to the Job ID that you can click to monitor the progress on the Jobs page. For more information, see Create and Manage Jobs.

For data flows created using Oracle Object Storage connections, the data from the source CSV file is loaded into the target Oracle Autonomous Database. You can also export data from an Oracle Autonomous Database table to a CSV file in Oracle Object Storage.

Schedule Data Flows or Workflows

You can schedule workflows and data flows to run at specified time interval.

Each schedule allows you to specify a start date and time, and execution frequency.

To schedule a data flow or a workflow:

1. In the left pane, click **Schedules**.
2. Click **Create Schedule**.
3. From the **Resource** drop-down menu, select **Data Flow** or **Workflow** as appropriate.
4. From the **Resource Name** drop-down list, select a data flow or workflow that you want to schedule.
5. From the **Frequency** drop-down list, select the time frame in which you wish to execute the created schedule. You can schedule it On Startup, Simple, Daily, Hourly, Weekly, Monthly or Yearly. Based on the selected frequency, **Time** field appears allowing you to schedule the selected Data Flow or Workflow. For detailed information refer to the below table:

Frequency	Time Values
On Startup	NIL
Simple	Click Select Date Time icon next to **Date and Time** field, to select the required date and time in `MM/DD/YY` and `HH:mm` format.
Daily	Click the clock icon next to **Time** field, to select the required time in `HH:mm` format.
Hourly	Click up and down arrows next to **Time** field, to select the required time in `Minutes:Seconds` format.
Weekly	• **Run Every** parameter appears with check-boxes for all the days of the week. • Click the clock icon next to **Time** field, to select the required time in `HH:mm` format.
Monthly (day of the month)	• Select the required date from the **Monthly Date** drop-down box. • Click the clock icon next to **Time** field, to select the required time in `HH:mm` format.

Frequency	Time Values
Monthly (week day)	• From the **Monthly date** drop down box select the required value. • From the **Week day** drop-down box, select the required weekday in which you wish to schedule the job.
Yearly	• From the **Month** drop down box, select the required month in which you wish to schedule the job. • From the **Monthly date** drop down box, select the required date in which you wish to schedule the job. • Click the clock icon next to **Time** field, to select the required time in `HH:mm` format.

6. In the **Number of Attempts on Failure** text box, enter the required value or click the up or down arrows next to the text box, to select the required value. This value denotes the number of retry attempts that should happen after the schedule failure.
7. **Stop Execution After** field denotes the time after which the schedule has to stop after executing it. It can be in **Hours**, **Minutes** or **Seconds**. Select the value and its unit.
8. Select the status of the scheduled Data Flow or Workflow from the **Status** options. It can be **Active**, **Inactive** or **Active for Period**. Through **Active for Period** option you can configure the exact time frame of the created schedule. It's Starting and Ending Date with time, time interval in which you wish to run the schedule and exceptions, if any.
9. After configuring the above details, click **Save**.
   The created schedules get listed in the **Schedules** page along with the following details:
   - Resource
   - Status
   - Frequency
   - Valid or Invalid

   Click the Actions menu (⋮) of the respective schedule to perform the following operations:
   - Click **Edit** to edit the details of the created schedule.
   - Click **Disable** to disable the created schedule. The status of a schedule becomes inactive when you disable it and gets displayed in the Scheduled list. Select the schedule and click **Enable**, to enable it whenever required. You can enable it again for a specific period, through **Enable for Period** option.
   - Click **Validate** to validate the created schedule.
   - Click **Export** to export the schedule. See Export Objects.
   - Click **Delete** to delete the created schedule. Upon confirmation the created schedule is deleted.

To schedule a data flow or a workflow specific to a project:

- Click the Project title displayed on the **Projects** page.
  Project Details page appears.
- In the left pane, click **Schedules**.
  Schedules page appears. It displays all the schedules pertaining to the selected project.

> **Note:**
>
> This page is project specific and displays only the Data Flows and Workflows pertaining to the selected project.

- To create a new schedule, click **Create Schedule**.
- Enter all the required details and click **Save**.
  The new schedule is created and is added to the existing list of schedules for the project.

Monitor Status of Data Loads, Data Flows, and Workflows

When you run a data load, data flow, or workflow Oracle Data Transforms runs jobs in the background to complete the request. You can view the status of the job in the panel on the bottom right of the Data Load Details, the Data Flow Editor, and the Workflow Editor page.

This panel includes the following:

- A Refresh icon ( ) to refresh the displayed status.
- The current status. You can see any of the following statuses:
  - Not Started – If no data load, data flow, or workflow has been executed yet.
  - Running – When the job starts.
  - Done – When the job completes.
    The status of the last job run persists in this panel till you run a new data load, data flow, or workflow.

> **Note:**
>
> For data loads created using OCI GoldenGate, the status panel shows the link to the GoldenGate Deployment Console and the status of the Extract and Replicat processes.

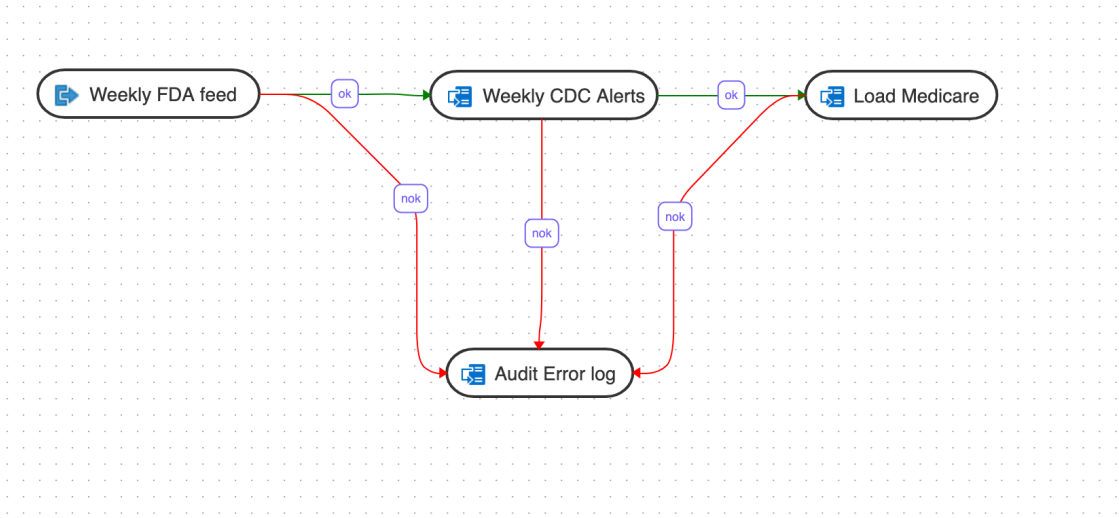
- The name of the job and a link that takes you to the Jobs Details page where you can see detailed information about the job and monitor the execution status. For more information, see Create and Manage Jobs.
  When you run a data load, multiple jobs run in the background to complete the request. This panel shows a link for each job that executes when running a data load.

Introduction to Workflows

A workflow is made up of multiple flows organized in a sequence in which they must be executed.

Each data flow is executed as a step. You can also add workflows, SQL queries, data loads, as well as variables as steps within a workflow. When you execute a workflow, a step either succeeds or fails. Depending on whether the first step succeeds or fails, you can choose the next step that must be executed.

Here is an example of a workflow:



In this example, the workflow performs the following actions:

1. Execute the "Weekly FDA feed" data flow.
2. If the "Weekly FDA feed" data flow execution is successful, execute the "Weekly CDC alerts" data flow.
3. If the "Weekly CDC alerts" data flow execution is successful, execute the "Load Medicare" data flow.
4. If any of the above data flow fails, then execute the "Audit_error_log" data flow.

Topics

- Create a New Workflow
  You can add data flows, workflows, variables, or data loads in a workflow.

Create a New Workflow

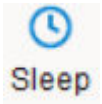
You can add data flows, workflows, variables, or data loads in a workflow.

To create a new workflow for your project:

1. On the **Home** page, click the required Project title.
   You are navigated to the **Project Details** page.
2. In the left pane, click **Workflows**.
   The Workflow page appears.
3. On the **Workflow** page, click **Create Workflow**.
   The **Create Workflow** page appears.
4. In the **Name** field provide a name for the new workflow and click **Create**.
   The new workflow is created and listed in the **Workflow** page with its **Name** and **Folder** details. Click the Actions icon (⋮) next to the workflow to edit, rename, copy, change folder, start, export, or delete it.
5. Click the Workflow to configure the **Workflow Details**.
6. From the left panel drag the data flows, workflows, variables, or data loads that you want to run in the workflow. To use variables in the workflow, see Use Variables in a Workflow for more information.

If you have connected to any other Data Transforms instance, you can also add data loads that you have created in that Data Transforms instance. See Create a Data Transforms Connection for Remote Data Load and Run a Data Load for more information.

7. Select either the ok (green arrow) icon, the not ok (red arrow) or the ok/not ok (black arrow) in the toolbar.
This defines the mode that subsequent links drawn on the canvas will be set to.
   - ok (green) defines the success path.
   - not ok (red) defines the failure path.
   - ok/not ok (black) defines a path that is followed on success or failure.
8. Use the Sleep icon (

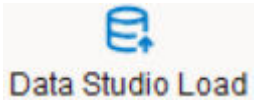


) to add a delay in the workflow. Drag the Sleep icon on the canvas and connect it in the flow with either the ok (green), not ok (red) or ok/not ok (black) links. This will add a delay at that point in the flow.
9. If you wan to add a SQL or PL/SQL query as a step in the workflow, do the following:
   - Drag the SQL icon (

) on the canvas.
   - Double click the SQL step in the editor to open the step properties page.
   - Select the **Attributes** tab.
   - From the **Connection** drop-down select the connection you want to run the query on.

     > **Note:**
     >
     > The drop-down lists only Oracle database connections.

   - In the **SQL** textbox add the query that you want to run.
10. If you want to add a Data Studio Data Load as a step in the workflow, use the Data Studio Load icon (



) on the canvas. See Define a Data Studio Data Load in a Work Flow for detailed instructions on how to use a Data Studio Data Load to load data to Autonomous Database.
11. Select the step and click the **Connector** icon () next to it to connect it with the next step.
12. After defining all the required Workflow details,
    - Click  , to save the created/designed workflow.

- Click , to align the nodes of the designed workflow.
- Click , to execute the created workflow.
  If you have added variables to the workflow, the **Variable values** page appears that displays the list of variables that you have added to the workflow. You can choose to use the current value, the default value, or set a custom value for each variable. Note that the custom value is applied only to the current run of the workflow. The custom value does not persist for any subsequent sessions.
- Select a single step in the canvas and click the **Execute Step** icon ( ), to execute only the selected data flow or workflow.
  To check the status of the workflow, see the Status panel on the right below the Properties Panel. For details about the Status panel, see Monitor Status of Data Loads, Data Flows, and Workflows. This panel shows the link to the Job ID that you can click to monitor the execution status on the Jobs page. For more information about jobs see Create and Manage Jobs.
- Click , to maximize or minimize the created workflow diagram in the design canvas.
  The newly created workflows get listed in the Project Details page. The following details are displayed:
  - Name of the Workflow
  - Folder corresponding to the workflow - Default Folder

  To view the details of a workflow, click the name of the workflow and you are navigated to the Workflow Details page.

In the left pane, you can search for the required Data Flow or Workflow using the Name filter. In the **Name** text box, enter the name of the required Data Flow or Workflow.

Select a step in the canvas and check the **Properties Panel** available at the right side of the design canvas to know the following properties of the selected step in a created Data Flow or Workflow:

- Name
- Linked Object

> **Note:**
>
> You cannot edit this field.

- Step -
  - First Step - Select the **First Step** check-box, to execute the selected step as the first step for execution in a Data Flow or Workflow.

> **Note:**
>
> You can select only a single step as a first step for execution in a Data Flow or Workflow.

- Number of attempts on Failure
- Time between attempts in seconds(s)

- Log steps in journal - Select any of the following options from the drop-down box:
  - Always
  - Never
  - Error

- Define a Data Studio Data Load in a Work Flow
  You can add a Data Studio Data Load as a step in a workflow and run the data load at scheduled intervals.

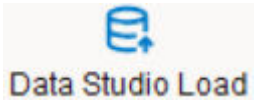
Define a Data Studio Data Load in a Work Flow

You can add a Data Studio Data Load as a step in a workflow and run the data load at scheduled intervals.

To do this, you need to create a connection, import data entities, create a data flow and then schedule the data load run.

To define and run a Data Studio Data Load within a workflow:

1. On the **Home** page, click the required Project title.
   You are navigated to the **Project Details** page.
2. In the left pane, click **Workflows**.
   The Workflow page appears.
3. On the **Workflow** page, click **Create Workflow**.
   The **Create Workflow** page appears.
4. In the **Name** field provide a name for the new workflow and click **Create**.
   The new workflow is created and listed in the **Workflow** page with its **Name** and **Folder** details. Click the Actions icon (⋮) next to the workflow to edit, rename, copy, change folder, start, export, or delete it.
5. Click the Workflow to configure the **Workflow Details**.
6. To add a Data Studio Data Load as a step in the workflow, drag the Data Studio Load icon (

) on the canvas.
7. Double click the Data Studio Load step in the editor to open the step properties page.
8. Select the **Attributes** tab.
9. From the **Connection** drop-down select the connection you want to load the data from.

   > **Note:**
   > The drop-down lists only Oracle database connections.

10. Select a Data Studio Load from the drop down.
11. Click , to save the created/designed workflow.

12. Click , to execute the created workflow.
13. To check the status of the workflow, see the Status panel on the right below the Properties Panel. For details about the Status panel, see Monitor Status of Data Loads, Data Flows, and Workflows. This panel shows the link to the Job ID that you can click to monitor the execution status on the Jobs page. For more information about jobs see Create and Manage Jobs.
The newly created workflows get listed in the Project Details page.

Create and Use Variables

A variable is an object that stores a single value, which can be a string, a number or a date.

Topics

- Create a Variable
You can specify the value of the variables when you create the variable or set the value using a query expression.
- Use Variables in a Data Flow
You can use variables in a data flow by using them as a filter definition. You can refer to variables in a data flow using the format `#<variable_name>`. During execution the variable is substituted by the value.
- Use Variables in a Workflow
The left panel of the Workflow Details page lists the variables that you can use as steps within a workflow.

Create a Variable

You can specify the value of the variables when you create the variable or set the value using a query expression.

The Variables page lists all the existing variables. Click the **Actions** icon ( ) of a variable to edit, refresh, or delete it. All the available variables are also listed in the left panel of the Workflow Details page. You can use variables as steps in data flows and workflows.

To create a variable:

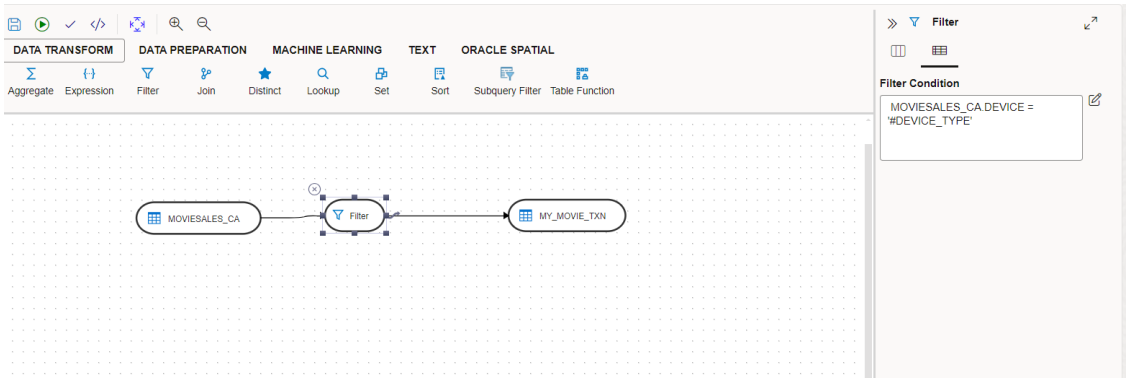
1. On the **Home** page, click the required Project title. You are navigated to the **Project Details** page.
2. In the left pane, click **Variables**. The Create Variable page appears.
3. Enter a name for the variable.
4. From the **Data Type** drop down, select one of the following:
    - **Short text** - The value can be alphanumeric, can contain special characters, and cannot exceed 255 characters.
    - **Long text** - The value can be alphanumeric, can contain special characters, and cannot exceed 64000 characters.
    - **Numeric** - The value can be a numeric value and can be preceded by a minus (-) sign. The value cannot exceed 10 digits. Possible value range is -999999999 to 9999999999.
    - **Date** - The value is a date format.

5. In the **Default Value** field enter the value you want to assign to the variable. The allowed value depends on the **Data Type** you select.
6. For variables that you modify, you can click the Actions menu next to the variable to open the Refresh Variable page to update the value. From the **Keep History** drop down select how you want the variable value to be displayed in the Refresh Variable page:
   - **All Values** - You can see the history of all the values held by this variable.
   - **Latest Value** - You can see only the latest value specified for the variable.
   - **No History** - Data Transforms does not keep the history of the values held by this variable.
   - **Secure Value** - This is useful when the variable contains passwords or other sensitive data. The value is not displayed in the Refresh Variable page.
7. Enter a description for the variable.
8. If you want the variable value to be set by a query, click the **Refresh** tab. Select the connection type and schema where you want to execute the command. Enter the query and click **Validate** to check the syntax of your expression.
9. Click **Save**.
   The newly created variable is listed in the Variables page as well as in the Variables node in the left panel of the Workflow Details page.

Use Variables in a Data Flow

You can use variables in a data flow by using them as a filter definition. You can refer to variables in a data flow using the format `#<variable_name>`. During execution the variable is substituted by the value.

When you execute the data flow, the **Variable values** page appears that displays the list of variables that you have added to the data flow. You can choose to use the current value, the default value, or set a custom value for each variable. Note that the custom value is applied only to the current run of the data flow. The custom value does not persist for any subsequent sessions.
Here is an example of the use of a variable in a data flow:



In this example, the data flow uses the variable DEVICE_TYPE as a step with the filter condition set as follows:

```
MOVIESALES_CA.DEVICE = '#DEVICE_TYPE'
```

The variable is quoted because the substituted value needs to be quoted. Since the variable value is substituted during execution, you can use it for many use cases such as in the transformation expression in the mapping.

Use Variables in a Workflow

The left panel of the Workflow Details page lists the variables that you can use as steps within a workflow.

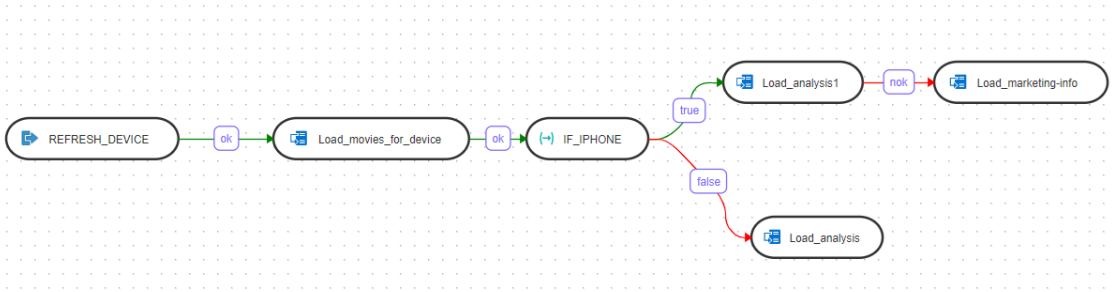
You can drag and drop the variables you want to use in the workflow on the design canvas. The Properties Panel available on the right side of the design canvas displays various details about the step such as the following.

- **Name**, **Type**, and the **Step** sequence number. These are non-editable values.
- Click the link under **Linked Object** to view and update the details about the variable.
- **Operation**: Select any one from the following:
  - **Set Variable** - There are two functions for this step:
    Update sets the current value of a variable.

    Increment increases or decreases a numeric value by the specified amount.
  - **Refresh Variable** - This variable step refreshes the variable by running the query specified in the variable definition.
  - **Evaluate Variable**: This variable step type compares the value of the variable with a given value according to an operator. If the condition is met, then the evaluation step is true, otherwise it is false.
- **Number of attempts on Failure**
- **Time between attempts in seconds(s)**
- **Log steps in journal** - You can select from Always, Never, or Error.

When you execute the workflow, the **Variable values** page appears that displays the list of variables that you have added to the workflow. You can choose to use the current value, the default value, or set a custom value for each variable. Note that the custom value is applied only to the current run of the workflow. The custom value does not persist for any subsequent sessions.

Here is an example of the use of a variable in a workflow:



In this example, a variable called DEVICE _TYPE with the value "iphone" is used as a step in the workflow. The workflow performs the following actions:

1. Execute the "REFRESH_DEVICE data flow.
2. If the "REFRESH_DEVICE" data flow execution is successful, execute the "Load_movies_for_device" data flow.

3. If the "Load_movies_for_device" data flow execution is successful, execute the "IF_IPHONE" variable evaluation step.
4. If the variable value is equal to ‘iphone’, execute path for “true”.
5. If the variable value is not equal to ‘iphone’, execute path for “false”.

Use Machine Learning (ML) Models

Data Transforms supports the use of ML Model in a data flow. Learn how to create and use Machine Learning (ML) Models in data flows.

Topics

- Create an ML Model Data Entity in the Data Flow editor
  To use ML models in Data Transforms you need to create two data flows. You need to first build the ML model data entity using the Data Flow editor, and then you can use the data entity in a data flow to mine data from a source connection and load it into a target server.
- ML Model Data Entity Properties
  The **Properties** tab of the **Add Data Entity** wizard provides data mining options that you can use to define the ML Model data entity.
- Use ML Model in a Data Flow
  You can use the **Prediction Model** database function to run ML Model algorithms on source data and load the output to a target database.

Create an ML Model Data Entity in the Data Flow editor

To use ML models in Data Transforms you need to create two data flows. You need to first build the ML model data entity using the Data Flow editor, and then you can use the data entity in a data flow to mine data from a source connection and load it into a target server.

To build an ML Model data entity in the Data Flow editor,

1. Drag the data entity that you want to build the ML Model on onto the Design Canvas.
2. Select the component and click the Add Data Entity icon present on the top right corner of the target component.
3. **Add Data Entity** page appears allowing you to configure the following details of the target component:
   General tab
   - In the **Name** text box, enter the name of the newly created Data Entity.
   - From the **Entity Type** drop-down, select **ML Model** as the data entity type.
     When you select this entity type the user interface changes as follows:
     – The **Connection** drop down only lists the Oracle connections that you have created.
     – The Add Data Entity wizard displays the **Properties** tab where you can select the Type of Learning, Function, Algorithm, and configure parameters to define the ML Model. See ML Model Data Entity Properties for more information.
   - From the **Connection Type** drop-down, select the required connection from which you wish to add the newly created Data Entity. For ML Model data entities, the **Connection Type** drop-down only lists **Oracle** as the option.
   - The Connection drop-down is populated with the connections you have created with the associated connection type. From the **Connection** drop-down, select the server name where you wish to keep the ML model data entity.

- In the **Schema** drop-down, all schema corresponding to the selected connection are listed in two groups.
  - New Database Schema (ones that you've not imported from before) and
  - Existing Database Schema (ones that you've imported from before and are potentially replacing data entities).

  From the **Schema** drop-down, select the required schema.
- In the **Tags** text box, enter a tag of your choice. You can use tags to filter the Data Entities displayed in the Data Entity Page.
- If you want to mark this data entity as a feature group, expand **Advanced Options** and click the **Treat as Feature Group** checkbox.
- Click **Next**.

Properties tab

- Select the **Type of Learning**, **Function**, and **Algorithm** you want to use to build this data entity. For more information about the options, see ML Model Data Entity Properties.
- Based on the options selected, the **Parameters** section is populated with the list of parameters that are marked as "Importance" and "High". You can add other required parameters using the  icon.
  You must specify a value for each parameter so that the data flow can run successfully.

Columns tab

- Click the  Add Columns icon, to add new columns to the newly created Data Entity.
  A new column is added to the displayed table.
- The table displays the following columns:
  - Name
  - Data Type - Click the cell to configure the required Data Type.
  - Scale
  - Length
  - Actions - Click the cross icon to delete the created column.
- To delete the columns in bulk, select the columns and click the  Delete icon.
- To search for the required column details, in the Search text box enter the required column name and click enter. The details of the required column are displayed.
- Click **Next**.

Preview Data Entity tab

It displays a preview of all the created columns and their configured details. If the data entity belongs to an Oracle database, you can also view statistics of the table. See View Statistics of Data Entities for more information.

4. Click **Save** to save the configuration and exit the wizard.

5. Save and execute the data flow.
   The new Data Entity is created. displayed in the Data Entities page.

ML Model Data Entity Properties

The **Properties** tab of the **Add Data Entity** wizard provides data mining options that you can use to define the ML Model data entity.

This topic assumes prior knowledge of Oracle Machine Learning concepts such as data mining functions and algorithms. For more information, see Oracle Machine Learning for SQL API Guide.

Type of Learning	Function	Algorithm
Supervised	Classification	Decision Tree
		Explicit Semantic Analysis
		Generalized Linear Models
		Naive Bayes
		Random Forest
		Neural Network
		Support Vector Machines
	Regression	Generalized Linear Models
		Neural Network
		Support Vector Machines
	Time Series	Exponential Smoothing
	Attribute Importance	Minimum Description Length
Unsupervised	Association	Apriori
	Attribute Importance	CUR matrix decomposition
	Anomaly Detection	One Class Support Vector Machines
	Clustering	Expectation Maximization
		k-Means
		Orthogonal Partitioning Clustering
	Feature Extraction	Explicit Semantic Analysis
		Non-Negative Matrix Factorization
		Singular Value Decomposition

Use ML Model in a Data Flow

You can use the **Prediction Model** database function to run ML Model algorithms on source data and load the output to a target database.

Before you use an ML Model in a data flow, you need build the ML Model. For instructions on how to create an ML model, see Create an ML Model Data Entity in the Data Flow editor.

To use an ML Model in a data flow:

1. Follow the instructions in Create a Data Flow to create a new data flow.
2. In the Data Flow Editor, drag the tables that you want to use as a source in the data flow and drop them on the design canvas.
3. From the Database Functions toolbar, click **Machine Learning** and drag the **Prediction Model** transformation component drop it on the design canvas.

4. Click the **Prediction Model** transformation component to view its properties.
5. In the **General** tab, specify the following:
   - **Connection** - The drop-down lists all the available Oracle connections. Select the Oracle connection that you want to use.
   - **Schema** - Select the schema.
   - **ML Model** - The drop-down lists all the available ML models. See Create an ML Model Data Entity in the Data Flow editor for instructions on how to build an ML Model.
6. In the **Column Mapping** tab, map the source column that you want to embed to the INPUT attribute of the operator. The only column available in the column mappings is `prediction parameters`. Drag a text column from the available columns to the Expression column.
7. Drag the table that you want to use as a target in the data flow and drop it on the design canvas.
8. Save and execute the data flow.
   Data Transforms will run the prediction model on the source data and write the output to the target table.

Create and Manage Jobs

When you execute a data load, data flow or workflow Oracle Data Transforms creates a job to complete the process in the background.

A job is made up of multiple steps that corresponds to an execution task.

The **Jobs** page by default, lists the execution job sessions that are running and completed for the present day as `Date` parameter is set to `Today` by default. You can click on any of the instance IDs from the list to see its details.

To work with jobs in Data Transforms, you can do the following:

- Create a Job from the Home page
- Create a Job from the Projects page
- View and Delete a Job
- Configure Purge Interval for Jobs
- Search for a Job

Create a Job from the Home page

To create a new job,

- On the Home page, click **Jobs** on the left pane to access the Jobs page.
- Click **Create Job**.
  The **Create Job** page slides-in.
- From the **Resource Type** drop-down, select the type of resource for which you wish to create a new job. It can be - Data Flow, Workflow or Schema.
- All the resources associated with the selected Resource Type get listed in the **Resource** field. From the **Resource** drop-down, select the required resource for which you wish to create a new job.
- Click **Create**.

A new job is created and is added to the existing list of jobs in the Jobs page.

For each session, you can see the Job Session **ID**, **Name**, **Status**, **Start Time** and **End Time** for each of the jobs that are executed.

Create a Job from the Projects page

To view the job sessions pertaining to a specific project:

- Click the Project tile displayed in the Projects page.
  Project Details page appears.
- In the left pane, click **Jobs**.
  Jobs page appears. It displays all the jobs and their details pertaining to the selected project.

> **Note:**
>
> This page is project specific and displays only the jobs pertaining to the selected project.

- To create a new job, click **Create Job**.
- Enter all the required details and click **Create**.
  A new job is created and is added to the existing list of jobs for the project.

View and Delete a Job

Click the Actions menu (⋮) to **View Details** of the jobs, **Delete** and **Rerun** the jobs, if required. If any of the steps failed, you can click on them to see the error details.

To delete a job, select the check box of the respective job session and click **Delete**. Upon confirmation the selected job gets deleted.

Configure Purge Interval for Jobs

You can configure the interval to purge older jobs to reduce the load on the sessions table. The default time interval is 30 days, which means that all jobs that are older than 30 days will be purged periodically. Note that only an ADMIN user (odiadmin) can configure the purge interval.

To configure the purge interval for jobs, click the gear icon (⚙) on the Jobs page. In the **Configurations** dialog box, set the required value in the **Delete jobs before (days)** field and click **Save**.

Search for a Job

You can also search for the required job session to know its details based on filters such as:

- **Name** - Name of the Job
- **Status** - Select the required session status from the list - All, Done, Error, Queued, Running, Waiting and Warning. Note that to use the Queued and Waiting filter options in the drop-down you must enable the **Show Queued/Waiting** Jobs option.
- **Show Queued/Waiting Jobs** – Enable this option to include the queued or waiting jobs in the list on the Jobs page. When this option is enabled, the **Status** drop-down includes the Queued and Waiting options that you can use to filter the list.

- **Date** - Select the date in which the required job session was executed - All, Today, Yesterday, Last Week, Last Month and Custom Range, which allows you to select specific From and To dates.

Export and Import Objects

You can move Data Transforms objects between environments. You require an Object Storage connection to store the exported file. You then need to use the same Object Storage connection to import the objects.

Topics

- Export Objects
  You can export Data Transforms artifacts such as projects, connections, data loads, data flows, workflows, and schedules from one environment and import them to another.
- Import Objects
  To import objects you need to provide the same Object Storage connection that you had used for the export operation.

Export Objects

You can export Data Transforms artifacts such as projects, connections, data loads, data flows, workflows, and schedules from one environment and import them to another.

The Export and Import options in Data Transforms allow you to do the following:

- Move code from development environments to production environments.
- Share code between development environments.
- Backup and restore objects.

To export objects from a Development environment you need to provide an Object Storage connection as an input. You need to use the same Object Storage connection when you want to run the import process in another environment. See Import Objects. The export operation exports all the metadata of the repository objects along with all its dependent parent and child objects, and stores the information in an XML format in the Object Storage bucket.

You can use the Export wizard to export the entire project or export individual objects at the following levels:

- Projects
- Connections
- Data Load
- Data Flow
- Workflow
- Schedule

When you export an individual object, all the dependent objects are automatically included in the export.

> **Note:**
>
> Individual variables export is not currently supported. However, when you export a Project, the variables within that project are also exported.

To export objects using the Export wizard:

1. From the left pane, click the Home tab. Click **Export**.
   The **Export** wizard appears.
2. In the **Export File Name** field, enter a name for the file. The exported data will be stored in the ZIP format with the *filename_timestamp_DTR* file naming convention.
3. Select an Object Storage connection from the drop-down to store the exported file. Note that you need to select the same Object Storage connection during the import operation.
4. Click **Next**.
5. Select the objects that you want to export. You can choose to export a single, multiple, or all the objects that are listed on the page.
6. Click **Next**.
7. Select the connections that you want to export.
   You may choose not to make any selection on this page. The connections that are associated with the objects that you selected in the previous screen are automatically included during the export operation.
8. Click **Next** to preview your selections.
9. Click **Export** to start the export operation.
   A confirmation prompt appears when the export operation starts. Click the link that is displayed in the prompt to monitor the progress of the export operation on the Job Details page.

To export individual objects:

1. Click the **Actions** menu of the project, connection, dataload, dataflow, workflow, or schedule that you want to export, and click **Export**.
2. In the **Export File Name** field, enter a name for the file. The exported data will be stored in the ZIP format with the *filename_timestamp_DTR* file naming convention.
3. Select an Object Storage connection from the drop-down to store the exported file. You need to select the same Object Storage connection during the import operation.
4. Click **Export** to start the export operation.
   A confirmation prompt appears when the export operation starts. Click the link that is displayed in the prompt to monitor the progress of the export operation on the Job Details page.

Import Objects

To import objects you need to provide the same Object Storage connection that you had used for the export operation.

Note the following about the import process:

- **Projects** - To import projects, you can choose whether you want to overwrite the data in the repository or merge the matching container objects with the objects that are being imported.
  - **Overwrite** - If there is a match in the repository, all the dependent child objects will be overwritten by default.
    If there is no match, the project along with the child objects will be copied into the repository.
  - **Merge** - If the project exists in the repository, but the dependent child objects have no match, all the child objects in the repository will be merged with the new child objects that are coming through the import process.

If the project exists in the repository, and the dependent child objects have a match, all the child objects in the repository will be overwritten with the new child objects that are coming through the import process.

If the project does not exist in the repository, the project along with the child objects will be copied into the repository.

- **Connections** - For connections that have a match in the repository, the import process updates only the advanced properties of the connection in the repository along with the connection name. Credentials such as the schema name, passwords, and connection URL are persisted as is in the new environment. You can use such connections without any modification in the target repository.
  For connections that have no match in the repository, credentials such as the schema name, passwords, and connection URL are not carried over during the import operation. Post import, you need to update the credential information for such connections. You also need to make sure that you create a schema user in the database along with the required tables created for the schema user, which you may use as a data entity in dataloads, dataflows, workflows, and schedules.
- **Data Loads**, **Data Flows**, and **Workflows** - For these objects if there is a match in the repository, all the objects will be overwritten by default.
- **Schedule** - The status for schedules that have no match in the repository are set to Inactive. You need to update the connection information that is associated with the schedule, and then set the status of the schedule to Active.

To import objects:

1. From the left pane, click the Home tab. Click **Import**.
   The **Import** dialog box appears.
2. From the **Object Storage Connection** drop-down, select the connection that you had used for the export operation. See Export Objects.
3. From the **Import File Name** drop-down, select the exported file.
4. Choose one of the following import options:
   - **Merge** - Choose this to merge the objects that are already in the repository with the objects that are coming through the import process. This option is only available for container objects such as Projects.
   - **Overwrite** - Choose this to replace the objects that are already in the repository with the objects that are coming through the import process.
5. Click **Import** to start the import operation.
   A confirmation prompt appears when the import operation starts. Click the link that is displayed in the prompt to monitor the progress of the import operation on the Job Details page.

Reference

Get additional information about using Oracle Data Tranforms on an Autonomous Database.

Topics:

- Enable Access to Private Data Sources from Autonomous Database
- Troubleshoot Mismatch with Server Cert DN Error
- Increase the Memory of ODI Agent

Enable Access to Private Data Sources from Autonomous Database

APPLIES TO: Data Transforms that is part of the suite of data tools built into Oracle Autonomous Database.

The Autonomous Database from where you are accessing Data Transforms should be configured to use a private endpoint to be able to communicate with private database sources. Otherwise when you try to create and test such a connection, you might get a "failed to connect" error.

If your Autonomous Database is configured to use a Private Endpoint, then you can only access private data sources from clients in the same Virtual Cloud Network (VCN). See Configuring Network Access with Private Endpoints for detailed instructions.

After you have configured network access using private endpoints, you need to do the following additional configurations for Oracle Data Transforms to communicate with private data sources. This chapter includes the following topics:

- Configuring DNS and Traffic Management
- Configuring RCE Proxy on the Private Endpoint

Configuring DNS and Traffic Management

You need to set up a fully qualified domain name (FQDN) and specify this FQDN in the connection host field when you create the connection. Refer to https://docs.oracle.com/en-us/iaas/Content/DNS/Concepts/views.htm for more information.

The VCN DNS resolver should know about the FQDN and be able to resolve it to a Type A IP address.

If you are not able to test a connection in Data Transforms on an Autonomous Database that is configured to use private endpoints, then do the following:

- Check whether the query to set the `ROUTE_OUTBOUND_CONNECTIONS` database property to `PRIVATE_ENDPOINT` is run. See Enhanced Security for Outbound Connections with Private Endpoints for more information.
- Make sure that the FQDN is used as the host so that the framework resolves the FQDN to the required IP address.
  Ensure that the data source VM has the FQDN set up with the required security rules.

If test connection fails even after using an FQDN, then do the following:

- From the OCI services menu, click **DNS Management** under **Networking**. Configure private DNS zones views and resolvers.
- Create and publish a record using the FQDN and the private IP (Type A) of the target database.

> **Note:**
>
> Do not use `*.oraclecloud.com` as the domain name when you set up the FQDN because it is a reserved domain name.

- Retry test connection.
  The test connection operation should complete successfully.

Configuring RCE Proxy on the Private Endpoint

For Autonomous Databases that are configured to use a private endpoint, the Autonomous Database service supports RCE to access resources in the customer subnet through a private route.

If you encounter a "failed to connect" error when you try to test a connection, contact your database administrator to check whether the database is RCE-enabled.

Note that RCE proxy can forward the reverse traffic only to the default OCI DNS resolver. If there is any overridden resolver, then connectivity will fail.

Troubleshoot Mismatch with Server Cert DN Error

APPLIES TO: ● Data Transforms that is part of the suite of data tools built into Oracle Autonomous Database.

For a connection on an Autonomous Database that is configured to use private endpoints, you must specify the distinguished name (DN) of the database server in the JDBC URL so that the Oracle Connection Manager (CMAN) can accept the request. Otherwise when you try to create and test such a connection, you might get a "Mismatch with Server Cert DN" error.

First, specify `ssl_server_dn_match=yes` in the JDBC URL for the CMAN to accept the request. Next, use the `ssl_server_cert_dn` parameter to specify the DN of the database server. Note that the order in which the keys are placed in `ssl_server_cert_dn` is important. The sequence should be `CN`, `O`, `L`, `ST`, `C`.

For example:

```
jdbc:oracle:thin:@(description= (retry_count=20)(retry_delay=3)
(address=(protocol=tcps)(port=1521)(host=xxxxx.adb.us-phoenix-1.oraclevcn.com))
(connect_data=(service_name=xxxxx.adb.oraclecloud.com))
(security=(ssl_server_dn_match=yes)(ssl_server_cert_dn="CN=adwc.uscom-
east-1.oraclecloud.com, O=Oracle Corporation, L=Redwood City, ST=California,
C=US")))
```

To get the values for `ssl_server_cert_dn`:

1. Login to the VM that is in the same subnet where your Autonomous Database lies.
2. Type the following command:
   ```
   openssl s_client -connect xxxxx.adb.us-phoenix-1.oraclevcn.com:1521 -showcerts
   ```
   You will find the values for the following under Server Certificate details:

   CN, O, L, ST, C

   For example, "`CN=adwc.uscom-east-1.oraclecloud.com, O=Oracle Corporation, L=Redwood City, ST=California, C=US`"

Increase the Memory of ODI Agent

APPLIES TO: ● Data Transforms that is available as a separate listing on Marketplace called Data Integrator: Web Edition.

If the data that you are loading from the source schema is huge, then you may want to increase the memory of the ODI Agent to avoid OutOfMemory exception errors.

To increase the memory of the ODI Agent:

1. Edit the `/u01/oracle/transforms_home/common/scripts/jettyServer.sh` file.

2. Add the `java -Xms1024m -Xmx4096m` parameter.
3. Restart the jetty server. Log in as OPC user and execute the following commands:

```
ssh -i <path to id_rsa> opc@<Instance IP>
sudo su
systemctl stop|start jettyserver.service
exit
```

Data Marketplace

Data Marketplace provides an innovative platform to access and share datasets effortlessly with Autonomous Database.

The Data Marketplace is a part of the Data Studio suite of tools that provides data scientists and business users with access to load and query datasets for their organizational requirements. You can access and discover data sets curated by data providers. You can integrate the shared data with the existing data, augmenting the data using AI capabilities.

This allows you as a producer to share data in the marketplace, which users of the system, either in the same tenancy or in a different tenancy, can then load the data locally in their database and consume it.

With the Data Marketplace, you can:

- Publish listings as a producer that enables you to access provide data to the intended audience based on how you configure the publishing. You can also view the statistical profile of your dataset.
- Load, analyze, and query the data as a consumer and leverage the AI capabilities of the Data Studio, like language detection, sentiment analysis, and key phrase extraction. You can also browse the datasets you can access based on your profile.

Topics

- The Data Marketplace Tool
  Data Marketplace provides an innovative platform to access and share datasets effortlessly.

The Data Marketplace Tool

Data Marketplace provides an innovative platform to access and share datasets effortlessly.

The Data Marketplace is a part of the Data Studio suite of tools that provides data scientists and business users with access to load, and query datasets for their organizational requirements. You can access and discover data sets curated by data providers. You can integrate the shared data with the existing data, augmenting the data using AI capabilities.

This allows you as a producer to *share* data in the marketplace, which users of the system (either in the same tenancy or different, depending on how you list your table) can then load the data locally in their database and *consume* it.

With the Data Marketplace, you can:

- Publish listings as a *producer* that enables you to access provide data to the intended audience based on how you configure the publishing. You can also view the statistical profile of your dataset.

- Load, analyze, and query the data as a *consumer* and leverage the AI capabilities of the Data Studio, like language detection, sentiment analysis, and key phrase extraction. You can also browse the datasets you can access based on your profile.

Data Marketplace Terminology

- **Dataset**: A unit of data you share in the Autonomous Database Marketplace as a producer. As a producer, you can create a dataset that supports tables and views.
- **Listings**: A list of tables available for consumption in the Data Marketplace. As a producer, you can publish a listing, specify the description, add a namespace, add tags, add links to documentation, and specify with whom you intend to share the data. Upon creation, the listings are displayed as cards on the Data Marketplace page.
- Access Data Marketplace Tool
  You can access shared datasets using the Data Marketplace tool, navigate to the Data Studio and launch the Data Marketplace tool.
- Publish Listings to Marketplace
  As a producer, you can create listings, configure the listing to prepare it for publishing, submit the listing for publishing to publish a table to Marketplace.
- Consume Dataset
  Depending on how you have made your listing available, consumers can access your dataset available on the Data Marketplace page.
- DBMS_DATAMART Package Reference
  This chapter provides information about the packages you use with the Data Marketplace Tool in Data Studio. The `DBMS_DATAMART` topic also covers the procedures included in the `DBMS_DATAMART` package. The `DBMS_DATAMART` package provides APIs to support common operations with Data Marketplace.

Access Data Marketplace Tool

You can access shared datasets using the Data Marketplace tool, navigate to the Data Studio and launch the Data Marketplace tool.

Prerequisites

You must run the following code as an `ADMIN` user to enable the Data Marketplace tool:

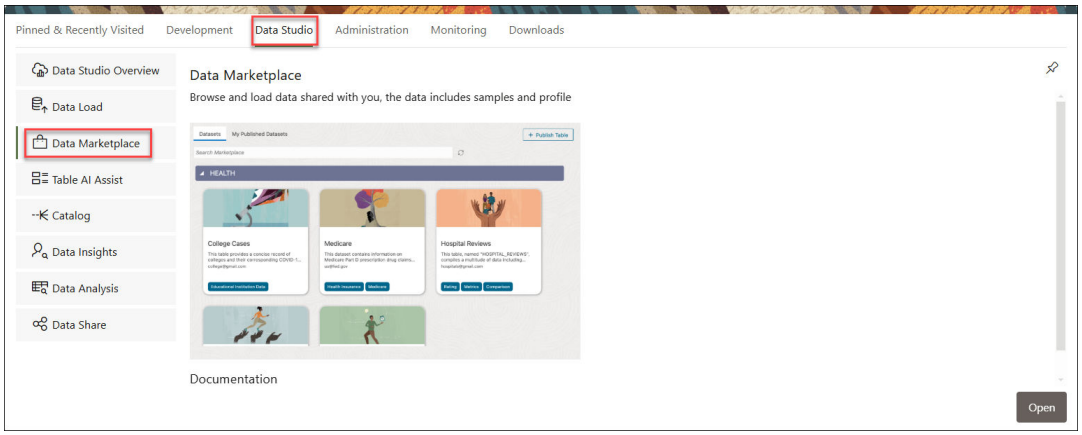
```
BEGIN
  DBMS_CLOUD_LINK_ADMIN.GRANT_READ( username =>'USER_NAME');
END;
```

> **Note:**
>
> - You must have the Data Marketplace tool enabled to view the listings. Otherwise, you will view an empty Marketplace page.
> - You can enable the Data Marketplace tool by running the above command or go to the Marketplace tab of the Data Studio Settings.

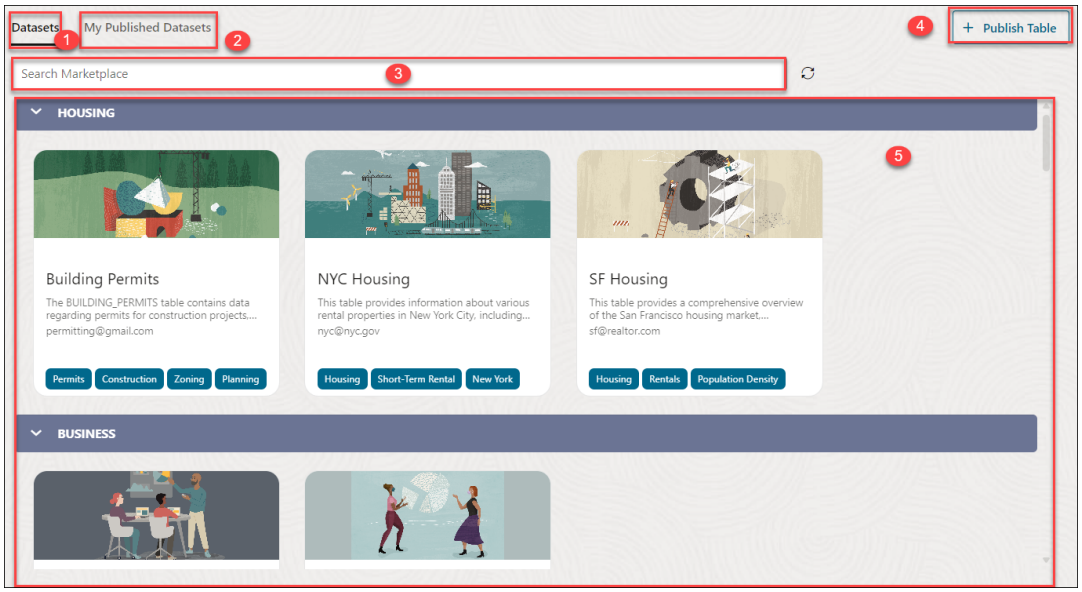
Follow these steps:

1. Navigate to the **Launchpad** on your Database Actions instance.

2. Click the **Data Studio** tab and select the **Data Marketplace** pane.



3. Clicking the Data Marketplace tab opens the Data Marketplace page.



You can view the following components on the Data Marketplace page:

- **Datasets tab**: This is the default tab selected when you open the Data Marketplace tool. You can view all the datasets shared with you available on this page.
- **My Published Datasets**: The datasets you publish are visible on this tab under the namespace or category you specify while publishing. You can **unregister** the listing by clicking any of the published datasets and selecting **Unregister**.
- **Search Marketplace Field**: You can search for the datasets from this field.

  Click the search field and enter the name of the Dataset you wish to view. You can also browse the datasets by the tags of the dataset. After you type the name and click enter, you can view the result in the display area. Click **Refresh** at the right end of the search field to refresh the page. Remove the search suggestion and click refresh to switch the view to its default view.

- **+Publish Table** icon: You can publish datasets using this icon. See Publish Listings to Marketplace for more details.
- **Display area:** The area below the search marketplace field displays the datasets returned by the search. You can expand the header which consists of the name of the category. It consists of the Dataset card with its name, description, and tags. Click the dataset card to view its details. See **Consume Dataset** for more information.

Publish Listings to Marketplace

As a producer, you can create listings, configure the listing to prepare it for publishing, submit the listing for publishing to publish a table to Marketplace.

Prerequisites

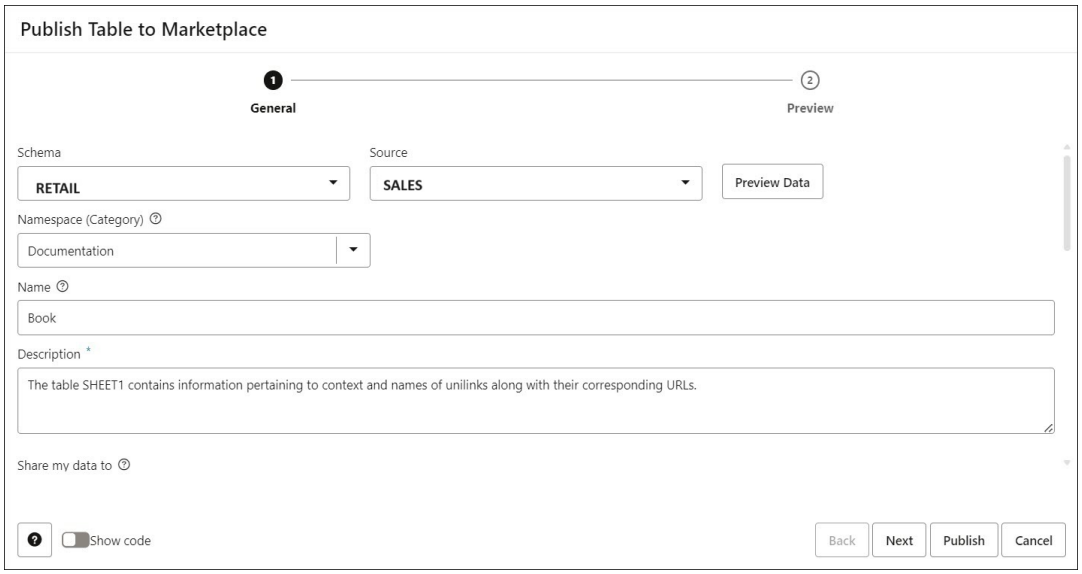
You must run the following code as an `ADMIN` user to publish datasets for a given region:

```
BEGIN
DBMS_CLOUD_LINK_ADMIN.GRANT_REGISTER( username =>'USER_NAME', scope =>
'MY$REGION');
END;
```

> **Note:**
>
> If you do not run the above command, the tool will raise an exception.

This topic describes how to publish listings to the Marketplace.

1. On the Data Marketplace page, select **+Publish Table**. You can view a **Publish Table to Marketplace** wizard.
2. On the **General** tab, specify the following details:



- **Schema**: Select the schema from the drop-down. For example, *RETAIL*.

- **Source**: You can view the source tables you load in the Autonomous Database using the Data Load tool. For example, *SALES*.

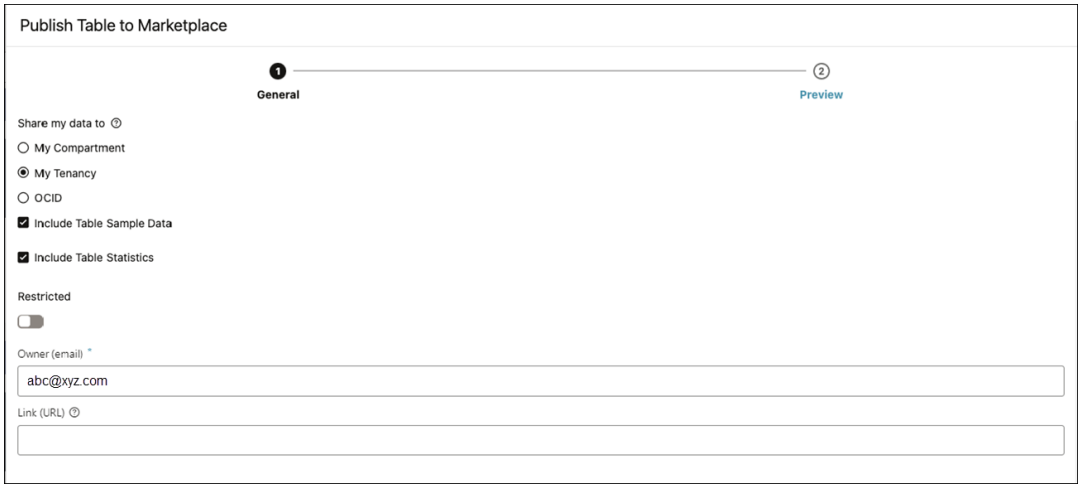
  > **Note:**
  >
  > *Note: If you have not loaded a table, the drop-down will be empty and the option will be greyed out. Make sure you have source data.*

  Click **Preview Data** to view the preview of the loaded table. You can also view the statistics of the table.

  See View Statistics of Data Entities for more information.

- **Namespace:** Specify the category of the data set you wish to group your data sets under. For example, *Documentation*.
- **Name:** Specify the name of the dataset. This should fall under the Namespace category. For example, *Book*.
- **Description:** This field is Auto filled using the AI capabilities of the Marketplace feature. It displays the description of the dataset shared in the listing. The description helps consumers understand the content of your dataset.

  > **Note:**
  >
  > – You can view the auto filled description only after you have configured and set AI profile using the Use DBMS_CLOUD_AI to Configure AI Profiles procedure.
  >
  > – After you have configured and set your AI profile, you must set the Data Studio Settings wizard on the Connections page.



- **Share my data to**: Specifies who and from where you are allowed to access the registered table.

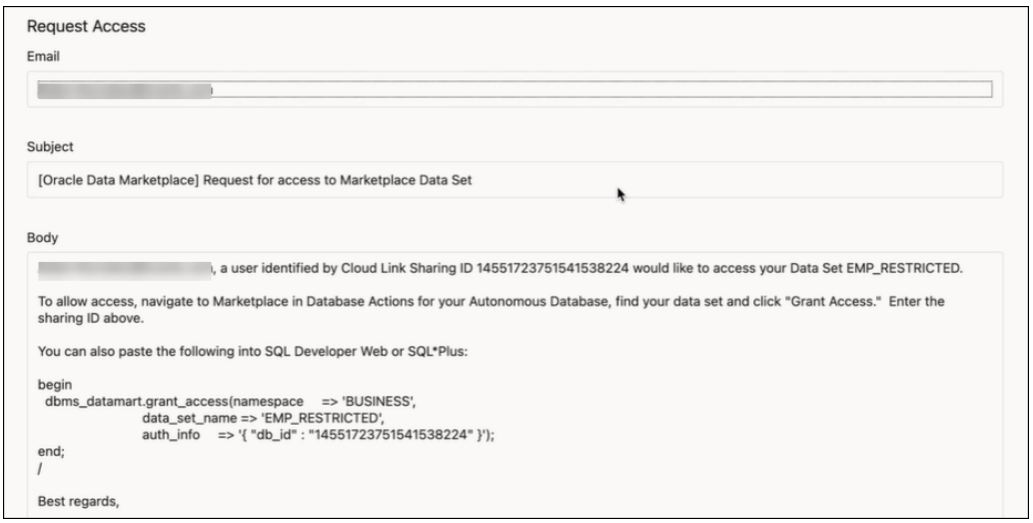
  You can choose from any of the following available options:

- **MY TENANCY**: You can grant remote data access to any resource, tenancy, compartment, or database in the tenancy of the Autonomous Database instance that is registering the data set.
- **MY COMPARTMENT**: You can grant remote data access to any resource, compartment, or database in the compartment of the Autonomous Database instance that is registering the data set.
- **OCID**: Access to the data set is allowed for the specific Autonomous Database instances identified by OCID.

> **Note:**
>
> Follow the steps provided in GET SHARING ID to obtain the OCID. It is alternatively known as `SHARING_ID`.

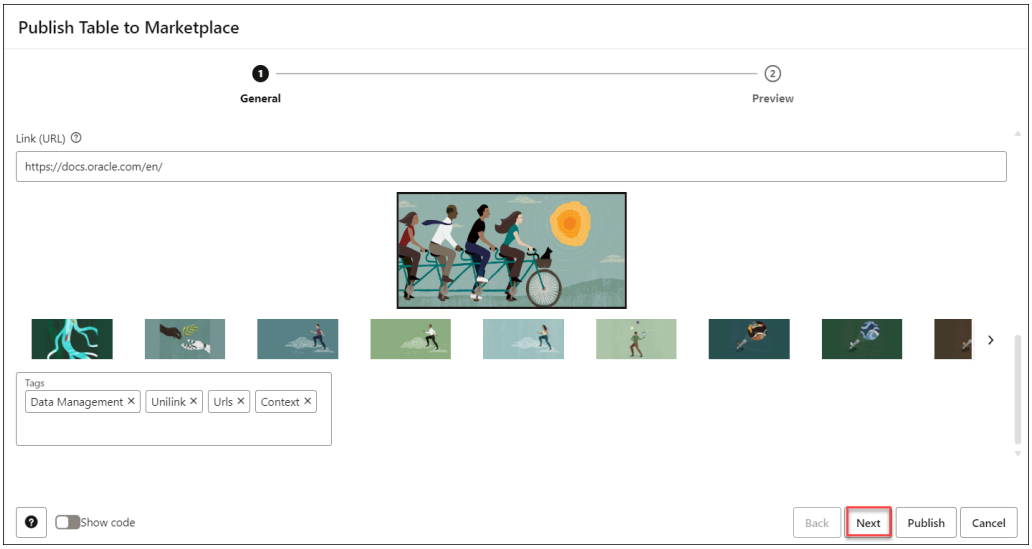
You can also include to share the Sample Data and Statistics of the table you select.

Restricted: When you publish a listing as restricted, the producer will receive an email notification with the Database ID for your listing.

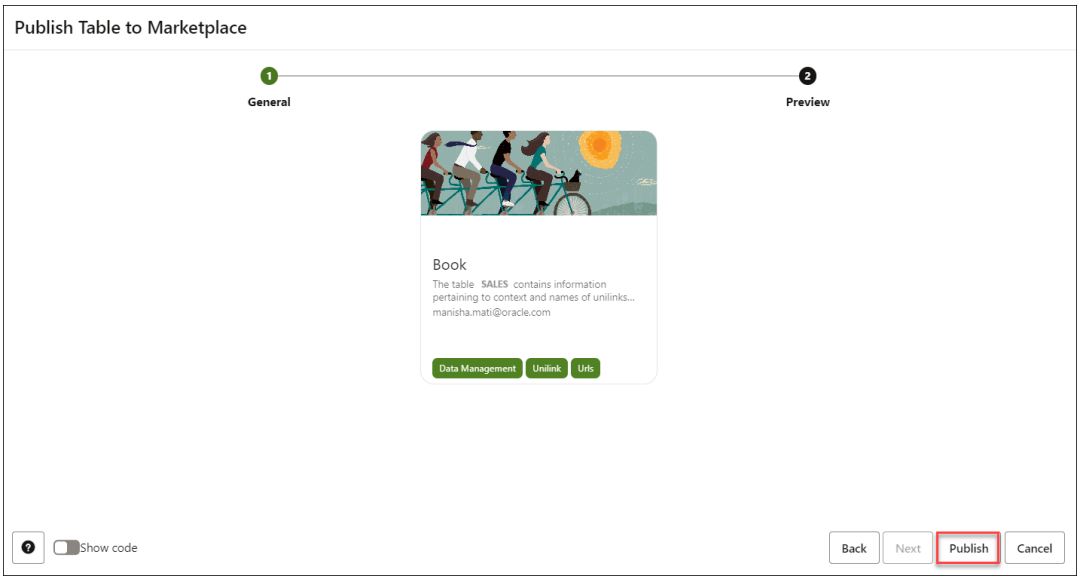


The producer shall grant access to the dataset for the consumer to consume the dataset by specifying the Database ID shared in the email.

- **Owner (email)**: Specify the email address of the producer, for example, *abc@xyz.com*.
- **Link (URL)**: Specify a URL that you can browse to discover more details on the data the provider shares. This field is optional.
- **Tags:** In the **Tags** text box, enter the tags of your choice. You can use tags to filter the Datasets on the search field.



3. Click **Next** to progress to the **Preview** tab. The preview tab displays the preview of the dataset card that will be displayed as a result of this publish.



4. Click **Publish** to publish the table to the Marketplace.

   You can view the published dataset in the **My Published Dataset** tab of the Data Marketplace page.

Consume Dataset

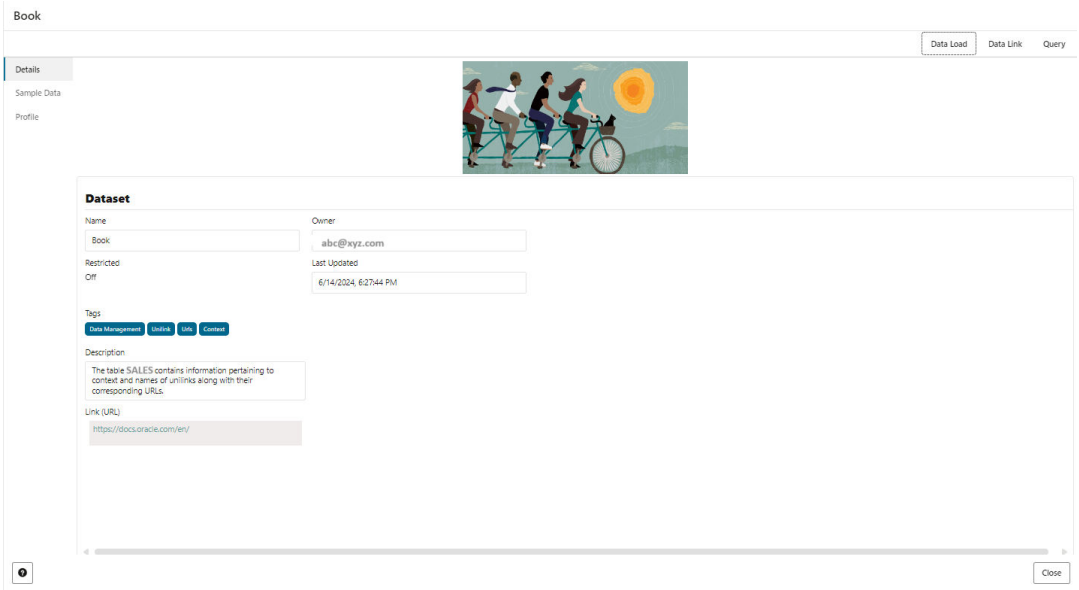
Depending on how you have made your listing available, consumers can access your dataset available on the Data Marketplace page.

As consumers you can view and consume the data in the dataset by following these steps:

1. On the Data Marketplace page select the listing whose data you wish to view.

   This opens the listing card.

2. You can view the following details of the Dataset:



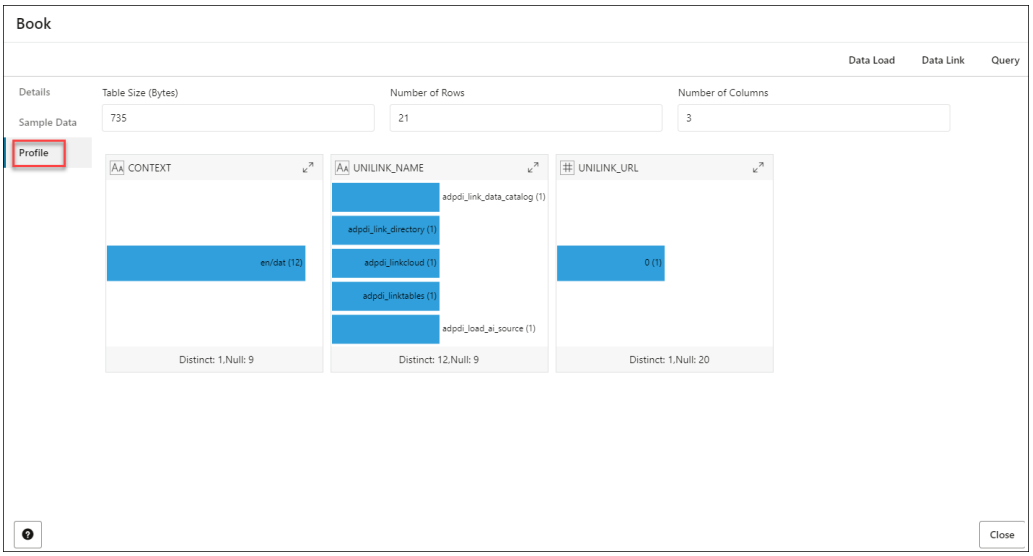
You can view the following fields on the Details tab:

- **Name:** Displays the name of the dataset.
- **Owner**: Displays the email address of the provider.
- **Restricted:** Specifies if the dataset shared is restricted or not.
- **Tags:** Displays the tags you selected while configuring the listing.
- **Description**: Displays description of the dataset shared in the listing.
- **Link(URL)**: Displays the URL the producer specifies during configuring the listing for publishing which displays more details on the data they share.

3. The Sample Data tab displays the preview of all the columns of the table selected in the dataset.



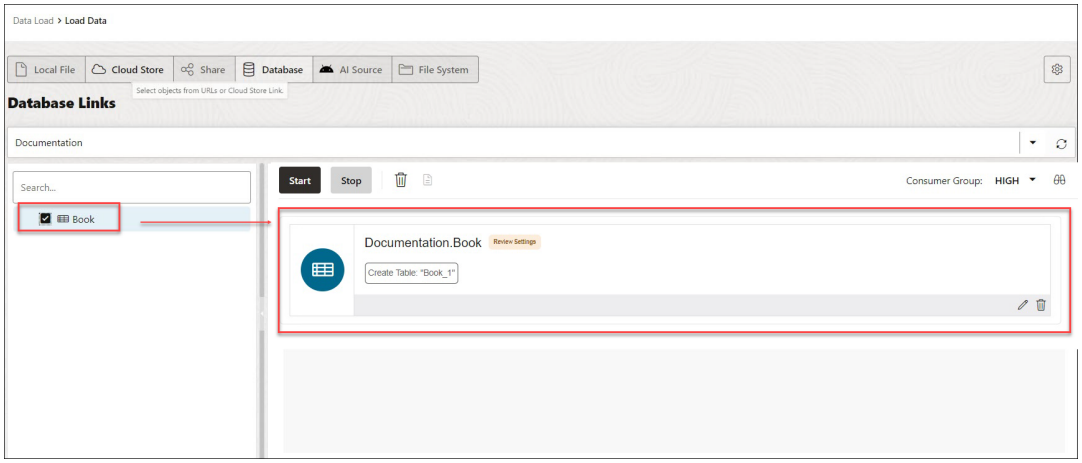
4. The Profile tab displays the statistical data of the dataset. You can view:
   - The total number of rows and columns in the data entity is displayed at the top.
   - The statistics panel displays the thumbnail graphs for each column with information about the Min, Max, Distinct, and Null values.
   - Two types of thumbnail representations are displayed based on the histogram:
     - A bar chart represents data for frequency and top-frequency histograms. The bar chart shows the first top 10 values for the number of rows in the table.
     - A table lists data for Hybrid and Height-Balanced histograms. The table displays the entire data and is scrollable. The table displays the range for the values and the percentage of rows in each range.
   - You can click expand thumbnail to view the statistics of the column in a new tab.
   - The detailed view of each chart also shows the type of histogram.



5. On the top header of the listing card details you can view the following icons:

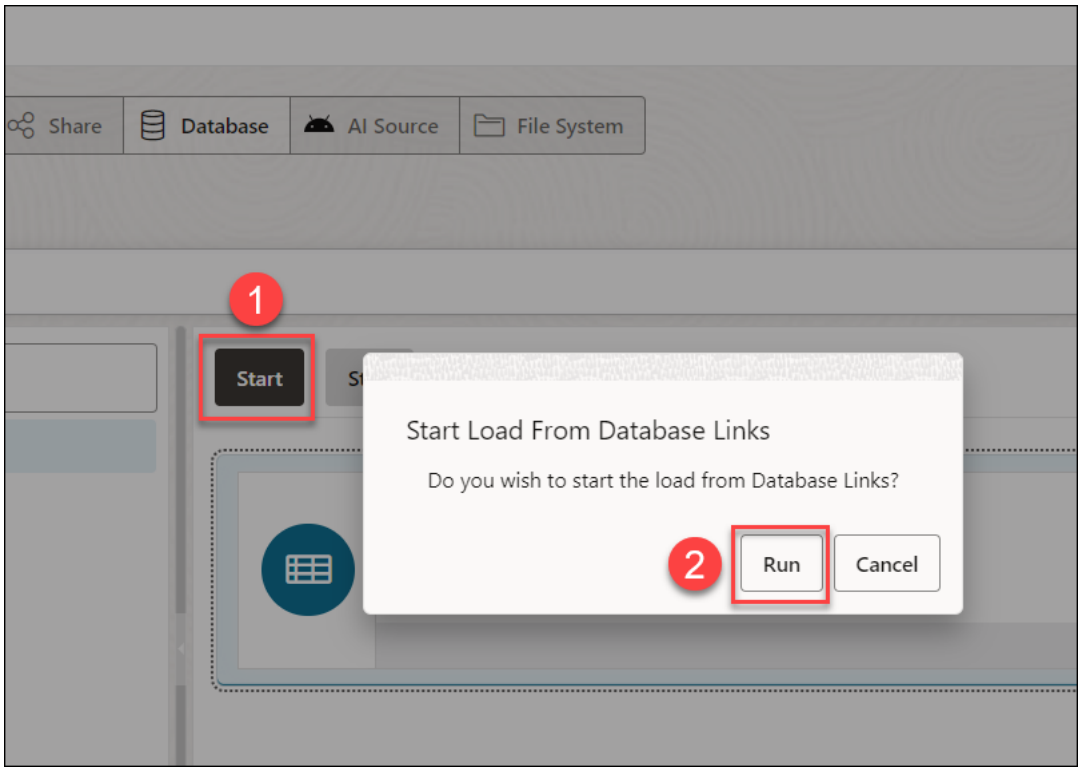
   Data Load: Selecting this option takes you to the Data Load page. You can create external tables derived from tables selected from the dataset.

   Drag and drop tables from the selected dataset.

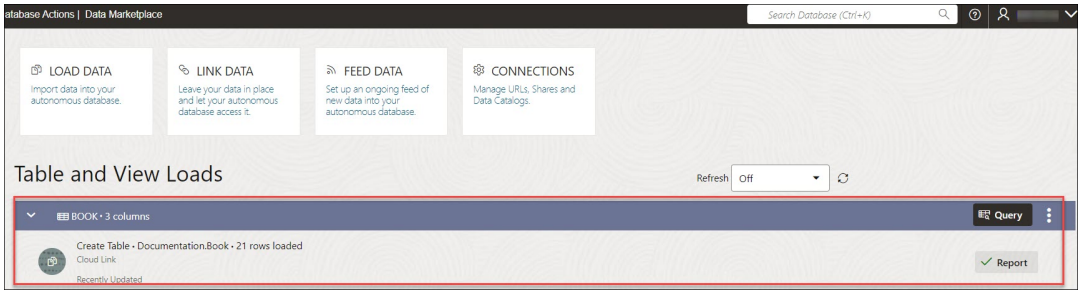


In this example, the only table selected is *BOOK*.

Create the external table by clicking **Start**, on the Load Data page, and then clicking **Run** on the Run Data Load Job dialog.



You can view the external table on the Dashboard under **Table and View Loads**.



- **Data Link**: Selecting this icon opens the Data Link page which is similar to the above operation of selecting the Data Load icon.
- **Query:** Selecting this icon opens the The Data Analysis Tool with the associated table as source input. This icon helps you analyze the data in the dataset and perform queries with the associated table. You can also create analytic views from the associated table in the dataset. See Data Analysis tool for more information.

DBMS_DATAMART Package Reference

This chapter provides information about the packages you use with the Data Marketplace Tool in Data Studio. The `DBMS_DATAMART` topic also covers the procedures included in the `DBMS_DATAMART` package. The `DBMS_DATAMART` package provides APIs to support common operations with Data Marketplace.

- Summary of DBMS_DATAMART Subprograms
  You can use Oracle Data Marketplace APIs to register and access datasets.

Summary of DBMS_DATAMART Subprograms

You can use Oracle Data Marketplace APIs to register and access datasets.

The `DBMS_DATAMART` package simplifies common operations with Oracle Data Marketplace, such as registering objects, requesting access, granting access, revoking access, unregistering and listing datasets.

The following table lists the `DBMS_DATAMART` subprograms and briefly describes them.

DBMS_DATAMART Package Subprograms

Subprogram	Description
register_object	This procedure registers datasets.
request_access	This procedure requests access to a specific dataset.
grant_access	This procedure grants access to a specific dataset.
revoke_access	This procedure revokes access to a specific dataset.
unregister	This procedure unregisters a dataset.
list_data_sets	This procedure returns the list of datasets.
get_stock_image_urls	This procedure returns the URL base from which you can select images.
get_reserved_namespaces	This procedure displays a JSON array of namespaces which are reserved by the system and cannot be used by you.
update_registration	This procedure updates an existing registration.

Register Object Procedure

This procedure registers datasets.

Syntax

```
procedure register_object(schema_name         in varchar2,
                          object_name         in varchar2,
                          namespace           in varchar2,
                          data_set_name       in varchar2,
                          scope               in clob default null,
                          thumbnail           in blob default null,
                          thumbnail_url       in varchar2 default null,
                          thumbnail_mime_type in varchar2 default null,
                          parameters          in clob default null);
```

Register Object Parameters

If you are accessing `Register Object` procedures, use the following parameters.

Parameter	Description
`schema_name`	The name of the schema.
`object_name`	The name of the object.
`namespace`	The namespace for the object.
`data_set_name`	The name of the dataset.
`scope`	This parameter indicates who can remotely access the data. The scope can be set to various levels, including to the region where the database resides, to individual tenancies, or to compartments. `MY$REGION` is not supported.
`thumbnail`	The image the tool displays in the cover of a dataset. You cannot use this parameter.
`thumbnail_url`	You cannot specify your own `thumbnail_url`. Oracle provides this parameter.
`thumbnail_mime_type`	It specifies Multipurpose Internet Mail Extensions (MIME) type of thumbnail.
`Parameters`	It specifies JSON parameters such as `tags`, `description`, `data_set_owner`, `authorization_required`, `offload_targets`, `offload_targets` and `store_sample`.

Request Access Procedure

This procedure requests access to the specified dataset.

Syntax

```
procedure request_access(namespace     in  varchar2,
                         data_set_name in  varchar2,
                         auth_info     out clob);
```

Request Access Parameters

If you are accessing `Request Access` procedure, use the following parameters.

Parameter	Description
`Namespace`	The namespace for the object.
`Dataset name`	The name of the dataset.
`Auth Info`	This parameter is generated populated as JSON on return. If the dataset provider has provided an email address and the requesting user has set up their Data Studio SMTP settings, the request will be submitted via email and this parameter will be generated with the `delivery` field set to smtp. Otherwise, the `delivery` field will be set to `browser` and you must use the rest of the metadata to submit their own request over email.

Grant Access Procedure

This procedure grants access to the specified dataset. The auth info parameter must contain the Sharing ID.

Syntax

```
procedure grant_access(namespace     in varchar2,
                       data_set_name in varchar2,
                       auth_info     in clob);
```

Revoke Access Procedure

This procedure revokes access to the specified dataset. The auth info parameter must contain the Sharing ID.

Syntax

```
procedure revoke_access(namespace     in varchar2,
                        data_set_name in varchar2,
                        auth_info     in clob);
```

Unregister Procedure

This procedure revokes access to the specified dataset. The auth info parameter must contain the Sharing ID.

Syntax

```
procedure unregister(namespace     in varchar2,
                     data_set_name in varchar2);
```

List Dataset Procedure

This procedure displays a list of datasets.

Syntax

```
function list_data_sets(registrations in boolean default false)
  return clob;
```

If `registrations` is `true`, it returns the list of datasets you register. If `false`, it returns the available data sets.

Get Stock Image URL Procedure

This procedure displays the URL from which you can select images.

Syntax

```
function get_stock_image_url return varchar2;
```

Get Reserved Namespaces Procedure

This procedure displays JSON array of namespaces that are reserved by the system. You cannot use these namespaces.

Syntax

```
function get_reserved_namespaces return clob;
```

Update Registration Procedure

This procedure updates an existing registration.

Syntax

```
procedure update_registration(namespace           in varchar2,
                              data_set_name       in varchar,
                              scope               in clob default null,
                              thumbnail           in blob default null,
                              thumbnail_url       in varchar2 default null,
                              thumbnail_mime_type in varchar2 default null,
                              parameters          in clob default null);
```

The parameters of this procedure are already described above.

View Data Pump Jobs and Import Data with Data Pump

Use the Data Pump page to view Data Pump jobs and use the wizard to quickly create and run Data Pump import jobs.

See The Data Pump Page for more information.

Link Data

You can link to data in remote databases, in cloud storage, or in connected file systems or the local file system. You can also use Database Links to access objects on another database.

When you link to data in a remote database or in cloud storage or in connected file systems or the local file system, the target object produced is an external table or a view. When you select

that target table or view on the Data Load, explore page in Database Actions, the source preview for the object shows the current data in the source object.

Linking to columns in a remote database, to files in cloud storage, or in connected file systems or the local file system can be useful when the source data is being continually updated. For example, if you have a database table with columns that are updated for each new sales transaction, and you have run a data load job that links columns of the source table to targets in your Oracle Autonomous Database, then those targets have the new sales data as it is updated in the source.

In addition, you can use Database Links to access objects on another database. Using Database Links, the other database need not be an Oracle Database system.

- Query External Data with Autonomous Database
  Describes packages and tools to query and validate data with Autonomous Database.
- Use Database Links with Autonomous Database
  You can create database links from an Autonomous Database to another Autonomous Database, to other Oracle databases, or to non-Oracle databases. In addition, you can create database links from other databases to an Autonomous Database.

Query External Data with Autonomous Database

Describes packages and tools to query and validate data with Autonomous Database.

External data is not managed by the database; however, you can use `DBMS_CLOUD` procedures to query your external data. Although queries on external data will not be as fast as queries on database tables, you can use this approach to quickly start running queries on your external source files and external data. Depending on the type of external table, you can validate external data using the `DBMS_CLOUD` validation procedures. The data validation procedures let you validate the source files for an external table so that you can identify problems and either correct the data in the external table or exclude invalid data before you use the data.

> **Note:**
>
> If you are not using `ADMIN` user, ensure the user has the necessary privileges for the operations the user needs to perform. See Manage User Privileges on Autonomous Database - Connecting with a Client Tool for more information.

- Query External Data
- Query External Data with ORC, Parquet, or Avro Source Files
  Autonomous Database makes it easy to access ORC, Parquet, or Avro data stored in object store using external tables. ORC, Parquet, and Avro sources have metadata embedded in them and the `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` procedure can utilize this metadata to simplify the creation of external tables.
- Query External Tables with Implicit Partitioning
  On Autonomous Database you can create implicit partitioned external tables from Hive style partitioned data or from simple folder partitioned data stored on your Cloud Object Store.
- Query External Tables with Partitioning Specified in Source Files
  If you want to query multiple data files in the Object Store as a single external table and the files can be represented as multiple logical partitions, it is highly recommended to use an external partitioned table. Using an external partitioned table preserves the logical partitioning of your data files for query access.

- Query External Partitioned Data (with Partitioning Clause)
  If you want to query multiple data files in the Object Store as a single external table and the files can be represented as multiple logical partitions, then it is highly recommended to use an external partitioned table. Using an external partitioned table preserves the logical partitioning of your data files for query access. Use the procedure `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` to create an external partitioned table.
- Query Data that Uses Table Hyperlink Source
  You can create an external table over an Autonomous Database Table Hyperlink.
- Query Hybrid Partitioned Data
  If you want to query internal data and multiple data files in the Object Store as single logical table you can use a hybrid partitioned table to represent the data as single object. Use the procedure `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE` to create a hybrid partitioned table.
- Query External Data Pump Dump Files
  You can also query Oracle Data Pump dump files in the Cloud by creating an external table using `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`.
- Query Big Data Service Hadoop (HDFS) Data from Autonomous Database
  You can create database links to Oracle Big Data Service from Autonomous Database.
- Query External Data with Data Catalog
  Oracle Cloud Infrastructure Data Catalog is the metadata management service for Oracle Cloud that helps you discover data and support data governance. It provides an inventory of assets, a business glossary, and a common metastore for data lakes.
- Query External Data with AWS Glue Data Catalog
  Autonomous Database supports a system for synchronizing with an Amazon AWS Glue Data Catalog instance.
- Query Apache Iceberg Tables
  Autonomous Database supports querying Apache Iceberg tables.
- Validate External Data
  To validate any external table, you can use the procedure `DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE`.
- Validate External Partitioned Data
  To validate an external partitioned table, you can use the procedure `DBMS_CLOUD.VALIDATE_EXTERNAL_PART_TABLE`. This procedure includes a parameter that lets you specify a specific partition to validate.
- Validate Hybrid Partitioned Data
  To validate a hybrid partitioned table, you can use the procedure `DBMS_CLOUD.VALIDATE_HYBRID_PART_TABLE`. This procedure includes a parameter that lets you specify a specific partition to validate.
- View Logs for Data Validation
  To validate an external table, use the procedures `DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE`, `DBMS_CLOUD.VALIDATE_EXTERNAL_PART_TABLE`, and `DBMS_CLOUD.VALIDATE_HYBRID_PART_TABLE`.

Query External Data

To query data in files in the Cloud, you need to first store your object storage credentials in your Autonomous Database, and then create an external table using the PL/SQL procedure `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`.

You can also use the procedure DBMS_CLOUD.CREATE_EXTERNAL_TABLE to query external data in attached file systems or in the local file system.

The procedure DBMS_CLOUD.CREATE_EXTERNAL_TABLE supports external files in the supported cloud object storage services, including:

- Oracle Cloud Infrastructure Object Storage
- Azure Blob Storage or Azure Data Lake Storage
- Amazon S3
- Amazon S3-Compatible, including: Oracle Cloud Infrastructure Object Storage, Google Cloud Storage, and Wasabi Hot Cloud Storage.
- GitHub Repository

The source file in this example, channels.txt, has the following data:

```
S,Direct Sales,Direct
T,Tele Sales,Direct
C,Catalog,Indirect
I,Internet,Indirect
P,Partners,Others
```

1. Store your object store credentials using the procedure DBMS_CLOUD.CREATE_CREDENTIAL.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'DEF_CRED_NAME',
       username => 'adb_user@example.com',
       password => 'password' );
   END;
   /
   ```

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials you can then use the same credential name for creating external tables.

   See CREATE_CREDENTIAL Procedure for information about the username and password parameters for different object storage services.

2. Create an external table on top of your source files using the procedure DBMS_CLOUD.CREATE_EXTERNAL_TABLE.

   The procedure DBMS_CLOUD.CREATE_EXTERNAL_TABLE supports external files in the supported cloud object storage services. The credential is a table level property; therefore, the external files must be on the same object store.

   For example:

   ```
   BEGIN
      DBMS_CLOUD.CREATE_EXTERNAL_TABLE(
       table_name =>'CHANNELS_EXT',
       credential_name =>'DEF_CRED_NAME',
   ```

```
    file_uri_list =>'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/channels.txt',
    format => json_object('delimiter' value ','),
    column_list => 'CHANNEL_ID VARCHAR2(2), CHANNEL_DESC VARCHAR2(20),
CHANNEL_CLASS VARCHAR2(20)' );
END;
/
```

The parameters are:

- `table_name`: is the external table name.
- `credential_name`: is the name of the credential created in the previous step. The `credential_name` parameter must conform to Oracle object naming conventions. See Database Object Naming Rules for more information.

  The `credential_name` parameter is not used when you specify a directory or a Table Hyperlink URL with `file_uri_list`. See Creating and Managing Directories on Autonomous Database and About Table Hyperlinks on Autonomous Database for more information.
- `file_uri_list`: is a comma delimited list of the source files you want to query.
- `format`: defines the options you can specify to describe the format of the source file.

  If the data in your source files is encrypted, decrypt the data by specifying the `format` parameter with the `encryption` option. See Decrypt Data While Importing from Object Storage for more information on decrypting data.
- `column_list`: is a comma delimited list of the column definitions in the source files.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

> **Note:**
>
> Autonomous Database supports a variety of source file formats, including compressed data formats. See DBMS_CLOUD Package Format Options and the `DBMS_CLOUD compression` format option to see the supported compression types.

You can now run queries on the external table you created in the previous step. For example:

```
SELECT count(*) FROM channels_ext;
```

By default the database expects all rows in the external data file to be valid and match both the target data type definitions as well as the format definition of the files. If there are any rows in the source files that do not match the format options you specified, the query reports an error. You can use `DBMS_CLOUD` parameters, like *rejectlimit,* to suppress these errors. As an alternative, you can also validate the external table you created to see the error messages and the rejected rows so that you can change your format options accordingly. See Validate External Data for more information.

For detailed information about the parameters, see CREATE_EXTERNAL_TABLE Procedure.

See DBMS_CLOUD URI Formats for more information on the supported cloud object storage services.

- External Table Metadata Columns
  The external table metadata helps you determine where data is coming from when you perform a query.

External Table Metadata Columns

The external table metadata helps you determine where data is coming from when you perform a query.

The external tables you create with `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`, `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE`, or `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE` include two invisible columns `file$path` and `file$name`. These columns help identify which file a record is coming from.

- `file$path`: Specifies the file path text up to the beginning of the object name.
- `file$name`: Specifies the object name, including all the text that follows the final "`/`".

For example:

```
SELECT genre_id, name, file$name, file$path FROM ext_genre
     WHERE rownum <= 2;

genre_id      name        file$name     file$path
--------      ---------   -----------   ----------------------
1             Action      genre.csv     https://objectstorage.us-
ashburn-1.oraclecloud.com/n/namespace-string/b/moviestream_gold/o/genre
2             Adventure   genre.csv     https://objectstorage.us-
ashburn-1.oraclecloud.com/n/namespace-string/b/moviestream_gold/o/genre
```

See Invisible Columns for more information on invisible columns.

Query External Data with ORC, Parquet, or Avro Source Files

Autonomous Database makes it easy to access ORC, Parquet, or Avro data stored in object store using external tables. ORC, Parquet, and Avro sources have metadata embedded in them and the `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` procedure can utilize this metadata to simplify the creation of external tables.

You don't need to know the structure of the data, `DBMS_CLOUD` can examine the file and convert ORC, Parquet, or Avro contents into the equivalent Oracle columns and data types. You only need to know the location of the data in object store, specify its type, ORC, Parquet, or Avro, and have credentials to access the source file on your object store.

> **Note:**
>
> The steps to use external tables are very similar for ORC, Parquet, and Avro. These steps show working with a Parquet format source file.

The source file in this example, `sales_extended.parquet`, contains Parquet format data. To query this file in Autonomous Database, do the following:

1. Store your object store credentials, to access the object store, using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'adb_user@example.com',
    password => 'password' );
END;
/
```

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials you can then use the same credential name for creating external tables.

   See CREATE_CREDENTIAL Procedure for information about the `username` and `password` parameters for different object storage services.

2. Create an external table for ORC, Parquet, or Avro on top of your source files using the procedure `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`.

   The procedure `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` supports external files in the supported cloud object storage services, including: Oracle Cloud Infrastructure Object Storage, Azure Blob Storage or Azure Data Lake Storage, Amazon S3, and Amazon S3-Compatible, including: Oracle Cloud Infrastructure Object Storage, Google Cloud Storage, and Wasabi Hot Cloud Storage. The credential is a table level property; therefore, the external files must be on the same object store.

   By default, the columns created in the external table automatically map their data types to Oracle data types for the fields found in the source files and the external table column names match the source field names.

```
BEGIN
   DBMS_CLOUD.CREATE_EXTERNAL_TABLE(
    table_name =>'sales_extended_ext',
    credential_name =>'DEF_CRED_NAME',
    file_uri_list =>'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-
string/b/bucketname/o/sales_extended.parquet',
    format =>  '{"type":"parquet", "schema": "first"}'
 );
END;
/
```

   The parameters are:

   - `table_name`: is the external table name.
   - `credential_name`: is the name of the credential created in the previous step. The `credential_name` parameter must conform to Oracle object naming conventions. See Database Object Naming Rules for more information.
   - `file_uri_list`: is a comma delimited list of the source files you want to query. The URI format for Dedicated Endpoints is supported in commercial (OC1) realms. See

Object Storage Dedicated Endpoints, and DBMS_CLOUD URI Formats for more information.

- `format`: defines the options to describe the format of the source file. For a Parquet file, use the `format` parameter to specify the `type parquet`. For an Avro file use the `format` parameter to specify the `type avro`. For an ORC file use the `format` parameter to specify the `type orc`.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

By default the `format schema` parameter is set and the columns and data types are derived automatically and the fields in the source match the external table columns by name. Source data types are converted to the external table column Oracle data types according to the `DBMS_CLOUD` mapping for ORC, Parquet, or Avro data types. The valid `schema` parameter values are:

- `first`: Analyze the schema of the first ORC, Parquet, or Avro file that `DBMS_CLOUD` finds in the specified `file_uri_list` (`first` is the default value for `schema`).
- `all`: Analyze all the schemas for all the ORC, Parquet, or Avro files found in the `file_uri_list`. Because these are simply files captured in an object store, there is no guarantee that each file's metadata is the same. For example, File1 may contain a field called "address", while File2 may be missing that field. Examining each file to derive the columns is a bit more expensive but may be required if the first file does not contain all the required fields.

> **Note:**
>
> If the `column_list` parameter is specified, then you provide the column names and data types for the external table and the `schema` value, if specified is ignored. Using `column_list` you can limit the columns in the external table. If `column_list` is not specified then the `schema` default value is `first`.

3. You can now run queries on the external table you created in the previous step:

```
DESC sales_extended_ext;
Name           Null? Type
-------------- ----- --------------
PROD_ID              NUMBER(10)
CUST_ID              NUMBER(10)
TIME_ID              VARCHAR2(32767)
CHANNEL_ID           NUMBER(10)
PROMO_ID             NUMBER(10)
QUANTITY_SOLD        NUMBER(10)
AMOUNT_SOLD          NUMBER(10,2)
GENDER               VARCHAR2(32767)
CITY                 VARCHAR2(32767)
STATE_PROVINCE       VARCHAR2(32767)
INCOME_LEVEL         VARCHAR2(32767)
```

```
SELECT prod_id, quantity_sold, gender, city, income_level
     FROM sales_extended_ext
     WHERE ROWNUM < 8;
```

```
     PROD_ID     QUANTITY_SOLD     GENDER    CITY              INCOME_LEVEL
1    13          1                 M         Adelaide          K: 250,000 - 299,999
2    13          1                 M         Dolores           L: 300,000 and above
3    13          1                 M         Cayuga            F: 110,000 - 129,999
4    13          1                 F         Bergen op Zoom    C: 50,000 - 69,999
5    13          1                 F         Neuss             J: 190,000 - 249,999
6    13          1                 F         Darwin            F: 110,000 - 129,999
7    13          1                 M         Sabadell          K:250,000 - 299,999
```

This query shows values for rows in the external table. If you want to query this data frequently, after examining the data you can load it into a table with `DBMS_CLOUD.COPY_DATA`.

See CREATE_EXTERNAL_TABLE Procedure for Avro, ORC, or Parquet Files and COPY_DATA Procedure for Avro, ORC, or Parquet Files for more information.

See DBMS_CLOUD URI Formats for information on supported cloud object storage services.

- Query External Data with ORC, Parquet, or Avro Source Files and Explicitly Set the Text Column Size
  When you access ORC, Parquet, or Avro data stored in object store using external tables on Autonomous Database, you can either automatically or explicitly set the maximum size of text columns.

Query External Data with ORC, Parquet, or Avro Source Files and Explicitly Set the Text Column Size

When you access ORC, Parquet, or Avro data stored in object store using external tables on Autonomous Database, you can either automatically or explicitly set the maximum size of text columns.

By default the text column size is set based on the value of `MAX_STRING_SIZE`.

The source file in this example, `sales_extended.parquet`, contains Parquet format data. To query this file in Autonomous Database and set the maximum text column size, do the following:

1. Store your object store credentials, to access the object store, using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'adb_user@example.com',
    password => 'password' );
END;
/
```

Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials you can then use the same credential name for creating external tables.

See CREATE_CREDENTIAL Procedure for information about the `username` and `password` parameters for different object storage services.

2. Create an external table for ORC, Parquet, or Avro on top of your source files using the procedure `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` and specify the `maxvarchar` format parameter.

   The procedure `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` supports external files in the supported cloud object storage services, including: Oracle Cloud Infrastructure Object Storage, Azure Blob Storage or Azure Data Lake Storage, Amazon S3, and Amazon S3-Compatible, including: Oracle Cloud Infrastructure Object Storage, Google Cloud Storage, and Wasabi Hot Cloud Storage. The credential is a table level property; therefore, the external files must be on the same object store.

   By default, the columns created in the external table automatically map their data types to Oracle data types for the fields found in the source files and the external table column names match the source field names.

```
BEGIN
   DBMS_CLOUD.CREATE_EXTERNAL_TABLE(
    table_name =>'sales_extended_ext',
    credential_name =>'DEF_CRED_NAME',
    file_uri_list =>'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-
string/b/bucketname/o/sales_extended.parquet',
    format =>  '{"type":"parquet", "schema": "first", "maxvarchar":"extended"}'
 );
END;
/
```

   The parameters are:

   - `table_name`: is the external table name.
   - `credential_name`: is the name of the credential created in the previous step. The `credential_name` parameter must conform to Oracle object naming conventions. See Database Object Naming Rules for more information.
   - `file_uri_list`: is a comma delimited list of the source files you want to query. The URI format for Dedicated Endpoints is supported in commercial (OC1) realms. See Object Storage Dedicated Endpoints, and DBMS_CLOUD URI Formats for more information.
   - `format`: defines the options to describe the format of the source file. For a Parquet file, use the `format` parameter to specify the `type parquet`. For an Avro file use the `format` parameter to specify the `type avro`. For an ORC file use the `format` parameter to specify the `type orc`.

     The option `maxvarchar` with value `extended` specifies that text columns are created as `varchar(32767)` on an Autonomous Database instance with extended string size. The possible values are `standard` with `varchar(4000)`, `extended` with `varchar(32767)`, and `auto`. The default `maxvarchar` value is `auto`. With this value, the maximum text size is based on the value of `MAX_STRING_SIZE`.

   In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

   By default the `format schema` parameter is set and the columns and data types are derived automatically and the fields in the source match the external table columns by name.

Source data types are converted to the external table column Oracle data types according to the `DBMS_CLOUD` mapping for ORC, Parquet, or Avro data types. The valid `schema` parameter values are:

- `first`: Analyze the schema of the first ORC, Parquet, or Avro file that `DBMS_CLOUD` finds in the specified `file_uri_list` (`first` is the default value for `schema`).
- `all`: Analyze all the schemas for all the ORC, Parquet, or Avro files found in the `file_uri_list`. Because these are simply files captured in an object store, there is no guarantee that each file's metadata is the same. For example, File1 may contain a field called "address", while File2 may be missing that field. Examining each file to derive the columns is a bit more expensive but may be required if the first file does not contain all the required fields.

> **Note:**
>
> If the `column_list` parameter is specified, then you provide the column names and data types for the external table and the `schema` value, if specified is ignored. Using `column_list` you can limit the columns in the external table. If `column_list` is not specified then the `schema` default value is `first`.

3. You can now run queries on the external table you created in the previous step:

```
DESC sales_extended_ext;
Name           Null? Type
-------------- ----- --------------
PROD_ID              NUMBER(10)
CUST_ID              NUMBER(10)
TIME_ID              VARCHAR2(32767)
CHANNEL_ID           NUMBER(10)
PROMO_ID             NUMBER(10)
QUANTITY_SOLD        NUMBER(10)
AMOUNT_SOLD          NUMBER(10,2)
GENDER               VARCHAR2(32767)
CITY                 VARCHAR2(32767)
STATE_PROVINCE       VARCHAR2(32767)
INCOME_LEVEL         VARCHAR2(32767)
```

```
SELECT prod_id, quantity_sold, gender, city, income_level
     FROM sales_extended_ext
     WHERE ROWNUM < 8;

     PROD_ID     QUANTITY_SOLD     GENDER    CITY              INCOME_LEVEL
1    13          1                 M         Adelaide          K: 250,000 - 299,999
2    13          1                 M         Dolores           L: 300,000 and above
3    13          1                 M         Cayuga            F: 110,000 - 129,999
4    13          1                 F         Bergen op Zoom    C: 50,000 - 69,999
5    13          1                 F         Neuss             J: 190,000 - 249,999
6    13          1                 F         Darwin            F: 110,000 - 129,999
7    13          1                 M         Sabadell          K:250,000 - 299,999
```

This query shows values for rows in the external table. If you want to query this data frequently, after examining the data you can load it into a table with `DBMS_CLOUD.COPY_DATA`.

If you specify the `format` option `maxvarchar` as `standard`, the `varchar2()` text columns are created with size `4000`. For example:

```
BEGIN
   DBMS_CLOUD.CREATE_EXTERNAL_TABLE(
    table_name =>'sales_extended_ext',
    credential_name =>'DEF_CRED_NAME',
    file_uri_list =>'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-
string/b/bucketname/o/sales_extended.parquet',
    format =>  '{"type":"parquet", "schema": "first", "maxvarchar":"standard"}'
 );
END;
/
```

The `format` parameter with the option `maxvarchar` set to the value `standard` specifies that text columns are created as `varchar(4000)`.

```
DESC sales_extended_ext;
Name           Null? Type
-------------- ----- --------------
PROD_ID              NUMBER(10)
CUST_ID              NUMBER(10)
TIME_ID              VARCHAR2(4000)
CHANNEL_ID           NUMBER(10)
PROMO_ID             NUMBER(10)
QUANTITY_SOLD        NUMBER(10)
AMOUNT_SOLD          NUMBER(10,2)
GENDER               VARCHAR2(4000)
CITY                 VARCHAR2(4000)
STATE_PROVINCE       VARCHAR2(4000)
INCOME_LEVEL         VARCHAR2(4000)
```

See CREATE_EXTERNAL_TABLE Procedure for Avro, ORC, or Parquet Files and COPY_DATA Procedure for Avro, ORC, or Parquet Files for more information.

See DBMS_CLOUD URI Formats for information on supported cloud object storage services.

Query External Tables with Implicit Partitioning

On Autonomous Database you can create implicit partitioned external tables from Hive style partitioned data or from simple folder partitioned data stored on your Cloud Object Store.

- About External Tables with Implicit Partitioning
  On Autonomous Database use the `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` procedure to create implicit partitioned external tables from Hive style partitioned data or from simple folder partitioned data stored on your Cloud Object Store.
- Query External Implicit Partitioned Data with Hive Format Source File Organization
  Use the `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` procedure to create an implicit partitioned external table from data in Object Store that was generated from Hive data.

- Query External Implicit Partitioned Non-Hive Style Data
  Use the `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` procedure to create an implicit partitioned external table from data in object storage that has been generated from non-Hive data.

About External Tables with Implicit Partitioning

On Autonomous Database use the `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` procedure to create implicit partitioned external tables from Hive style partitioned data or from simple folder partitioned data stored on your Cloud Object Store.

Use the `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` procedure to create implicit partitioned external tables. By passing the appropriate options to this procedure, the partitions are derived from the source data. The partitioned external tables support runtime discovery of partition columns and their values. Runtime discovery of changes in the underlying object store structure, such as adding or removing objects, simplifies the maintenance process by eliminating the need for additional synchronization procedures required by `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE`. This ensures that the data is up to date at query runtime.

With implicit partitioning, Autonomous Database automatically determines the columns a table is "partitioned on" based on the object storage source hierarchical file structure. There is no need to explicitly declare a partitioning scheme. Implicit partitioning provides partition-table-like performance benefits without the need to explicitly define a partitioned external table using the `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` procedure.

Implicit partitioned external tables support the following naming styles for partitioned objects in object store:

- Hive Naming Format: Data in object storage that has been generated from Hive has the following format.

  ```
  OBJBUCKET/<table>/<partcol1>=<value1>/<partcol2>=<value2>/file.parquet
  ```

  For an example, see Query External Implicit Partitioned Data with Hive Format Source File Organization
- Non-Hive "Pathtail" Naming Format: A second format that is often found in data lakes is similar to the hive format, but the folder name does not include the partition columns.

  ```
  OBJBUCKET/<table>/<value1>/<value2>/file.parquet
  ```

  For an example, see Query External Implicit Partitioned Non-Hive Style Data.

For more information on naming formats, see About External Tables with Source File Partitioning.

Using the `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` procedure, implicit partitioned external tables can be created in the following ways:

- Set the partition type to hive with a given list of partition columns
  In this case, `implicit_partition_type` is set to `hive` and `implicit_partition_columns` provides a list of partition columns.

  For example:

  ```
  DBMS_CLOUD.CREATE_EXTERNAL_TABLE (
     table_name         => 'mysales',
     credential_name    => 'mycredential',
     file_uri_list      => 'https://swiftobjectstorage.us-
  ```

```
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/*.parquet',
   column_list       => 'product varchar2(100), units number, country
varchar2(100),year (number), month varchar2(2)',
   format            => '{"type":"parquet",
                         "implicit_partition_type":"hive",
                         "implicit_partition_columns":
["country","year","month"]}');
```

- Set the partition type to hive without providing a list of partition columns
  In this case, `implicit_partition_type` is set to `hive` and `implicit_partition_columns` is not provided. The partition columns are automatically detected by searching for '=' in the path specified by `file_uri_list`. The `column_name` is to the left of the '=' and the value is on the right side. If a `column_name` is not found in the path, an error is thrown.

  For example:

```
DBMS_CLOUD.CREATE_EXTERNAL_TABLE (
   table_name        => 'mysales',
   credential_name   => 'mycredential',
   file_uri_list     => 'https://swiftobjectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/*.parquet',
   column_list       => 'product varchar2(100), units number, country
varchar2(100),year (number), month varchar2(2)',
   format            => '{"type":"parquet",
"implicit_partition_type":"hive"');
```

- Provide a list of partition columns without specifying the type
  In this case, `implicit_partition_type` is not set and `implicit_partition_columns` provides a list of columns.

  For example:

```
DBMS_CLOUD.CREATE_EXTERNAL_TABLE (
   table_name        => 'mysales',
   credential_name   => 'mycredential',
   file_uri_list     => 'https://swiftobjectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/*.parquet',
   column_list       => 'product varchar2(100), units number, country
varchar2(100),year (number), month varchar2(2)',
   format            => '{"type":"parquet",
                         "implicit_partition_columns":
["country","year","month"]}');
```

See CREATE_EXTERNAL_TABLE Procedure for further information.

Query External Implicit Partitioned Data with Hive Format Source File Organization

Use the `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` procedure to create an implicit partitioned external table from data in Object Store that was generated from Hive data.

The sample source files in the follow example, use this naming format:

```
OBJBUCKET/<table>/<partcol1>=<value1>/<partcol2>=<value2>/file.parquet
```

Consider the following sample source files:

```
OBJBUCKET/sales/country=USA/year=2024/month=01/sales-2024-01.parquet
```

```
OBJBUCKET/sales/country=USA/year=2024/month=02/sales-2024-02.parquet

OBJBUCKET/sales/country=USA/year=2024/month=02/sales-2024-03.parquet
```

To create an implicit partitioned external table with data stored in this sample Hive format, do the following:

1. Store Object Store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL (
       credential_name => 'DEF_CRED_NAME',
       username => 'adb_user@example.com',
       password => 'password' );
   END;
   /
   ```

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials you can then use the same credential name for creating external tables.

   See CREATE_CREDENTIAL Procedure for information about the `username` and `password` parameters for different object storage services.

2. Create an implicit partitioned external table on top of your source files using the procedure `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE`.

   In this example, since the format option `implicit_partition_columns` is not provided, partition columns are automatically detected and implicit partitioning is enabled by setting the format option `implicit_partition_type` to `hive`.
   However, if a list of columns is given, then those columns are used as partition columns and Autonomous Database doesn't try to discover the columns.

   To detect the partition columns Autonomous Database starts searching from the beginning of the path, specified by `file_uri_list`, for '='. When found, the left part of '=' up to the last '/' is taken as column (e.g. "country") and the right part of '=', until the first '/' as value (e.g. "USA"). The search continues for '=' until after the '/' that follows the first partition value, detecting the second '=' in the path, and so on.

   ```
   DBMS_CLOUD.CREATE_EXTERNAL_TABLE (
      table_name        => 'mysales',
      credential_name   => 'DEF_CRED_NAME',
      file_uri_list     => 'https://swiftobjectstorage.us-
   phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/*.parquet',
      column_list       => 'product varchar2(100), units number, country
   varchar2(100),year (number), month varchar2(2)',
      format            => '{"type":"parquet",
   "implicit_partition_type":"hive"');
   ```

   The parameters are:

- `table_name`: is the external table name.
- `credential_name`: is the name of the credential created in the previous step.
- `file_uri_list`: is a comma-delimited list of source file URIs. There are two options for this list:
  - Specify a comma-delimited list of individual file URIs without wildcarding.
  - Specify a single file URI with wildcards, where the wildcards can only be after the last slash "/". The character "*" can be used as the wildcard for multiple characters, the character "?" can be used as the wildcard for a single character.
- `column_list`: is a comma delimited list of column names and data types for the external table. The list includes the columns inside the data file and those derived from the object name (from names in the file path specified by `file_uri_list`).

  The `column_list` is not required when the data files are structured files (Parquet, Avro, or ORC).
- `format`: defines the options you can specify to describe the format of the source file. The `implicit_partition_type` option specifies the data format type as hive.

  If the data in your source file is encrypted, decrypt the data by specifying the `encryption` format option. See Decrypt Data While Importing from Object Storage for more information on decrypting data.

  See DBMS_CLOUD Package Format Options for more information.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

See CREATE_EXTERNAL_TABLE Procedure for detailed information about the parameters. See DBMS_CLOUD URI Formats for more information on the supported cloud object storage services.

3. You can now run queries on the external partitioned table you created in the previous step.

   Your Autonomous Database takes advantage of the partitioning information of your external partitioned table, ensuring that the query only accesses the relevant data files in Object Store.

   For example:

   ```
   SELECT product, units FROM mysales WHERE year='2024' and month='02'
   ```

   This SQL statement only queries data in the partition for month 02 of year 2024.

Query External Implicit Partitioned Non-Hive Style Data

Use the `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` procedure to create an implicit partitioned external table from data in object storage that has been generated from non-Hive data.

The sample source files in the follow example, use this naming format:

```
OBJBUCKET/<table>/<value1>/<value2>/file.parquet
```

Consider the following sample source files:

```
OBJBUCKET/sales/USA/2024/01/sales-2024-01.parquet

OBJBUCKET/sales/USA/2024/01/sales-2024-02.parquet
```

```
OBJBUCKET/sales/USA/2024/01/sales-2024-03.parquet
```

To create an implicit partitioned external table with data stored in this sample Hive format, do the following:

1. Store Object Store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL (
       credential_name => 'DEF_CRED_NAME',
       username => 'adb_user@example.com',
       password => 'password' );
   END;
   /
   ```

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials you can then use the same credential name for creating external tables.

   See CREATE_CREDENTIAL Procedure for information about the `username` and `password` parameters for different object storage services.

2. Create an implicit partitioned external table on top of your source files using the procedure `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE`.

   In this example, implicit partitioning is enabled by providing the partition columns with the format option `implicit_partition_columns`. Since the folder name does not include the partition columns, the partition values in the path, specified by `file_uri_list`, can only be discovered if an explicit list of columns is given with the format option `implicit_partition_columns`. In order to detect the respective column values in the path the order of columns must be the same as the order of values.

   The partition columns are automatically detected.

   ```
   DBMS_CLOUD.CREATE_EXTERNAL_TABLE (
      table_name        => 'mysales',
      credential_name   => 'DEF_CRED_NAME',
      file_uri_list     => 'https://swiftobjectstorage.us-
   phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/*.parquet',
      column_list       => 'product varchar2(100), units number, country
   varchar2(100),year (number), month varchar2(2)',
      format            => '{"type":"parquet", "implicit_partition_columns":
   ["country","year","month"]}');
   ```

   The parameters are:

   - `table_name`: is the external table name.
   - `credential_name`: is the name of the credential created in the previous step.
   - `file_uri_list`: is a comma-delimited list of source file URIs. There are two options for this list:

  - Specify a comma-delimited list of individual file URIs without wildcarding.
  - Specify a single file URI with wildcards, where the wildcards can only be after the last slash "/". The character "*" can be used as the wildcard for multiple characters, the character "?" can be used as the wildcard for a single character.

- `column_list`: is a comma delimited list of column names and data types for the external table. The list includes the columns inside the data file and those derived from the object name (from names in the file path specified by `file_uri_list`).

  The `column_list` is not required when the data files are structured files (Parquet, Avro, or ORC).

- `format`: defines the options you can specify to describe the format of the source file. The `implicit_partition_type` option is unset. Since `implicit_partition_columns` is given, the type is automatically detected as non-hive.

  If the data in your source file is encrypted, decrypt the data by specifying the `encryption` format option. See Decrypt Data While Importing from Object Storage for more information on decrypting data.

  See DBMS_CLOUD Package Format Options for more information.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

See CREATE_EXTERNAL_TABLE Procedure for detailed information about the parameters.

See DBMS_CLOUD URI Formats for more information on the supported cloud object storage services.

3. You can now run queries on the external partitioned table you created in the previous step.

   Your Autonomous Database takes advantage of the partitioning information of your external partitioned table, ensuring that the query only accesses the relevant data files in Object Store.

   For example:

```
SELECT product, units FROM mysales WHERE year='2024'
```

   This SQL statement only queries data in the partition for year 2024.

Query External Tables with Partitioning Specified in Source Files

If you want to query multiple data files in the Object Store as a single external table and the files can be represented as multiple logical partitions, it is highly recommended to use an external partitioned table. Using an external partitioned table preserves the logical partitioning of your data files for query access.

Using partitioned external tables has the potential to dramatically improve query performance by only accessing the data required for the query. For example, you may have two years of daily partitions stored in separate objects on Cloud Object Store. When you use partitioned external tables, a query for a single day only needs to access that day's source data. When you use partitioned external tables the database automatically partition prunes, and in this example only needs to scan a very small fraction of the data.

There are two ways to create an external partitioned table with the `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` procedure:

- Using the `file_url_path` value in combination with the `format` parameter: Autonomous Database analyzes Cloud Object Store file path information supplied with this parameter to determine the partition columns and data types (or you can manually specify the partition columns and data types).

  This type of partitioning provides a synchronization routine to handle changes when external partition files are added or removed.

- Using the `partitioning_clause` parameter: Autonomous Database uses the explicit partitioning clause you supply to create an external partitioned table.

  This type of partitioning does not support a synchronization routine.

  See Query External Partitioned Data (with Partitioning Clause) for a description of this type of external table.

- About External Tables with Source File Partitioning
  On Autonomous Database you can create partitioned external tables from Hive style partitioned data or from simple folder partitioned data stored on your Cloud Object Store.
- Query External Partitioned Data with Hive Format Source File Organization
  Use `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` to create an external partitioned table and generate the partitioning information from the Cloud Object Store file path.
- Query External Partitioned Data with Folder Format Source File Organization
  Use `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` to create an external partitioned table and generate the partitioning information from the Cloud Object Store file path.
- Refresh External Partitioned Tables with Updated or Deleted Source Files
  You can use `DBMS_CLOUD.SYNC_EXTERNAL_PART_TABLE` to refresh an external partitioned table. Use this procedure when new partitions are added or when partitions are removed from the object store source.

About External Tables with Source File Partitioning

On Autonomous Database you can create partitioned external tables from Hive style partitioned data or from simple folder partitioned data stored on your Cloud Object Store.

Using source file partitioning, instead of supplying a complete partition specification the procedure derives partitioning information from the file path for certain file patterns. For example, consider the following data file specifications:

- Hive style: for example: `sales/country=USA/year=2020/month=01/file1.csv`
- Simple folder partitioning style: for example: `sales/USA/2020/01/file1.parquet`

Using one of these common partitioning formats greatly simplifies both the creation and management of partitioned external tables. In addition, even though partition columns may not appear in the data file, they can still be queried using SQL. Partitioning data also improves query performance by dramatically reducing the amount of data scanned. In this example, when you query ‘USA’ data, the query can skip scanning the files for other countries.

Hive Format Partitioned Data in Cloud Object Store

Hive offers a standard metadata format for big data processing engines. Partitioned data in Cloud Object Store that is generated in Hive format is represented in a `folder/subfolder` format. For example, on Cloud Object Store a Hive format data file is stored as follows:

table`/`*partition1*`=`*partition1_value*`/`*partition2*`=`*partition2_value*`/data_file.csv`

Files saved in Hive partitioned format provide partition information in the data file path name. The data file path name includes information about the object contents, including partition column names and partition column values (the data file does not include the partition columns and their associated values).

For example, consider an external partitioned `SALES` table created from Hive format data on Cloud Object Store:

```
.../sales/country=USA/year=2020/month=01/file1.csv
.../sales/country=USA/year=2020/month=01/file2.csv
.../sales/country=USA/year=2020/month=02/file3.csv
.../sales/country=USA/year=2020/month=03/file1.csv
.../sales/country=FRA/year=2020/month=03/file1.csv
```

The Hive format partition information shows the data files in Cloud Object Store are partitioned by `country`, `year`, and `month` and the values for these partition columns are also specified within the Hive format path name for each data file (the path name includes values for the partitioned columns: `country`, `year`, and `month`).

The column names in the path will be used by the API to simplify the table definition.

Simple Folder Format Partitioned Data in Cloud Object Store

Partitioned data in Cloud Object Store that is generated in folder format is represented in a `folder/subfolder` format, similar to Hive format partitioned data, but the information in the path shows the column values and does not include the column names. Also, with folder format partitioned data the partition order specified in the object name is significant, and must match the order in the table columns.

For example, on Cloud Object Store a folder format data file is stored as follows:

table/partition1_value/partition2_value/.parquet*

The path includes both partition column values, in partition column order, and the data files. Autonomous Database allows you to create an external partitioned table from folder format data and you can perform a query using the specified partitions.

Files saved in folder partitioned format provide the data partition column values in the file name. Unlike Hive, the paths do not include the column name, therefore the column names must be provided. The order of partition columns is important and the order in the file name for column partition names must match the order in the `partition_columns` parameter.

About Querying Partitioned Data in Cloud Object Store

When you query external partitioned data in Hive format, the query engine understands and utilizes the partitioning information from the file path name. For example, consider an external partitioned `SALES` table where the source file, `sales/country=USA/year=2020/month=02/file3.csv` on Object Store includes the following sales data:

```
tents, 291
canoes, 22
backpacks, 378
```

The `country` values in the path name, and the time period values for `month` and `year` are not specified as columns within the data file. The partition column values are specified only in the path name with values shown: USA, 2020, and 02. After you create an external partitioned

table with this data file you can use the partition columns and their values when you run a query on the external partitioned table.

For example:

```
SELECT year, month, product, units
     FROM SALES WHERE year='2020' AND month='02' AND country='USA'
```

The benefit of creating an external partitioned table with data generated as Hive format partitioned data is that the query engine is optimized to partition prune the data to select the correct partition and the query only selects data from one partition and only needs to search a single data file. Thus, the query would only require a scan of the `file3.csv` file (`/sales/country=USA/year=2020/month=02/file3.csv`). For large amounts of data, such partition pruning can provide significant performance improvements.

Using standard Oracle Database external tables, the partition column must be available as a column within the data file to use it for queries or partition definitions. Without the special handling that is available with external partitioned tables on Autonomous Database, this would be a problem if you want to use data stored in Hive format on Cloud Object Store, as you would need to regenerate the data files to include the partition as a column in the data file.

About Creating Partitioned External Tables

When you use unstructured data stored in Hive format on Cloud Object Store and you create an external partitioned table, the columns and their types cannot be derived from the source file. Thus, the columns and their data types must be specified with the `column_list` parameter. To create the partitioned external tables, use the procedure `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` to specify the partition columns and their types as follows:

- The root for the file list is specified in the path name with the `file_uri_list` parameter. For example, `http://.../sales/*`
- The column names and data types are specified with the `column_list` parameter.
- The option `partition_columns` in the `format` parameter specifies the partition columns.
- The generated DLL includes the columns specified in the path name.

For this example, when the external table is created the `country`, `year`, and `month` columns are added in the `column_list` parameter. The external table is created with the `country`, `year`, and `month` columns, which are not in the data files, and list partitions are created enabling partition pruning.

When you use structured data, such as Parquet, Avro, or ORC files stored in folder format on Cloud Object Store, the columns and their data types are known. and you do not need to specify the column list as is required with unstructured data. To create the partitioned external tables, use the procedure `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` to specify the partition columns and their types as follows:

- The root for the file list is specified in the path name with the `file_uri_list` parameter. For example, `http://.../sales/*`
- The `column_list` parameter is not required for structured files. If you do not specify the column list, you must define the partition columns and their data types when you create the external partitioned table. Use the option `partition_columns` in the `format` parameter to specify the partition columns and their data types.
- The generated DLL includes the columns specified in the path name.

See Query External Partitioned Data with Hive Format Source File Organization and Query External Partitioned Data with Folder Format Source File Organization for complete examples.

- External Partitioning: CSV Source Files with Hive-style Folders
  Shows how to create external partitioned tables with CSV source files stored on Cloud Object Store in Hive-style folders.
- External Partitioning: CSV Source Files with Simple Folders
  Shows how to create external partitioned tables with CSV source files stored on Cloud Object Store in simple folder format.
- External Partitioning: Parquet Source Files with Hive-style Folders
  Shows how to create external partitioned tables with Parquet source files stored on Cloud Object Store in Hive-style folders.
- External Partitioning: Parquet with Simple Folders
  Shows how to create external partitioned tables with Parquet source files stored on Cloud Object Store in simple folder format.

External Partitioning: CSV Source Files with Hive-style Folders

Shows how to create external partitioned tables with CSV source files stored on Cloud Object Store in Hive-style folders.

Source file list:

```
.../sales/country=USA/year=2020/month=01/file1.csv
.../sales/country=USA/year=2020/month=01/file2.csv
.../sales/country=USA/year=2020/month=02/file3.csv
.../sales/country=USA/year=2020/month=03/file1.csv
.../sales/country=FRA/year=2020/month=03/file1.csv
```

API:

```
DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE (
   table_name        => 'mysales',
   credential_name   => 'mycredential',
   file_uri_list     => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/.../sales/*.csv',
   column_list       => 'product varchar2(100), units number, country
varchar2(100), year number, month varchar2(2)',
   field_list        => 'product, units', --[Because country, year and month
are not in the file, they are not listed in the field list]
   format            => '{"type":"csv","partition_columns":["country",
"year", "month"]}');
```

> **Note:**
>
> The `partition_columns` in the `format` parameter must match the column names found in the path (for example, the `country` column matches "`country=…`")

External Partitioning: CSV Source Files with Simple Folders

Shows how to create external partitioned tables with CSV source files stored on Cloud Object Store in simple folder format.

Source file list:

```
.../sales/USA/2020/01/file1.csv
.../sales/USA/2020/01/file2.csv
.../sales/USA/2020/02/file3.csv
.../sales/USA/2020/03/file1.csv
.../sales/FRA/2020/03/file1.csv
```

API:

```
DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE(
   table_name        => 'mysales',
   credential_name   => 'mycredential',
   file_uri_list     =>  'https://objectstorage.us-
phoenix-1.oraclecloud.com/.../sales/*.csv',
   column_list       => 'product varchar2(100), units number, country
varchar2(100), year number, month varchar2(2)',
   field_list        => 'product, units', --[Because country, year and month
are not in the file, they are not listed in the field list]
   format            => '{"type":"csv","partition_columns":["country",
"year", "month"]}');
```

> **Note:**
>
> The API call is the same as in the previous example, but the order of the `partition_columns` in the `format` parameter is significant because the column name is not in the file path.

External Partitioning: Parquet Source Files with Hive-style Folders

Shows how to create external partitioned tables with Parquet source files stored on Cloud Object Store in Hive-style folders.

Source file list:

```
.../sales/USA/2020/01/file1.parquet
.../sales/USA/2020/01/file2.parquet
.../sales/USA/2020/02/file3.parquet
.../sales/USA/2020/03/file1.parquet
.../sales/FRA/2020/03/file1.parquet
```

API:

```
DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE (
   table_name            => 'mysales',
     credential_name     => 'mycredential',
     file_uri_list       => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/.../sales/*.parquet',
     format              =>
         json_object(    'type' value 'parquet',
               'schema' value 'first',
               'partition_columns' value
```

```
                  json_array(
                      json_object('name' value 'country', 'type' value
'varchar2(100)'),
                      json_object('name' value 'year', 'type' value 'number'),
                      json_object('name' value 'month', 'type' value
'varchar2(2)')
                  )
          )

);
```

> **Note:**
>
> The `column_list` parameter is not specified. As shown, for each partition column specify both the name and data type in the `format` parameter `partition_columns`.

External Partitioning: Parquet with Simple Folders

Shows how to create external partitioned tables with Parquet source files stored on Cloud Object Store in simple folder format.

Source file list:

```
.../sales/USA/2020/01/file1.parquet
.../sales/USA/2020/01/file2.parquet
.../sales/USA/2020/02/file3.parquet
.../sales/USA/2020/03/file1.parquet
.../sales/FRA/2020/03/file1.parquet
```

API:

```
DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE (
   table_name        => 'mysales',
   credential_name   => 'mycredential',
   file_uri_list     => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/.../sales/*.parquet',
   format            =>
     json_object(    'type' value 'parquet',
                  'schema' value 'first',
           'partition_columns' value
             json_array(
                json_object('name' value 'country', 'type' value
'varchar2(100)'),
                json_object('name' value 'year', 'type' value 'number'),
                json_object('name' value 'month', 'type' value 'varchar2(2)')
             )
          )

);
```

> **Note:**
>
> The `column_list` parameter is not specified. You must include both the name and data type for the partition columns. In addition, the order of the `partition_columns` in the format clause matters because the column name is not in the file path.

Query External Partitioned Data with Hive Format Source File Organization

Use `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` to create an external partitioned table and generate the partitioning information from the Cloud Object Store file path.

Consider the following sample source files in Object Store:

```
custsales/month=2019-01/custsales-2019-01.csv

custsales/month=2019-02/custsales-2019-02.csv

custsales/month=2019-03/custsales-2019-03.csv
```

With this naming, the values for `month` are captured within the object name.

To create a partitioned external table with data stored in this sample Hive format, do the following:

1. Store Object Store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL (
       credential_name => 'DEF_CRED_NAME',
       username => 'adb_user@example.com',
       password => 'password' );
   END;
   /
   ```

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials you can then use the same credential name for creating external tables.

   See CREATE_CREDENTIAL Procedure for information about the `username` and `password` parameters for different object storage services.

2. Create an external partitioned table on top of your source files using the procedure `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE`.

   The procedure `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` supports external partitioned files in the supported cloud object storage services. The credential is a table level property; therefore, the external files must all be on the same cloud object store.

For example:

```
BEGIN
    DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE(
        TABLE_NAME => 'sales_sample',
        CREDENTIAL_NAME => 'DEF_CRED_NAME',
        FILE_URI_LIST => 'https://objectstorage.us-
ashburn-1.oraclecloud.com/n/namespace-string/b/moviestream_landing/o/
sales_sample/*.parquet',
        FORMAT => '{"type":"parquet", "schema":
"first","partition_columns":[{"name":"month","type":"varchar2(100)"}]}');
END;
/
```

The parameters are:

- `table_name`: is the external table name.
- `credential_name`: is the name of the credential created in the previous step.
- `file_uri_list`: is a comma-delimited list of source file URIs. There are two options for this list:
  - Specify a comma-delimited list of individual file URIs without wildcarding.
  - Specify a single file URI with wildcards, where the wildcards can only be after the last slash "/". The character "*" can be used as the wildcard for multiple characters, the character "?" can be used as the wildcard for a single character.
- `column_list`: is a comma delimited list of column names and data types for the external table. The list includes the columns inside the data file and those derived from the object name (from names in the file path).

  The `column_list` is not required when the data files are structured files (Parquet, Avro, or ORC).
- `format`: defines the options you can specify to describe the format of the source file. The `partition_columns format` parameter specifies the names of the partition columns.

  If the data in your source file is encrypted, decrypt the data by specifying the `encryption` format option. See Decrypt Data While Importing from Object Storage for more information on decrypting data.

  See DBMS_CLOUD Package Format Options for more information.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

The `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` call would result in the following table definition:

```
CREATE TABLE "ADMIN"."SALES_SAMPLE"
   ( "DAY_ID" TIMESTAMP (6),
     "GENRE_ID" NUMBER(19,0),
     "MOVIE_ID" NUMBER(19,0),
     "CUST_ID" NUMBER(19,0),
     "APP" VARCHAR2(4000 BYTE) COLLATE "USING_NLS_COMP",
     "DEVICE" VARCHAR2(4000 BYTE) COLLATE "USING_NLS_COMP",
     "OS" VARCHAR2(4000 BYTE) COLLATE "USING_NLS_COMP",
```

```
      "PAYMENT_METHOD" VARCHAR2(4000 BYTE) COLLATE "USING_NLS_COMP",
      "LIST_PRICE" BINARY_DOUBLE,
      "DISCOUNT_TYPE" VARCHAR2(4000 BYTE) COLLATE "USING_NLS_COMP",
      "DISCOUNT_PERCENT" BINARY_DOUBLE,
      "ACTUAL_PRICE" BINARY_DOUBLE,
      "MONTH" VARCHAR2(100 BYTE) COLLATE "USING_NLS_COMP"
   ) DEFAULT COLLATION "USING_NLS_COMP"
   ORGANIZATION EXTERNAL
    ( TYPE ORACLE_BIGDATA
      DEFAULT DIRECTORY "DATA_PUMP_DIR"
      ACCESS PARAMETERS
      ( com.oracle.bigdata.fileformat=parquet
com.oracle.bigdata.filename.columns=["month"]
com.oracle.bigdata.file_uri_list="https://objectstorage.us-
ashburn-1.oraclecloud.com/n/namespace-string/b/moviestream_landing/o/
sales_sample/*.parquet"
com.oracle.bigdata.credential.schema="ADMIN"
com.oracle.bigdata.credential.name=CRED_OCI
com.oracle.bigdata.trimspaces=notrim
  )
    )
   REJECT LIMIT 0
  PARTITION BY LIST ("MONTH")
 (PARTITION "P1" VALUES (('2019-01'))
      LOCATION
       ( 'https://objectstorage.us-ashburn-1.oraclecloud.com/n/namespace-
string/b/moviestream_landing/o/sales_sample/month=2019-01/*.parquet'
       ),
 PARTITION "P2" VALUES (('2019-02'))
      LOCATION
       ( 'https://objectstorage.us-ashburn-1.oraclecloud.com/n/namespace-
string/b/moviestream_landing/o/sales_sample/month=2019-02/*.parquet'
       ))
  PARALLEL ;
```

See CREATE_EXTERNAL_PART_TABLE Procedure for detailed information about the parameters.

See DBMS_CLOUD URI Formats for more information on the supported cloud object storage services.

3. You can now run queries on the external partitioned table you created in the previous step.

   Your Autonomous Database takes advantage of the partitioning information of your external partitioned table, ensuring that the query only accesses the relevant data files in the Object Store.

   For example:

   ```
   SELECT movie_id, month FROM sales WHERE month='2019-02'
   ```

   The external partitioned tables you create with `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` include two invisible columns `file$path` and `file$name`. These columns help identify which file a record is coming from. See External Table Metadata Columns for more information.

If there are any rows in the source files that do not match the format options you specified, the query reports an error. You can use `DBMS_CLOUD` parameters, like *rejectlimit,* to suppress these errors. As an alternative, you can also validate the external partitioned table you created to see the error messages and the rejected rows so that you can change your format options accordingly. See Validate External Data and Validate External Partitioned Data for more information.

Query External Partitioned Data with Folder Format Source File Organization

Use `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` to create an external partitioned table and generate the partitioning information from the Cloud Object Store file path.

When you create an external table with folder format data files, you have two options for specifying the types of the partition columns:

- You can manually specify the columns and their data types with `column_list` parameter. See Query External Partitioned Data with Hive Format Source File Organization for an example using the `column_list` parameter.
- You can let `DBMS_CLOUD` derive the data file columns and their types from information in structured data files such as Avro, ORC, and Parquet data files. In this case, you use the `partition_columns` option with the `format` parameter to supply the column names and their data types for the partition columns and you do not need to supply the `column_list` or the `field_list` parameters.

Consider the following sample source files in Object Store:

```
.../sales/USA/2020/01/sales1.parquet

.../sales/USA/2020/02/sales2.parquet
```

To create a partitioned external table with the Cloud Object Store file path defining the partitions from files with this sample folder format, do the following:

1. Store your object store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL (
       credential_name => 'DEF_CRED_NAME',
       username => 'adb_user@example.com',
       password => 'password' );
   END;
   /
   ```

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials you can then use the same credential name for creating external tables.

   See CREATE_CREDENTIAL Procedure for information about the `username` and `password` parameters for different object storage services.

2. Create an external partitioned table on top of your source files using the procedure `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE`.

The procedure DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE supports external partitioned files in the supported cloud object storage services. The credential is a table level property; therefore, the external files must all be on the same cloud object store.

For example:

```
BEGIN DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE(
    table_name => 'MYSALES',
    credential_name => 'DEF_CRED_NAME',
    file_uri_list => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-
string/b/bucketname/o/sales/*.parquet',
    format =>
        json_object('type' value 'parquet', 'schema' value 'first',
                    'partition_columns' value
                          json_array(
                                json_object('name' value 'country', 'type' value
'varchar2(100)'),
                                json_object('name' value 'year', 'type' value 'number'),
                                json_object('name' value 'month', 'type' value
'varchar2(2)')
                          )
         )
    );
END;
/
```

The DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE parameters for structured data files, such as for a Parquet data file does not require the column_list or the field_list parameters. The column names and data types are derived for the columns from the first parquet file that the procedure scans (and therefore all files must have the same shape). The generated column list includes the columns derived from the object name and these column have the data types specified with the partition_columns format parameter.

The parameters are:

- table_name: is the external table name.
- credential_name: is the name of the credential created in the previous step.
- file_uri_list: is a comma-delimited list of source file URIs. There are two options for this list:
  - Specify a comma-delimited list of individual file URIs without wildcarding.
  - Specify a single file URI with wildcards, where the wildcards can only be after the last slash "/". The character "*" can be used as the wildcard for multiple characters, the character "?" can be used as the wildcard for a single character.
- column_list: is a comma delimited list of column names and data types for the external table. The list includes the columns that are inside the file as well as those derived from the object name.

  The column_list is not required when the data files are structured files (Parquet, Avro, or ORC).
- field_list: Identifies the fields in the source files and their data types. The default value is NULL meaning the fields and their data types are determined by the column_list parameter.

The `field_list` is not required when the data files are structured files (Parquet, Avro, or ORC).

- `format`: defines the options you can specify to describe the format of the source file. The `partition_columns format` parameter specifies the names of the partition columns. See DBMS_CLOUD Package Format Options for more information.

  If the data in your source file is encrypted, decrypt the data by specifying the `encryption` format option. See Decrypt Data While Importing from Object Storage for more information on decrypting data.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

See CREATE_EXTERNAL_PART_TABLE Procedure for detailed information about the parameters.

See DBMS_CLOUD URI Formats for more information on the supported cloud object storage services.

If there are any rows in the source files that do not match the format options you specified, the query reports an error. You can use `DBMS_CLOUD` parameters, like `rejectlimit` to suppress these errors. As an alternative, you can also validate the external partitioned table you created to see the error messages and the rejected rows so that you can change your format options accordingly. See Validate External Data and Validate External Partitioned Data for more information.

3. You can now run queries on the external partitioned table you created in the previous step.

   Your Autonomous Database takes advantage of the partitioning information of your external partitioned table, ensuring that the query only accesses the relevant data files in the Object Store. For example, the following query only reads data files from one partition.

   For example:

   ```
   SELECT year, month, product, units
   FROM SALES WHERE year='2020' AND month='02' AND country='USA'
   ```

   The external partitioned tables you create with `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` include two invisible columns `file$path` and `file$name`. These columns help identify which file a record is coming from. See External Table Metadata Columns for more information.

Refresh External Partitioned Tables with Updated or Deleted Source Files

You can use `DBMS_CLOUD.SYNC_EXTERNAL_PART_TABLE` to refresh an external partitioned table. Use this procedure when new partitions are added or when partitions are removed from the object store source.

> **Note:**
>
> Only use `DBMS_CLOUD.SYNC_EXTERNAL_PART_TABLE` when the partitioned external table is created with `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` and the `file_url_path` parameter.

To refresh a partitioned external table:

1. Refresh an external partitioned table on top of your Cloud Object Store source files using the procedure `DBMS_CLOUD.SYNC_EXTERNAL_PART_TABLE`.

   For example:

```
BEGIN
   DBMS_CLOUD.SYNC_EXTERNAL_PART_TABLE(table_name => 'MYSALES');
END;
/
```

   `DBMS_CLOUD.SYNC_EXTERNAL_PART_TABLE` uses the credential information from the corresponding `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` with the specified `table_name` to access Cloud Object Store.

   See SYNC_EXTERNAL_PART_TABLE Procedure for detailed information about the parameter.

2. The partition information is now refreshed and you can run queries on the updated external partitioned table. with new partitions available or with partitions that were removed no longer available in the external partitioned table.

Query External Partitioned Data (with Partitioning Clause)

If you want to query multiple data files in the Object Store as a single external table and the files can be represented as multiple logical partitions, then it is highly recommended to use an external partitioned table. Using an external partitioned table preserves the logical partitioning of your data files for query access. Use the procedure `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` to create an external partitioned table.

There are two ways to create an external partitioned table on Autonomous Database:

- The first version is for. See Query External Partitioned Data with Hive Format Source File Organization for information on this type of external partitioned table usage.
- The second version is for. When you a create a partitioned external table in this way you include a partitioning clause in the `partitioning_clause` parameter. The partitioning clause that you include depends upon the data files in the Cloud and the type of partitioning you use. This section describes this type of external partitioned table usage.

1. Store your object store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL (
    credential_name => 'DEF_CRED_NAME',
    username => 'adb_user@example.com',
    password => 'password' );
END;
/
```

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your

object store credentials change. Once you store the credentials you can then use the same credential name for creating external tables.

See CREATE_CREDENTIAL Procedure for information about the `username` and `password` parameters for different object storage services.

2. Create an external partitioned table on top of your source files using the procedure `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE`.

   The procedure `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` supports external partitioned files in the supported cloud object storage services. The credential is a table level property; therefore, the external files must be on the same object store.

   For example:

```
BEGIN
  DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE(
      table_name =>'PET1',
      credential_name =>'DEF_CRED_NAME',
      format => json_object('delimiter' value ',', 'recorddelimiter' value 'newline',
'characterset' value 'us7ascii'),
      column_list => 'col1 number, col2 number, col3 number',
      partitioning_clause => 'partition by range (col1)
         (partition p1 values less than (1000) location
             ( ''https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/
bucketname/o/file_11.txt''),
          partition p2 values less than (2000) location
             ( ''https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/
bucketname/o/file_21.txt''),
          partition p3 values less than (3000) location
             ( ''https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/
bucketname/o/file_31.txt'') )'
     );
END;
/
```

   The parameters are:

   - `table_name`: is the external table name.
   - `credential_name`: is the name of the credential created in the previous step.
   - `partitioning_clause`: is the complete partitioning clause, including the location information for individual partitions.
   - `format`: defines the options you can specify to describe the format of the source file. format: defines the options you can specify to describe the format of the source file.

     If the data in your source files is encrypted, decrypt the data by specifying the `format` parameter with the `encryption` option. See Decrypt Data While Importing from Object Storage for more information on decrypting data.
   - `column_list`: is a comma delimited list of the column definitions in the source files.

   In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

   You can now run queries on the external partitioned table you created in the previous step. Your Autonomous Database takes advantage of the partitioning information of your

external partitioned table, ensuring that the query only accesses the relevant data files in the Object Store. For example, the following query only reads data files from partition P1:

```
SELECT * FROM pet1 WHERE col1 < 750;
```

The external partitioned tables you create with `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` include two invisible columns `file$path` and `file$name`. These columns help identify which file a record is coming from. See External Table Metadata Columns for more information.

If there are any rows in the source files that do not match the format options you specified, the query reports an error. You can use `DBMS_CLOUD` parameters, like *rejectlimit,* to suppress these errors. As an alternative, you can also validate the external partitioned table you created to see the error messages and the rejected rows so that you can change your format options accordingly. See Validate External Data and Validate External Partitioned Data for more information.

See CREATE_EXTERNAL_PART_TABLE Procedure for detailed information about the parameters.

See DBMS_CLOUD URI Formats for more information on the supported cloud object storage services.

Query Data that Uses Table Hyperlink Source

You can create an external table over an Autonomous Database Table Hyperlink.

Consider creating external tables over an Autonomous Database Table Hyperlink in the following scenarios:

- When you need to use a REST client to create an external table and then utilize the SQL API to perform analytics on the datasets. This may include running SQL queries to analyze the data or aggregating results based on your requirements.
- When you need to query data from databases across regions.

There are two ways to create an external table in an Autonomous Database using a Table Hyperlink:

- Use a single Autonomous Database Table Hyperlink to create an external table. After the table is created, you can read data from the table using the `SELECT` statement. You can also apply filters and clauses when retrieving data from the table. For example, you can filter the data with the `WHERE` clause or sort it using the `ORDER BY` clause.

  Example to create an external table over a single Table Hyperlink:

```
DBMS_CLOUD.CREATE_EXTERNAL_TABLE (
   table_name => 'ext_tab_parurl',
   file_uri_list => 'https://dataaccess.adb.us-
chicago-1.oraclecloudapps.com/adb/p/K6XExample/data'
);
END;
/
```

After the table is created, you can query the table using the `SELECT` statement, and you can also filter or sort the output. Example to read data from the `ext_tab_parurl` table in descending order:

```
SELECT * FROM ext_tab_parurl
   ORDER BY height DESC;

COUNTY    SPECIES      HEIGHT
------    ----------   -------
First     Chestnut      51
Main      Alder         45
Main      Douglas-fir   34
Main      Pine          29
Main      Cedar         21
```

- Use a comma-delimited list of Autonomous Database Table Hyperlink URLs, you must ensure that all included Table Hyperlinks must have the same column names, column order, and column data types in the same schema.

  Example to create an external table over a list of Table Hyperlinks:

  ```
  DBMS_CLOUD.CREATE_EXTERNAL_TABLE (
     table_name => 'ext_tab_parurl_1',
     file_uri_list => 'https://dataaccess.adb.us-
  chicago-1.oraclecloudapps.com/adb/p/K6XExample/data',
                      'https://dataaccess.adb.us-
  chicago-1.oraclecloudapps.com/adb/p/K7XExample/data'
  );
  END;
  /
  ```

  Query the table using the `SELECT` statement:

  ```
  SELECT * FROM ext_tab_parurl_1;

  COUNTY   SPECIES      HEIGHT
  ------   ----------   -------
  First     Chestnut      51
  Main      Alder         45
  Main      Douglas-fir   34
  Main      Pine          29
  Main      Cedar         21
  Main      Aspen         35
  First     Maple         16
  ```

The parameters are:

- `table_name`: is the external table name.
- `file_uri_list`: is either a single Table Hyperlinks or a comma-delimited list of Table Hyperlinks.

See CREATE_EXTERNAL_TABLE Procedure for more information.

Query Hybrid Partitioned Data

If you want to query internal data and multiple data files in the Object Store as single logical table you can use a hybrid partitioned table to represent the data as single object. Use the procedure `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE` to create a hybrid partitioned table.

If your data, internal or external, can be represented in finer granularity as multiple logical partitions then it is highly recommended to create a hybrid partitioned table with multiple internal and external partitions, preserving the logical partitioning of your data for query access.

When you a create a hybrid partitioned table, you include a partitioning clause in the `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE` statement. The partitioning clause that you include depends upon your data files and the type of partitioning you use. See Creating Hybrid Partitioned Tables for more information.

1. Store your object store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL (
       credential_name => 'DEF_CRED_NAME',
       username => 'adb_user@example.com',
       password => 'password' );
   END;
   /
   ```

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials you can then use the same credential name for creating hybrid partitioned tables.

   See CREATE_CREDENTIAL Procedure for information about the `username` and `password` parameters for different object storage services.

2. Create a hybrid partitioned table on top of your source files using the procedure `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE`.

   The procedure `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE` supports external partitioned files in the supported cloud object storage services. The credential is a table level property; therefore, the external files must be on the same object store.

   For example:

```
BEGIN
  DBMS_CLOUD.CREATE_HYBRID_PART_TABLE(
      table_name =>'HPT1',
      credential_name =>'DEF_CRED_NAME',
      format => json_object('delimiter' value ',', 'recorddelimiter' value 'newline',
'characterset' value 'us7ascii'),
      column_list => 'col1 number, col2 number, col3 number',
      partitioning_clause => 'partition by range (col1)
           (partition p1 values less than (1000) external location
              ( ''https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/
```

```
bucketname/o/file_11.txt'') ,
            partition p2 values less than (2000) external location
              ( ''https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/
bucketname/o/file_21.txt'') ,
            partition p3 values less than (3000) )'
     );
END;
/
```

The parameters are:

- `table_name`: is the hybrid partitioned table name.
- `credential_name`: is the name of the credential created in the previous step.
- `partitioning_clause`: is the complete partitioning clause, including the location information for individual partitions.
- `format`: defines the options you can specify to describe the format of the source file.

  If the data in your source files is encrypted, decrypt the data by specifying the `format` parameter with the `encryption` option. See Decrypt Data While Importing from Object Storage for more information on decrypting data.
- `column_list`: is a comma delimited list of the column definitions in the source files.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

You can now run queries on the hybrid partitioned table you created in the previous step. Your Autonomous Database takes advantage of the partitioning information of your hybrid partitioned table, ensuring that the query only accesses relevant data files in the Object Store. For example, the following query only reads data files from partition P1:

```
SELECT * FROM hpt1 WHERE col1 < 750;
```

The hybrid partitioned tables you create with `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE` include two invisible columns `file$path` and `file$name`. These columns help identify which file a record is coming from. See External Table Metadata Columns for more information.

If there are any rows in the source files that do not match the format options you specified, the query reports an error. You can use `DBMS_CLOUD` parameters, like *rejectlimit,* to suppress these errors. As an alternative, you can also validate the hybrid partitioned table you created to see the error messages and the rejected rows so that you can change your format options accordingly. See Validate External Data and Validate Hybrid Partitioned Data for more information.

See CREATE_HYBRID_PART_TABLE Procedure detailed information about the parameters.

See DBMS_CLOUD URI Formats for more information on the supported cloud object storage services.

Query External Data Pump Dump Files

You can also query Oracle Data Pump dump files in the Cloud by creating an external table using `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`.

The source files to create this type of external table must be exported from the source system using the `ORACLE_DATAPUMP` access driver in External Tables. See Unloading and Loading Data with the ORACLE_DATAPUMP Access Driver for details on exporting using the `ORACLE_DATAPUMP` access driver.

To create an external table you first move the Oracle Data Pump dump files that were exported using the `ORACLE_DATAPUMP` access driver to your Object Store and then use `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` to create the external table.

The source files in this example are the Oracle Data Pump dump files, `exp01.dmp` and `exp02.dmp`.

1. Store your object store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'DEF_CRED_NAME',
       username => 'adb_user@example.com',
       password => 'password' );
   END;
   /
   ```

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials you can then use the same credential name for creating external tables.

   See CREATE_CREDENTIAL Procedure for information about the `username` and `password` parameters for different object storage services.

2. Create an external table on top of your source files using the procedure `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`.

   For example:

   ```
   BEGIN
      DBMS_CLOUD.CREATE_EXTERNAL_TABLE(
       table_name =>'CHANNELS_EXT',
       credential_name =>'DEF_CRED_NAME',
       file_uri_list =>'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
   namespace-string/b/bucketname/o/exp01.dmp,
                        https://objectstorage.us-phoenix-1.oraclecloud.com/n/
   namespace-string/b/bucketname/o/exp02.dmp'
       format => json_object('type' value 'datapump', 'rejectlimit' value
   '1'),
       column_list => 'CHANNEL_ID NUMBER, CHANNEL_DESC VARCHAR2(20),
   CHANNEL_CLASS VARCHAR2(20)' );
   END;
   /
   ```

The parameters are:

- `table_name`: is the external table name.
- `credential_name`: is the name of the credential created in the previous step.
- `file_uri_list`: is a comma delimited list of the Data Pump dump files you want to query.
- `format`: defines the options you can specify to describe the format of the source file defines the options you can specify to describe the format of the source file. When you specify the type `'datapump'`, the only other valid format parameter is `'rejectlimit'`.
- `column_list`: is a comma delimited list of the column definitions in the source files.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

You can now run queries on the external table you created in the previous step. For example:

```
SELECT count(*) FROM channels_ext;
```

By default the database expects all rows in the external data file to be valid and match both the target data type definitions as well as the format definition of the files. As part of validation, `DBMS_CLOUD` makes sure all the necessary dump file parts are there and also checks that the dump files are valid and not corrupt (for example `exp01.dmp`, `exp02.dmp`, and so on). You can use the `DBMS_CLOUD` format option `rejectlimit` to suppress these errors. As an alternative, you can also validate the external table you created to see the error messages and the rejected rows. See Validate External Data for more information.

For detailed information about the parameters, see CREATE_EXTERNAL_TABLE Procedure.

See DBMS_CLOUD URI Formats for more information on the supported cloud object storage services.

Query Big Data Service Hadoop (HDFS) Data from Autonomous Database

You can create database links to Oracle Big Data Service from Autonomous Database.

Big Data Service provides enterprise-grade Hadoop as a service, with end-to-end security, high performance, and ease of management and upgradeability. After deploying the Oracle Cloud SQL Query Server to Big Data Service, you can easily query data available on Hadoop clusters at scale from Autonomous Database using SQL. This allows you to blend data coming from your data lake with data in your Autonomous Database.

Oracle Cloud SQL Query Server is an Oracle SQL engine that can be accessed using database links in Autonomous Database. You create links to Big Data Service using `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`. See Integrating Cloud SQL with Autonomous Database for more information.

Query External Data with Data Catalog

Oracle Cloud Infrastructure Data Catalog is the metadata management service for Oracle Cloud that helps you discover data and support data governance. It provides an inventory of assets, a business glossary, and a common metastore for data lakes.

Autonomous Database can leverage this metadata to dramatically simplify management for access to your data lake's object store. Instead of manually defining external tables to access your data lake, use the external tables that are defined and managed automatically. These tables will be found in Autonomous Database protected schemas that are kept up to date with changes in Data Catalog.

For more information about Data Catalog, please refer to Data Catalog documentation.

- About Querying with Data Catalog
  By synchronizing with Data Catalog metadata, Autonomous Database automatically creates external tables for each logical entity harvested by Data Catalog. These external tables are defined in database schemas that are fully managed by the metadata synchronization process. Users can immediately query data without having to manually derive the schema (columns and data types) for external data sources and manually create external tables.
- Concepts Related to Querying with Data Catalog
  An understanding of the following concepts is necessary for querying with Data Catalog.
- Synchronization Mapping
  The synchronization process creates and updates Autonomous Database schemas and external tables based on Data Catalog data assets, folders, logical entities, attributes and relevant custom overrides.
- Typical Workflow with Data Catalog
  There is a typical workflow of actions performed by users who want to query with Data Catalog.
- Example: MovieStream Scenario
  In this scenario, Moviestream is capturing data in a landing zone on object storage. Much of this data, but not necessarily all, is then used to feed an Autonomous Database. Prior to feeding Autonomous Database, the data is transformed, cleansed and subsequently stored in the "gold" area.
- Example: Partitioned Data Scenario
  This scenario illustrates how to create external tables in Autonomous Database that are based on Data Catalog logical entities harvested from partitioned data in Object Store.

About Querying with Data Catalog

By synchronizing with Data Catalog metadata, Autonomous Database automatically creates external tables for each logical entity harvested by Data Catalog. These external tables are defined in database schemas that are fully managed by the metadata synchronization process. Users can immediately query data without having to manually derive the schema (columns and data types) for external data sources and manually create external tables.

Synchronization is dynamic, keeping the Autonomous Database up-do-date with respect to changes to the underlying data, reducing administration cost as it automatically maintains hundreds to thousands of tables. It also allows multiple Autonomous Database instances to share the same Data Catalog, further reducing management costs and providing a common set of business definitions.

The Data Catalog folders/buckets are containers that sync with Autonomous Database schemas. Logical entities within those folders/buckets map to Autonomous Database external tables. These schemas and external tables are automatically generated and maintained through the sync process:

- Folders/Buckets map to database schemas that are for organizational purposes only.
- The organization is meant to be consistent with the data lake and minimize confusion when accessing data thru different paths.

- Data Catalog is the source of truth for the tables contained within schemas. Changes made in the Data Catalog update the schema's tables during a subsequent sync.

To use this capability, a Database Data Catalog Administrator initiates a connection to a Data Catalog instance, selects which data assets and logical entities to synchronize, and runs the sync. The sync process creates schemas and external tables based on the selected Data Catalog harvested data assets and logical entities. As soon as the external tables are created, Data Analysts can start querying their data without having to manually derive the schema for external data sources and create external tables.

> **Note:**
>
> The `DBMS_DCAT` Package is available for performing the tasks required to query Data Catalog object store data assets. See DBMS_DCAT Package.

Concepts Related to Querying with Data Catalog

An understanding of the following concepts is necessary for querying with Data Catalog.

Data Catalog
Data Catalog harvests data assets that point to the object store data sources you want to query with Autonomous Database. From Data Catalog you can specify how the data is organized during harvesting, supporting different file organization patterns. As part of the Data Catalog harvesting process, you can select the buckets and files you want to manage within the asset. For further information, see Data Catalog Overview.

Object Stores
Object Stores have buckets containing a variety of objects. Some common types of objects found in these buckets include: CSV, parquet, avro, json, and ORC files. Buckets generally have a structure or a design pattern to the objects they contain. There are many different ways to structure data and many different ways of interpreting these patterns.
For example, a typical design pattern uses top-level folders that represent tables. Files within a given folder share the same schema and contain data for that table. Subfolders are often used to represent table partitions (for example, a subfolder for each day). Data Catalog refers to each top-level folder as a logical entity, and this logical entity maps to an Autonomous Database external table.

Connection
A connection is an Autonomous Database connection to a Data Catalog instance. For each Autonomous Database instance there can be connections to multiple Data Catalog instances. The Autonomous Database credential must have rights to access Data Catalog assets that have been harvested from object storage.

Harvest
A Data Catalog process that scans object storage and generates the logical entities from your data sets.

Data Asset
A data asset in Data Catalog represents a data source, which includes databases, Oracle Object Storage, Kafka, and more. Autonomous Database leverages Oracle Object Storage assets for metadata synchronization.

Data Entity
A data entity in Data Catalog is a collection of data such as a database table or view, or a single file and normally has many attributes that describe its data.

Logical Entity
In Data Lakes, numerous files typically comprise a single logical entity. For example, you may have daily clickstream files, and these files share the same schema and file type.
A Data Catalog logical entity is a group of Object Storage files that are derived during harvesting by applying filename patterns that have been created and assigned to a data asset.

Data Object
A data object in Data Catalog refers to data assets and data entities.

Filename Pattern
In a data lake, data may be organized in different ways. Typically, folders capture files of the same schema and type. You must register to Data Catalog how your data is organized. Filename patterns are used to identify how your data is organized. In Data Catalog, you can define filename patterns using regular expressions. When Data Catalog harvests a data asset with an assigned filename pattern, logical entities are created based on the filename pattern. By defining and assigning these patterns to data assets, multiple files can be grouped as logical entities based on the filename pattern.

Synchronize (Sync)
Autonomous Database performs synchronizations with Data Catalog to automatically keep its database up-to-date with respect to changes to the underlying data. Synchronization can be performed manually, or on a schedule.
The sync process creates schemas and external tables based on the Data Catalog data assets and logical entities. These schemas are protected, which means their metadata is managed by Data Catalog. If you want to alter the metadata, you must make the changes in Data Catalog. The Autonomous Database schemas will reflect any changes after the next sync is run. For further details, see Synchronization Mapping.

Synchronization Mapping

The synchronization process creates and updates Autonomous Database schemas and external tables based on Data Catalog data assets, folders, logical entities, attributes and relevant custom overrides.

Data Catalog	Autonomous Database	Mapping Description
Data asset and folder (object storage bucket)	Schema name	**Default values**: By default, the generated schema name in Autonomous Database has the following format: `DCAT$<dcat-con-id>_<data-asset-name>_<folder-name>` • `dcat-con-id` is the unique Data Catalog connection identifier. For details on specifying this identifier, see the `dcat_con_id` parameter in the DBMS_DCAT RUN_SYNC Procedure. • `data-asset-name` is the name of the Data Catalog data asset's name. • `folder-name` is the Data Catalog folder name. This folder maps to an object storage bucket. **Customizations**: The default `data-asset-name` and `folder-name` can be customized by defining custom properties, business names and display names to override these default names. • `data-asset-name` can be overidden by defining the `oracle-db-schema-prefix` custom property for the data asset in Data Catalog. • `folder-name` can be overidden by defining the `oracle-db-schema` custom property for the folder in Data Catalog, a business name or display name. The following attributes are used in order of precedence for generating the `folder-name`: 1. `oracle-db-schema` custom property 2. Business Name 3. Display Name **Examples**: • If the connection id is `DataModels`, the asset name is `ObjectStorage`, the folder name is `HR` and there are no custom property overrides, the derived schema name is: `DCAT$DATAMODELS_OBJECTSTORAGE_HR` • If the data asset name is MYASSET, the folder name is MYFOLDER, and there are no custom property overrides, the schema name is: `DCAT$MYASSET_MYFOLDER` • If the data asset has `oracle-db-prefix = FIRSTASSET`, and the folder has `oracle-db-schema = FIRSTFOLDER`, then the schema name is: `DCAT$FIRSTASSET_FIRSTFOLDER`

Data Catalog	Autonomous Database	Mapping Description
Logical entity	External table	Logical entities are mapped to external tables. If the logical entity has a partitioned attribute, it is mapped to a partitioned external table. The external table name is derived from the corresponding logical entity's Display Name, or Business Name. If `oracle-db-schema` is set, then its value overrides all the names and custom properties of the corresponding folders and data assets. For example, if `oracle-db-schema` for an entity is set to `EntitySchema`, then the table is created in schema `DCAT$ENTITYSCHEMA`.
Logical entity's attributes	External table columns	**Column names**: The external table column names are derived from the corresponding logical entity's attribute display names, or business names. For logical entities derived from Parquet, Avro, and ORC files, the column name is always the display name of the attribute as it represents the field name derived from the source files. For attributes corresponding to a logical entity derived from CSV files, the following attribute fields are used in order of precedence for generating the column name: 1. `oracle-db-column-name` 2. Business Name 3. Display Name **Column type**: The `oracle-db-column-type` custom property overrides the default column type that was derived by Data Catalog. For attributes corresponding to a logical entity derived from Avro files with `TIME_MICROS`, `TIME_MILLIS`, `TIMESTAMP_MICROS` or `TIMESTAMP_MILLIS` data types, you must set the `oracle-db-column-type` of the corresponding attribute in Data Catalog. **Column length**: The `oracle-db-column-length` custom property overrides the default column length for a string field that was derived by Data Catalog. **Column precision**: The `oracle-db-column-precision` custom property overrides the default precision for a number that was derived by Data Catalog. For attributes corresponding to a logical entity derived from Avro files with `TIME_MICROS`, `TIME_MILLIS`, `TIMESTAMP_MICROS` or `TIMESTAMP_MILLIS` data types, you must set the `oracle-db-column-precision` of the corresponding attribute in Data Catalog. **Column scale**: The `oracle-db-column-scale` custom property overrides the default scale for a number that was derived by Data Catalog.

Typical Workflow with Data Catalog

There is a typical workflow of actions performed by users who want to query with Data Catalog.

The Database Data Catalog Admin creates a connection between the Autonomous Database instance and a Data Catalog instance, then configures and runs a synchronization (sync) between the Data Catalog and Autonomous Database. The sync creates external tables and schemas in the Autonomous Database instance based on the synced Data Catalog contents. The Database Data Catalog Query Admin or Database Admin grants READ access to the generated external tables so that Data Analysts and other database users can browse and query the external tables.

The table below describes each action in detail. For a description of the different user types included in this table, see Data Catalog Users and Roles.

> **Note:**
>
> The `DBMS_DCAT` Package is available for performing the tasks required to query Data Catalog object store data assets. See DBMS_DCAT Package.

Action	Who is the user	Description
Create policies	Database Data Catalog Administrator	The Autonomous Database resource principal or Autonomous Database user credential must have the appropriate permissions to manage Data Catalog and to read from object storage. **More information**: Required Credentials and IAM Policies.
Create credentials	Database Data Catalog Administrator	Ensure database credentials are in place to access Data Catalog and to query object store. The user calls `DBMS_CLOUD.CREATE_CREDENTIAL` to create user credentials and/or `DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL` to enable resource principals. **More information**: DBMS_CLOUD CREATE_CREDENTIAL Procedure, Use Resource Principal with DBMS_CLOUD.

Action	Who is the user	Description
Create connections to Data Catalog	Database Data Catalog Administrator	To initiate a connection between an Autonomous Database instance and a Data Catalog instance the user calls `DBMS_DCAT.SET_DATA_CATALOG_CONN` to specify the target Data Catalog instance. Connections from an Autonomous Database instance to multiple Data Catalog instances is supported. The connection to the Data Catalog instance must use a database credential object with sufficient Oracle Cloud Infrastructure (OCI) privileges. For example, the Resource Principal Service Token for the Autonomous Database instance or an OCI user with sufficient privileges can be used. Once the connection has been made, the Data Catalog instance is updated with the `DBMS_DCAT` namespace and custom properties (if they do not already exist). The user can run a query to see the new connection including all current connections: `select * from all_dcat_connections;` **More information**: SET_DATA_CATALOG_CONN Procedure, UNSET_DATA_CATALOG_CONN Procedure.
Create a selective sync	Database Data Catalog Administrator	Create a sync job by selecting the Data Catalog objects to sync. The user can: • Select data assets/folders to sync. • Select individual logical entities to sync. • Preview the resulting external tables before syncing. • Change external tables (for example, the name) by modifying custom properties on Data Catalog. **More information**: See CREATE_SYNC_JOB Procedure, DROP_SYNC_JOB Procedure, Synchronization Mapping
Sync with Data Catalog	Database Data Catalog Administrator	The user initiates a sync operation. The sync is initiated manually through the `DBMS_DCAT.RUN_SYNC` procedure call, or automatically as part of a scheduled sync job. The sync operation creates, modifies and drops external tables and schemas according to the Data Catalog contents and sync selections. Manual configuration is applied using Data Catalog Custom Properties. **More information**: See RUN_SYNC Procedure, CREATE_SYNC_JOB Procedure, Synchronization Mapping
Monitor sync and view logs	Database Data Catalog Administrator	The user can view the sync status by querying the `USER_LOAD_OPERATIONS` view. After the sync process has completed, the user can view a log of the sync results, including details about the mappings of logical entities to external tables. **More information**: Monitoring and Troubleshooting Loads
Grant privileges	Database Data Catalog Query Administrator, Database Administrator	The database Data Catalog Query Administrator or database Administrator must grant READ on generated external tables to data analyst users. This allows the data analysts to query the generated external tables.

Action	Who is the user	Description
Browse and query external tables	Data Analyst	Data analysts are able to query the external tables through any tool or application that supports Oracle SQL. Data Analysts can review the synced schemas and tables in the DCAT$* schemas, and query the tables using Oracle SQL. **More information**: Synchronization Mapping
Terminate connections to Data Catalog	Database Data Catalog Administrator	To remove an existing Data Catalog association, the user calls the UNSET_DATA_CATALOG_CONN procedure. This action is only done when you no longer plan on using Data Catalog and the external tables that are derived from the catalog. This action deletes Data Catalog metadata, and drops synced external tables from the Autonomous Database instance. The custom properties on Data Catalog and OCI policies are not affected. **More information**: UNSET_DATA_CATALOG_CONN Procedure

Example: MovieStream Scenario

In this scenario, Moviestream is capturing data in a landing zone on object storage. Much of this data, but not necessarily all, is then used to feed an Autonomous Database. Prior to feeding Autonomous Database, the data is transformed, cleansed and subsequently stored in the "gold" area.

Data Catalog is used to harvest these sources and then provide a business context to the data. Data Catalog metadata is shared with Autonomous Database, allowing Autonomous Database users to query those data sources using Oracle SQL. This data may be loaded into Autonomous Database or queried dynamically using external tables.

For more information on using Data Catalog, see Data Catalog Documentation.

1. Object Store - Review buckets, folders and files
   a. Review the buckets in your object store.

      For example, below are the landing (`moviestream_landing`) and gold zone (`moviestream_gold`) buckets in object storage:

   b. Review the folders and files in the object store buckets.

      For example, below are the folders in the landing bucket (`moviestream_landing`) in object storage:
2. Data Catalog - Create filename patterns
   a. Inform Data Catalog how your data is organized using filename patterns. These are regular expressions used to categorize files. The filename patterns are used by the Data Catalog harvester to derive logical entities. The following two filename patterns are used to harvest the buckets in the MovieStream example. See Harvesting Object Storage Files as Logical Data Entities for further details on creating filename patterns.

Hive-style	Folder-style
`{bucketName:.*}/{logicalEntity:[^/]+}.db/{logicalEntity:[^/]+}/.*`	`{bucketName:[\w]+}/{logicalEntity:[^/]+}(?<!.db)/.*$`

Hive-style	Folder-style
• Creates logical entities for sources that contain ".db" as the first part of the object name. • To ensure uniqueness within the bucket, the resulting name is (db-name).(folder name)	• Creates a logical entity based on the folder name off of the root • To prevent duplication with Hive, object names that have ".db" in them are skipped.

b. To create filename patterns, go to the **Filename Patterns** tab for your Data Catalog and click **Create Filename Pattern**. For example, the following is the **Create Filename Pattern** tab for the `moviestream` Data Catalog:

3. Data Catalog - Data Asset Creation

   a. Create a data asset that is used to harvest data from your object store.

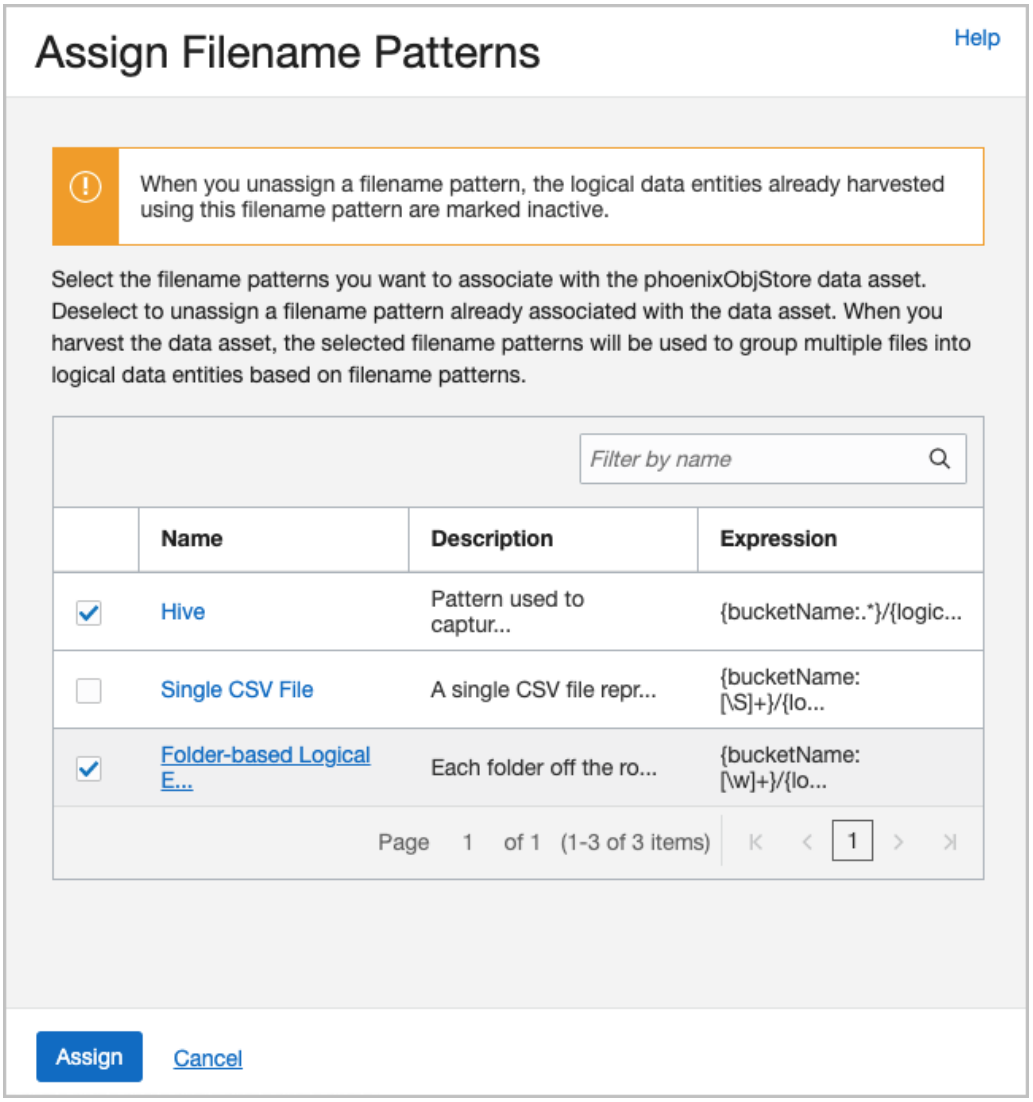
      For example, a data asset named `phoenixObjStore` is created in the `moviestream` Data Catalog:

   b. Add a connection to your data asset.

      In this example, the data asset connects to the compartment for the `moviestream` object storage resource.

   c. Now, associate your filename patterns with your data asset. Select **Assign Filename Patterns**, check the patterns you want and click **Assign**.

For example, here are the patterns assigned to the `phoenixObjStore` data asset:



4. Data Catalog - Harvest data from object store

   a. Harvest the Data Catalog data asset. Select the object store buckets containing the source data.

      In this example, the `moviestream_gold` and `moviestream_landing` buckets from object store are selected for harvesting.

   b. After running the job, you see the logical entities. Use the **Browse Data Assets** to review them.

      In this example, you are looking at the `customer-extension` logical entity and its attributes.
      If you have a glossary, Data Catalog recommends categories and terms to associate with the entity and its attributes. This provides a business context for the items. Schemas, tables and columns are oftentimes not self-explanatory.

      In our example, we want to differentiate between the different types of buckets and the meaning of their content:

      - what is a landing zone?
      - how accurate is the data?

- when was it last updated?
- what is the definition of a logical entity or its attribute

5. Autonomous Database - Connect to Data Catalog

   Connect Autonomous Database to Data Catalog. You need to ensure that the credential used to make that connection is using an OCI principal that is authorized to access the Data Catalog asset. For further information, see Data Catalog Policies and Access Cloud Resources by Configuring Policies and Roles.

   a. Connect to Data Catalog

```
-- Variables are used to simplify usage later
define oci_credential = 'OCI$RESOURCE_PRINCIPAL'
define dcat_ocid = 'ocid1.datacatalog.oc1.iad.aaaaaaaardp66bg....twiq'
define dcat_region='us-ashburn-1'
define uri_root =    'https://objectstorage.us-
ashburn-1.oraclecloud.com/n/mytenancy/b/landing/o'
define uri_private = 'https://objectstorage.us-
ashburn-1.oraclecloud.com/n/mytenancy/b/private_data/o'

-- Run as admin
-------
-- Enable resource principal support
-------
exec dbms_cloud_admin.enable_resource_principal();

-- Test to make sure credential was created. Returns a row if it was
successful
select *
from dba_credentials
where credential_name = 'OCI$RESOURCE_PRINCIPAL' and owner = 'ADMIN';

-- Query a private bucket to test the principal and privileges.
select *
from dbms_cloud.list_objects('&oci_credential', '&uri_private/');

--------
-- Set the credentials to use for object store and data catalog
-- Connect to Data Catalog
-- Review connection
---------
-- Set credentials
exec dbms_dcat.set_data_catalog_credential(credential_name =>
'&oci_credential');
exec dbms_dcat.set_object_store_credential(credential_name =>
'&oci_credential');

-- Connect to Data Catalog
begin
   dbms_dcat.set_data_catalog_conn (
     region => '&dcat_region',
     catalog_id => '&dcat_ocid');
end;
/
```

```
-- Review the connection
select * from all_dcat_connections;
```

b. Sync Data Catalog with Autonomous Database. Here, we'll sync all the object storage assets:

```
-- Sync Data Catalog with Autonomous Database
---- Let's sync all of the assets.
begin
 dbms_dcat.run_sync('{"asset_list":["*"]}');
end;
/

-- View log
select type, start_time, status, logfile_table from
user_load_operations;  -- Logfile_Table will have the name of the table
containing the full log.
select * from dbms_dcat$1_log;

-- View the new external tables
select * from dcat_entities;
select * from dcat_attributes;
```

c. Autonomous Database - Now start running queries against object store.

```
-- Query the Data !
select *from dcat$phoenixobjstore_moviestream_gold.genre
;
```

6. Change schemas for objects

   The default schema names are rather complicated. Let's simplify them by specifying both the asset and the folder's `Oracle-Db-Schema` custom attribute in Data Catalog. Change the data asset to `PHX` and the folders to `landing` and `gold` respectively. The schema is a concatenation of the two.

   a. From Data Catalog, navigate to the `moviestream_landing` bucket and change the asset to `landing` and `gold` respectively.

      Before change:

      After change:

   b. Run another sync.

Example: Partitioned Data Scenario

This scenario illustrates how to create external tables in Autonomous Database that are based on Data Catalog logical entities harvested from partitioned data in Object Store.

The following example is based on Example: MovieStream Scenario and has been adapted to demonstrate integrating with partitioned data. Data Catalog is used to harvest these sources and then provide a business context to the data. For further details about this example, see Example: MovieStream Scenario .

For more information on using Data Catalog, see Data Catalog Documentation.

1. Object Store - Review buckets, folders and files
   a. Review the buckets in your object store.

      For example, below are the landing (`moviestream_landing`) and gold zone (`moviestream_gold`) buckets in object storage:

   b. Review the folders and files in the object store buckets.

      For example, below are the folders in the landing bucket (`moviestream_landing`) in object storage:

2. Data Catalog - Create filename patterns
   a. Inform Data Catalog how your data is organized using filename patterns. These are folder prefixes or regular expressions used to categorize files. The filename patterns are used by the Data Catalog harvester to derive logical entities. When a folder prefix is specified, the Data Catalog automatically generates logical entities from the specified folder prefix in the object store. The following filename pattern is used to harvest the buckets in the MovieStream example. See Harvesting Object Storage Files as Logical Data Entities for further details on creating filename patterns.

Folder prefix	Description
`workshop.db/`	Creates logical entities for sources that contain "workshop.db" path in the object store.

   b. To create filename patterns, go to the **Filename Patterns** tab for your Data Catalog and click **Create Filename Pattern**. For example, the following is the **Create Filename Pattern** tab for the `moviestream` Data Catalog:

3. Data Catalog - Data Asset Creation
   a. Create a data asset that is used to harvest data from your object store.

      For example, a data asset named `amsterdamObjStore` is created in the `moviestream` Data Catalog:

   b. Add a connection to your data asset.

      In this example, the data asset connects to the compartment for the `moviestream` object storage resource.

   c. Now, associate your filename patterns with your data asset. Select **Assign Filename Patterns**, check the patterns you want and click **Assign**.

      For example, here are the patterns assigned to the `amsterdamObjStore` data asset:

4. Data Catalog - Harvest data from object store
   a. Harvest the Data Catalog data asset. Select the object store buckets containing the source data.

      In this example, the `moviestream_gold` and `moviestream_landing` buckets from object store are selected for harvesting.

   b. After running the job, you see the logical entities. Use the **Browse Data Assets** to review them.

      In this example, you are looking at the `sales_sample_parquet` logical entity and its attributes. Note that Data Catalog has identified the `month` attribute as partitioned.

5. Autonomous Database - Connect to Data Catalog

   Connect Autonomous Database to Data Catalog. You need to ensure that the credential used to make that connection is using an OCI principal that is authorized to access the

Data Catalog asset. For further information, see Data Catalog Policies and Access Cloud Resources by Configuring Policies and Roles.

a. Connect to Data Catalog

```
-- Variables are used to simplify usage later
define oci_credential = 'OCI$RESOURCE_PRINCIPAL'
define dcat_ocid = 'ocid1.datacatalog.oc1.eu-
amsterdam-1....leguurn3dmqa'
define dcat_region='eu-amsterdam-1'
define uri_root =    'https://objectstorage.us-
ashburn-1.oraclecloud.com/n/mytenancy/b/landing/o'
define uri_private = 'https://objectstorage.us-
ashburn-1.oraclecloud.com/n/mytenancy/b/private_data/o'

-- Run as admin
-------
-- Enable resource principal support
-------
exec dbms_cloud_admin.enable_resource_principal();

-- Test to make sure credential was created. Returns a row if it was
successful
select *
from dba_credentials
where credential_name = 'OCI$RESOURCE_PRINCIPAL' and owner = 'ADMIN';

-- Query a private bucket to test the principal and privileges.
select *
from dbms_cloud.list_objects('&oci_credential', '&uri_private/');

--------
-- Set the credentials to use for object store and data catalog
-- Connect to Data Catalog
-- Review connection
---------
-- Set credentials
exec dbms_dcat.set_data_catalog_credential(credential_name =>
'&oci_credential');
exec dbms_dcat.set_object_store_credential(credential_name =>
'&oci_credential');

-- Connect to Data Catalog
begin
   dbms_dcat.set_data_catalog_conn (
     region => '&dcat_region',
     catalog_id => '&dcat_ocid');
end;
/
-- Review the connection
select * from all_dcat_connections;
```

b. Sync Data Catalog with Autonomous Database. Here, we'll sync all the object storage assets:

```
-- Sync Data Catalog with Autonomous Database
---- Let's sync all of the assets.
begin
 dbms_dcat.run_sync('{"asset_list":["*"]}');
end;
/

-- View log
select type, start_time, status, logfile_table from
user_load_operations;  -- Logfile_Table will have the name of the table
containing the full log.
select * from dbms_dcat$1_log;

-- View the new external tables
select * from dcat_entities;
select * from dcat_attributes;
```

c. Autonomous Database - Now start running queries against object store.

```
-- Query the Data !
select count(*) from
DCAT$AMSTERDAMOBJSTORE_MOVIESTREAM_LANDING.SALES_SAMPLE_PARQUET;

-- Examine the generated partitioned table
select
dbms_metadata.get_ddl('TABLE','SALES_SAMPLE_PARQUET','DCAT$AMSTERDAMOBJS
TORE_MOVIESTREAM_LANDING') from dual;

CREATE TABLE
"DCAT$AMSTERDAMOBJSTORE_MOVIESTREAM_LANDING"."SALES_SAMPLE_PARQUET"
(    "MONTH" VARCHAR2(4000) COLLATE "USING_NLS_COMP",
     "DAY_ID" TIMESTAMP (6),
     "GENRE_ID" NUMBER(20,0),
     "MOVIE_ID" NUMBER(20,0),
     "CUST_ID" NUMBER(20,0),
...
)  DEFAULT COLLATION "USING_NLS_COMP"
ORGANIZATION EXTERNAL
 ( TYPE ORACLE_BIGDATA
   ACCESS PARAMETERS
   ( com.oracle.bigdata.fileformat=parquet
     com.oracle.bigdata.filename.columns=["MONTH"]
     com.oracle.bigdata.file_uri_list="https://swiftobjectstorage.eu-
amsterdam-1.oraclecloud.com/v1/tenancy/moviestream_landing/workshop.db/
sales_sample_parquet/*"
...
   )
 )
REJECT LIMIT 0
PARTITION BY LIST ("MONTH")
(PARTITION "P1"  VALUES (('2019-01'))
```

```
   LOCATION ( 'https://swiftobjectstorage.eu-
amsterdam-1.oraclecloud.com/v1/tenancy/moviestream_landing/workshop.db/
sales_sample_parquet/month=2019-01/*'),
PARTITION "P2"  VALUES (('2019-02'))
   LOCATION ( 'https://swiftobjectstorage.eu-
amsterdam-1.oraclecloud.com/v1/tenancy/moviestream_landing/workshop.db/
sales_sample_parquet/month=2019-02/*'),
...PARTITION "P24"  VALUES (('2020-12'))
   LOCATION ( 'https://swiftobjectstorage.eu-
amsterdam-1.oraclecloud.com/v1/tenancy/moviestream_landing/workshop.db/
sales_sample_parquet/month=2020-12/*'))
PARALLEL
```

6. Change schemas for objects

   The default schema names are rather complicated. Let's simplify them by specifying both the asset and the folder's `Oracle-Db-Schema` custom attribute in Data Catalog. Change the data asset to `PHX` and the folders to `landing` and `gold` respectively. The schema is a concatenation of the two.

   a. From Data Catalog, navigate to the `moviestream_landing` bucket and change the asset to `landing` and `gold` respectively.

      Before change:

      After change:

   b. Run another sync.

Query External Data with AWS Glue Data Catalog

Autonomous Database supports a system for synchronizing with an Amazon AWS Glue Data Catalog instance.

- About Querying with AWS Glue Data Catalog
  Autonomous Database allows you to synchronize with Amazon Web Service (AWS) Glue Data Catalog metadata. A database external table is automatically created by Autonomous Database for every table harvested by AWS Glue about data stored in Amazon Simple Storage Service (S3). Users can query data stored in S3 from Autonomous Database without having to manually derive the schema for the external data sources and create external tables.
- Concepts Related to Querying with AWS Glue Data Catalog
  An understanding of the following concepts is necessary for querying with Amazon Web Service (AWS) Glue data catalogs.
- Mapping Between Autonomous Database and AWS Glue
  During the synchronization process, external tables are created in Autonomous Database derived from the AWS Glue Data Catalog databases and tables over Amazon S3.
- User Workflow for Querying with AWS Glue Data Catalog
  The basic user workflow for querying AWS S3 data with AWS Glue Data Catalog involves connecting to AWS Glue Data Catalog, synchronizing with Autonomous Database to automatically create external tables, and then querying the S3 data.
- Example: Query with AWS Glue Data Catalog
  This example steps you through the process of running queries over datasets stored in Amazon Simple Storage Service (Amazon S3) using an AWS Glue Data Catalog.

About Querying with AWS Glue Data Catalog

Autonomous Database allows you to synchronize with Amazon Web Service (AWS) Glue Data Catalog metadata. A database external table is automatically created by Autonomous Database for every table harvested by AWS Glue about data stored in Amazon Simple Storage Service (S3). Users can query data stored in S3 from Autonomous Database without having to manually derive the schema for the external data sources and create external tables.

Amazon AWS Glue Data Catalog is a centralized metadata management service that helps data professionals discover data and supports data governance in AWS cloud. An Autonomous Database instance has the ability to synchronize automatically data catalog metadata with AWS Glue Data Catalog allowing database users to immediately use Autonomous Database to query data stored in the AWS cloud.

Synchronizing with AWS Glue Data Catalog has the same properties as synchronizing with OCI Data Catalog. Synchronization is dynamic, keeping the database up-do-date with respect to changes to the underlying data, reducing administration cost as it automatically maintains hundreds to thousands of tables.

Concepts Related to Querying with AWS Glue Data Catalog

An understanding of the following concepts is necessary for querying with Amazon Web Service (AWS) Glue data catalogs.

AWS Glue Data Catalog: Database
An AWS Glue database represents a collection of relational table definitions, organized in a logical group. Each AWS Glue data catalog instance manages multiple databases.

AWS Glue Data Catalog: Table
An AWS Glue table represents a relational table over data stored in the AWS cloud. An AWS Glue table defines the schema of the underlying data and consists of column information, partition information, serialization information, storage information, statistics, user-defined metadata and other metadata. Tables in AWS Glue data catalog can be created manually, or automatically using an AWS Glue crawler.

Glue Data Catalog: Crawler
You can use a crawler to populate the AWS Glue data catalog with tables. This is the primary method used by most AWS Glue users. A crawler can crawl multiple data stores in a single run. Upon completion, the crawler creates or updates one or more tables in your data catalog. Extract, transform, and load (ETL) jobs that you define in AWS Glue use these data catalog tables as sources and targets. The ETL job reads from and writes to the data stores that are specified in the source and target data catalog tables.
AWS Glue tables can be created manually by the user or automatically using a predefined or a custom crawler. Crawlers connect to the underlying data stores (for example, Amazon S3), invoke classifiers for deriving the schema of the data, and create AWS Glue tables for storing the inferred metadata. AWS Glue provides classifiers for common file types, such as CSV, JSON, Parquet, and AVRO.

Mapping Between Autonomous Database and AWS Glue

During the synchronization process, external tables are created in Autonomous Database derived from the AWS Glue Data Catalog databases and tables over Amazon S3.

AWS Glue organizes collected metadata in databases and tables. An AWS Glue database is a collection of relational table definitions. AWS Glue tables that describe the common schema and properties of the files associated with the table.

AWS Glue follows the relational model for representing attributes. For mapping hierarchical schemas to relational schemas, AWS Glue infers the schema of the semi-structured data and flattens the data to a relational schema using an ETL process.

The following table represents the mapping between OCI Data Catalog concepts and AWS Glue Data Catalog Concepts.

Table 3-5 OCI Data Catalog to AWS Glue Data Catalog Mapping

OCI Data Catalog	AWS Glue Data Catalog	Oracle Database
Data Asset	Database	Schema
Folder	(Bucket)	Schema
Logical Entity	Table	Table

User Workflow for Querying with AWS Glue Data Catalog

The basic user workflow for querying AWS S3 data with AWS Glue Data Catalog involves connecting to AWS Glue Data Catalog, synchronizing with Autonomous Database to automatically create external tables, and then querying the S3 data.

The Database Data Catalog Administrator creates a connection between the Autonomous Database instance and an AWS Glue Data Catalog instance, then configures and runs a synchronization (sync) between the AWS Glue Data Catalog and Autonomous Database. Autonomous Database automatically creates an external table for tables harvested by AWS Glue about data stored in S3.
The Database Data Catalog Query Admin or Database Admin grants READ access to the generated external tables so that Data Analysts and other database users can browse and query Autonomous Database without having to manually derive the schema for the external data sources and create external tables.

Users

The table below describes the different types of users that perform the user workflow actions.

User	Description
Database Data Catalog Administrator	Database user with `DCAT_SYNC` role.
Database Data Catalog Query Administrator	Database user able to grant access on automatically created external tables to other users.
Data Analyst	Database user on Autonomous Database querying data in AWS S3 either by querying automatically created external tables or by interacting directly with AWS Glue Data Catalog.
AWS Glue Data Catalog User	AWS user with access to an AWS Glue Data Catalog.
AWS S3 Object Storage User	AWS user with access to data stored in AWS S3

User Workflow

The table below describes each action included in the workflow and what type of user can perform the action.

> **Note:**
>
> The `DBMS_DCAT` package is available for performing the tasks required to query AWS S3 object storage using AWS Glue Data Catalog. See DBMS_DCAT Package.

Action	Who is the user	Description
Create policies	Database Data Catalog Administrator	The Autonomous Database user credential must have the appropriate permissions to access AWS Glue Data Catalog and to read from S3 object storage. **More information**: Required Credentials and IAM Policies.
Create credentials	Database Data Catalog Administrator	Ensure database credentials are in place to access AWS Glue Data Catalog and to query S3 object storage. The user calls `DBMS_CLOUD.CREATE_CREDENTIAL` to create user credentials. **Note:** Only Amazon Web Services (AWS) credentials are supported. AWS Amazon Resource Names (ARN) credentials are not supported. **More information**: DBMS_CLOUD CREATE_CREDENTIAL Procedure
Connect	Database Data Catalog Administrator	Establish a connection between an Autonomous Database instance and an AWS Glue Data Catalog instance. The connection uses the privileges of the AWS Glue Data Catalog User. Connections from an Autonomous Database instance to multiple AWS Glue Data Catalog instances is supported. To initiate a connection between an Autonomous Database instance and an AWS Glue Data Catalog instance the user: 1. Calls `DBMS_DCAT.SET_DATA_CATALOG_CREDENTIAL` using an AWS credential (CREATE_CREDENTIAL Procedure) to specify the target AWS Glue Data Catalog instance. 2. Calls the `DBMS_DCAT.SET_OBJECT_STORE_CREDENTIAL` using an AWS credential (CREATE_CREDENTIAL Procedure) to access the AWS S3 object storage. 3. Calls `DBMS_DCAT.SET_DATA_CATALOG_CONN` using an AWS Glue service endpoint Once the connection has been made, Autonomous Database stores the associated metadata, such as the AWS Glue catalog id, region, endpoint, and credential objects. **More information**: SET_DATA_CATALOG_CONN Procedure, UNSET_DATA_CATALOG_CONN Procedure, SET_DATA_CATALOG_CREDENTIAL Procedure, SET_OBJECT_STORE_CREDENTIAL Procedure.

Action	Who is the user	Description
Synchronize	Database Data Catalog Administrator	The user can manually start a synchronization with connected AWS Glue Data Catalogs using `DBMS_DCAT.RUN_SYNC` or create automatic synchronizations using `DBMS_DCAT.CREATE_SYNC_JOB`. Synchronization does the following: • Creates external tables in Autonomous Database that are derived from the connected AWS Glue Data Catalog. • Allows users to preview metadata for AWS Glue databases and tables. • Modifies or drops existing synchronized external tables if AWS Glue metadata has changed. **More information**: See RUN_SYNC Procedure, CREATE_SYNC_JOB Procedure, DROP_SYNC_JOB Procedure, Synchronization Mapping
Monitor Synchronization	Database Data Catalog Administrator	The user can view the sync status by querying the `USER_LOAD_OPERATIONS` view. After the sync process has completed, the user can view a log of the sync results, including details about the mappings to external tables. **More information**: Monitoring and Troubleshooting Loads
Grant privileges	Database Data Catalog Query Administrator, Database Administrator	The database Data Catalog Query Administrator or Database Administrator must grant READ privileges on generated external tables to data analyst users. This allows the data analysts to query the generated external tables.
Query	Data Analyst	Data analysts are able to review the synced schemas and tables in the GLUE$* schemas and query the external tables through any tool or application that supports Oracle SQL. Data in S3 is accessed using the privileges of the AWS S3 object storage user. **More information**: Example: Query with AWS Glue Data Catalog
Terminate connections	Database Data Catalog Administrator	To remove an existing Data Catalog association, the user calls the `UNSET_DATA_CATALOG_CONN` procedure. This action is only done when you no longer plan on using the connected AWS Glue Data Catalog and the external tables that are derived from the catalog. This action deletes AWS Glue Data Catalog metadata, and drops synced external tables from the Autonomous Database instance. **More information**: UNSET_DATA_CATALOG_CONN Procedure

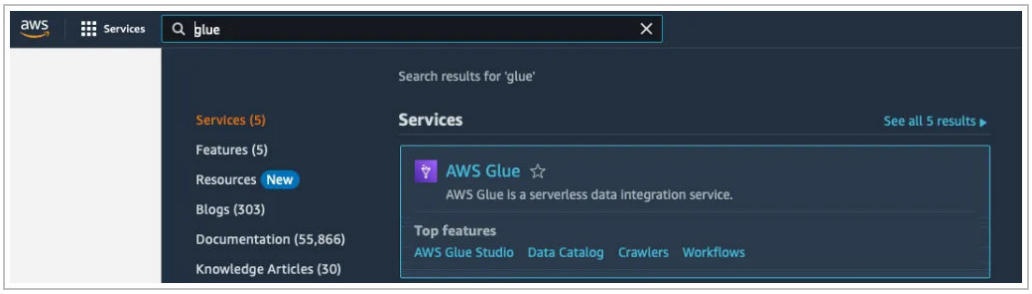
Example: Query with AWS Glue Data Catalog

This example steps you through the process of running queries over datasets stored in Amazon Simple Storage Service (Amazon S3) using an AWS Glue Data Catalog.

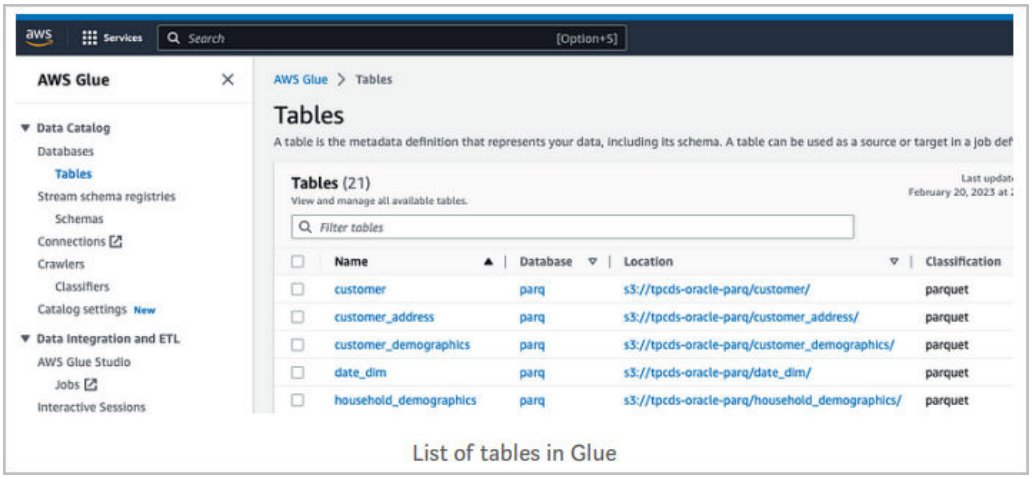
In this example, metadata in an AWS Glue Data Catalog is inspected to see what Amazon S3 objects have been previously crawled and exist in the data catalog. Autonomous Database is then associated with the AWS Glue Data Catalog and Amazon S3. The data catalog is synchronized with Autonomous Database to create external tables over the datasets stored in Amazon S3. The external tables are used to query the datasets in Amazon S3.

1. Inspect metadata in AWS Glue Data Catalog.

a. Launch the AWS Glue console.



b. Navigate to the data catalog, databases and tables to find existing objects.
In this example, some objects exist in Amazon S3 that AWS Glue has previously crawled and created tables for as shown below:



List of tables in Glue

2. Associate AWS Glue with Autonomous Database.

a. Create credentials in Autonomous Database.
The following procedure call includes the access ID and secret key to provide Autonomous Database with access to the underlying data in Amazon S3.

```
exec dbms_cloud.create_credential('CRED_AWS','<access id>', '<access
key>');
```

b. Associate the credentials with the AWS Glue Data Catalog and Amazon S3 object storage.
These procedure calls associate the data catalog and object storage, respectively, with the credentials.

```
exec dbms_dcat.set_data_catalog_credential('CRED_AWS');
exec dbms_dcat.set_object_store_credential('CRED_AWS');
```

c. Set up an AWS region where Glue is running.

```
exec dbms_dcat.set_data_catalog_conn(region => 'us-west-2',
catalog_type=>'AWS_GLUE');
```

3. Synchronize metadata to create external tables in Autonomous Database derived from AWS Glue databases and tables.

a. Now that the association is done, use the `all_glue_databases` view to find what databases are inside an AWS Glue Data Catalog.

```
select * from all_glue_databases order by name;
```

b. Use the `all_glue_tables` view to get a list of tables available for sync.

```
select * from all_glue_tables order by database_name, name;
```

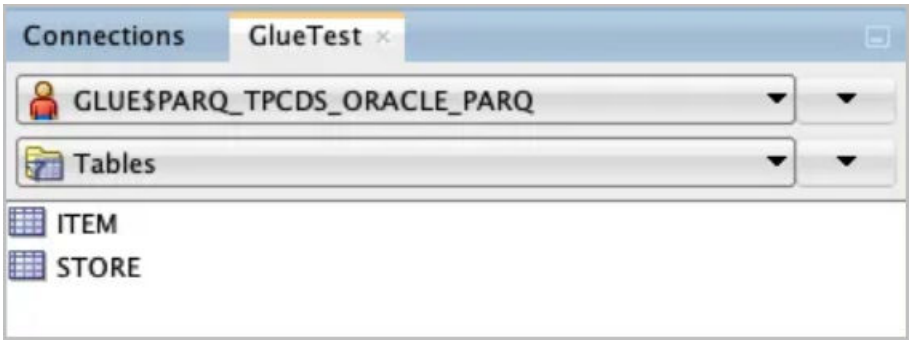
CATALO...	DATA...	NAME	TABLE_TYPE	CLASSIFICATION	DESCRIPTION	OWNER	CREATED_BY
221152...	parq	customer	EXTERNAL_TABLE	parquet	(null)	owner	arn:aws:sts::22115...
221152...	parq	customer_address	EXTERNAL_TABLE	parquet	(null)	owner	arn:aws:sts::22115...
221152...	parq	customer_demographics	EXTERNAL_TABLE	parquet	(null)	owner	arn:aws:sts::22115...
221152...	parq	date_dim	EXTERNAL_TABLE	parquet	(null)	owner	arn:aws:sts::22115...
221152...	parq	household_demographics	EXTERNAL_TABLE	parquet	(null)	owner	arn:aws:sts::22115...
221152...	parq	income_band	EXTERNAL_TABLE	parquet	(null)	owner	arn:aws:sts::22115...
221152...	parq	inventory	EXTERNAL_TABLE	parquet	(null)	owner	arn:aws:sts::22115...
221152...	parq	item	EXTERNAL_TABLE	parquet	(null)	owner	arn:aws:sts::22115...
221152...	parq	promotion	EXTERNAL_TABLE	parquet	(null)	owner	arn:aws:sts::22115...
221152...	parq	reason	EXTERNAL_TABLE	parquet	(null)	owner	arn:aws:sts::22115...
221152...	parq	ship_mode	EXTERNAL_TABLE	parquet	(null)	owner	arn:aws:sts::22115...
221152...	parq	store	EXTERNAL_TABLE	parquet	(null)	owner	arn:aws:sts::22115...
221152...	parq	store_returns	EXTERNAL_TABLE	parquet	(null)	owner	arn:aws:sts::22115...
221152...	parq	store_sales	EXTERNAL_TABLE	parquet	(null)	owner	arn:aws:sts::22115...
221152...	parq	test_glue	EXTERNAL_TABLE	(null)	(null)	221152299178	arn:aws:iam::22115...
221152...	parq	time_dim	EXTERNAL_TABLE	parquet	(null)	owner	arn:aws:sts::22115...
221152...	parq	warehouse	EXTERNAL_TABLE	parquet	(null)	owner	arn:aws:sts::22115...

c. Synchronize Autonomous Database with two tables, `store` and `item`, found in the `parq` database.

```
begin
dbms_dcat.run_sync(
    synced_objects => '
        {
            "database_list": [
                {
                "database": "parq",
                "table_list": ["store","item"]
                }
             ]
         }',
error_semantics => 'STOP_ON_ERROR');
end;
/
```

4. Inspect new objects in Autonomous Database and run a query on top of S3.

a. Use SQL Developer to view new objects created by the previous sync operation. The `GLUE$PARQ_TPCDS_ORACLE_PARQ` schema was generated and named automatically by the `dbms_dcat.run_sync` procedure call.



b. Run a SQL query over datasets stores in Amazon S3.

```
SELECT * FROM glue$parq_tpcds_oracle_parq.store;
```

Query Apache Iceberg Tables

Autonomous Database supports querying Apache Iceberg tables.

- About Querying Apache Iceberg Tables
  Autonomous Database supports querying of Apache Iceberg tables stored in Amazon Web Services (AWS) or in Oracle Cloud Infrastructure (OCI) Object Storage.
- Concepts Related to Querying Apache Iceberg Tables
  An understanding of the following concepts is helpful for querying Apache Iceberg tables.
- Examples: Querying Apache Iceberg Tables
  These examples show how to query Apache Iceberg tables on Amazon Web Services (AWS) and Oracle Cloud Infrastructure (OCI), using a data catalog and using direct URLs for the root manifest file.

About Querying Apache Iceberg Tables

Autonomous Database supports querying of Apache Iceberg tables stored in Amazon Web Services (AWS) or in Oracle Cloud Infrastructure (OCI) Object Storage.

Supported Configurations

These specific configurations are supported:

- Iceberg tables on AWS:
  - Iceberg tables registered with AWS Glue Data Catalog, created with Spark or Athena. For more information see Use the AWS Glue connector to read and write Apache Iceberg tables with ACID transactions and perform time travel and Using Iceberg tables.
  - Iceberg tables stored on AWS S3 by directly providing the URL for the root metadata file.
- Iceberg tables on OCI:
  - Iceberg tables generated with OCI Data Flow using a Hadoop Catalog.
    For more information, see Oracle Data Flow Examples and Using a Hadoop Catalog.
  - Iceberg tables stored on OCI Object Storage by directly providing the URL for the root metadata file.

Restrictions

- Partitioned Iceberg tables
  Oracle does not support Iceberg partitioned tables.
- Row-level updates of Iceberg tables
  Oracle does not support merge-on-read for Iceberg table updates. Queries encountering deleted files in the Iceberg metadata will fail. For more information about merge-on-read, see Enum RowLevelOperationMode.
- Schema evolution
  The schema for Oracle external tables is fixed and reflects the Iceberg schema version at external table creation time. Queries fail if the Iceberg metadata points to a different schema version compared to the one used to create the Iceberg table. If the Iceberg schema changes after the external table has been created, it is recommended that the external table be recreated.

Concepts Related to Querying Apache Iceberg Tables

An understanding of the following concepts is helpful for querying Apache Iceberg tables.

Iceberg Catalog

The Iceberg catalog is a service that manages table metadata, such table snapshots, the table schema and partitioning information. In order to query the latest snapshot for an Iceberg table, query engines must first access the catalog and obtain the location of the most recent metadata file. There are already a number of available catalog implementations, including AWS Glue, Hive, Nessie, and Hadoop. Autonomous Database supports the AWS Glue catalog and the HadoopCatalog implementation used by Spark.

For more information, see Optimistic Concurrency.

Metadata Files

The metadata file is a JSON document that keeps track of the table snapshots, partitioning scheme and schema information. The metadata file is the entry point to a hierarchy of manifest lists and manifest files. The manifests track the table's data files along with information including partitioning and column statistics. See the Iceberg Table Specification, for more information.

Transactions

Iceberg supports row-level updates to tables using either copy-on-write or merge-on-read. Copy-on-write generates new data files that reflect the updated rows, while merge-on-read generates new "delete files" that must be merged with the data files during reading. Oracle supports copy-on-write. Queries on iceberg tables fail if they encounter a delete file. For more information, see RowLevelOperationMode.

Schema Evolution

Iceberg supports schema evolution. Schema changes are reflected in the Iceberg metadata using a schema ID. Note that Oracle external tables have a fixed schema, determined by the most current schema version at table creation time. Iceberg queries fail when the queried metadata points to a different schema version compared to the one used at table creation time. For more information, see Schema Evolution.

Partitioning

Iceberg supports advanced partitioning options such as hidden partitioning and partition evolution that rely on processing/altering the table's metadata without costly data layout changes.

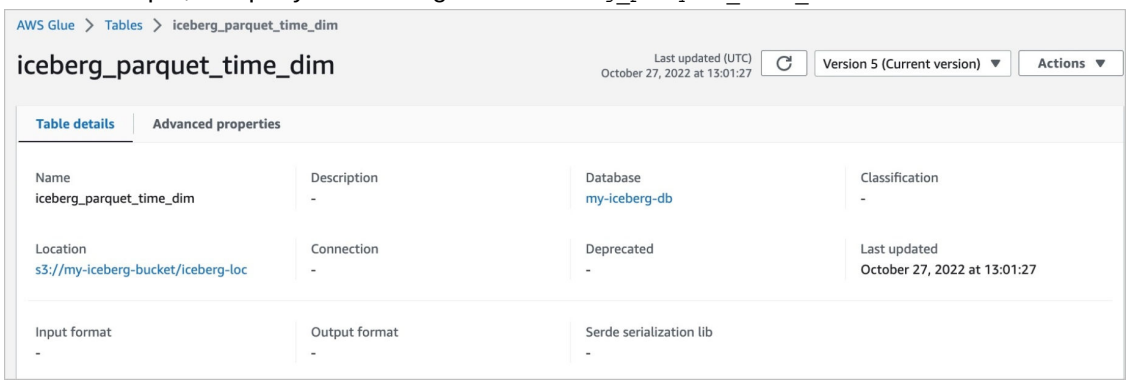
Examples: Querying Apache Iceberg Tables

These examples show how to query Apache Iceberg tables on Amazon Web Services (AWS) and Oracle Cloud Infrastructure (OCI), using a data catalog and using direct URLs for the root manifest file.

For detailed information on creating external tables for Apache Iceberg, see CREATE_EXTERNAL_TABLE Procedure for Apache Iceberg.

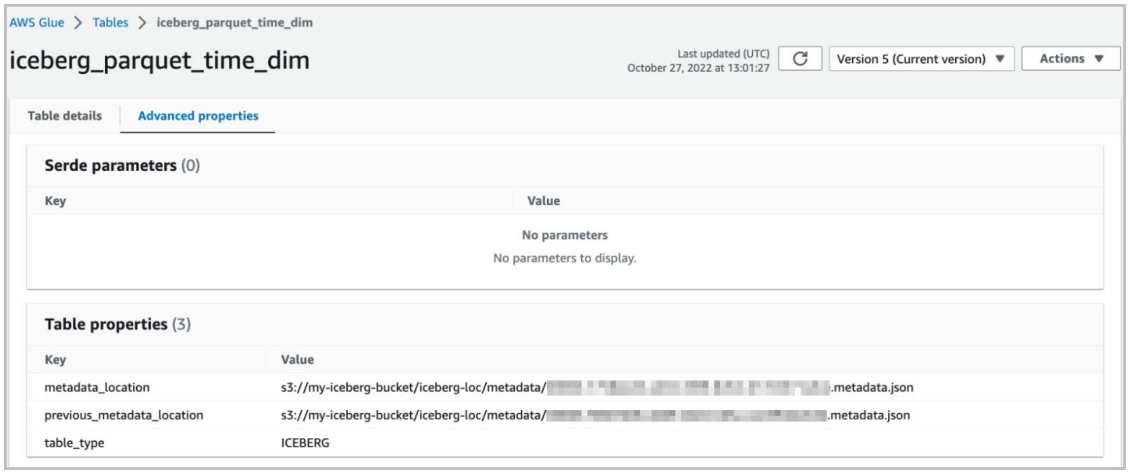
Query an Iceberg table on AWS using a Glue Data Catalog

In this example, we query the Iceberg table `iceberg_parquet_time_dim`.



The table belongs to Glue database `my-iceberg-db` and is stored in folder `s3://my-iceberg-bucket/iceberg-loc`.

The table details for `iceberg_parquet_time_dim` are shown here:



We can create an external table for `iceberg_parquet_time_dim` as follows:

```
BEGIN
DBMS_CLOUD.CREATE_EXTERNAL_TABLE (
    table_name       => 'iceberg_parquet_time_dim',
    credential_name  => 'AWS_CRED',
```

```
    file_uri_list    => '',
    format           =>
      '{"access_protocol":
        {
         "protocol_type": "iceberg",
         "protocol_config":
          {
           "iceberg_catalog_type": "aws_glue",
           "iceberg_glue_region": "us-west-1",
           "iceberg_table_path": "my-iceberg-db.iceberg_parquet_time_dim"
          }
        }
      }'
  );
END;
/
```

In the protocol_config section we specify that the table uses AWS Glue as the catalog type, and set the catalog's region to `us-west-1`.

The credential `AWS_CRED` is created using `dbms_cloud.create_credential` with an AWS API key. The Glue Data Catalog instance is determined by the account ID associated with `AWS_CRED` for region `us-west-1`, as there is a single Glue Data Catalog region for each account. The `iceberg_table_path` element in the `protocol_config` section uses a `$database_name.$table_name` path to specify the Glue table name and database name. Finally, the `column_list` and `field_list` parameters are left null because the table's schema is automatically derived from the Iceberg metadata.

For further information on creating the credential, see CREATE_CREDENTIAL Procedure. For information on AWS Glue resources, see Specifying AWS Glue resource ARNs.

Query an Iceberg table on AWS using the location of the root metadata file

If we know the location of the metadata file for an Iceberg table, we can create an external table without specifying a catalog, as follows:

```
BEGIN
DBMS_CLOUD.CREATE_EXTERNAL_TABLE (
    table_name       => 'iceberg_parquet_time_dim',
    credential_name  => 'AWS_CRED',
    file_uri_list    => 'https://my-iceberg-bucket.s3.us-west-1.amazonaws.com/
iceberg-loc/metadata/00004-1758ee2d-a204-4fd9-8d52-
d17e5371a5ce.metadata.json',
    format           =>'{"access_protocol":{"protocol_type":"iceberg"}}');
END;
/
```

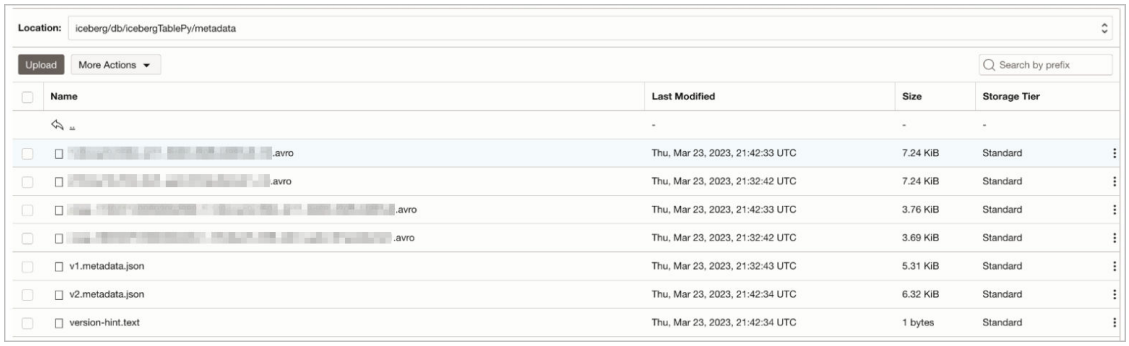
We use the `file_uri_list` parameter to specify the location of the metadata file in AWS S3 virtual-hosted-style URL format. For information on this format, see Methods for accessing an AWS S3 bucket.

In this example, the database accesses the metadata file directly, so there is no need to provide a `protocol_config` section in the `format` parameter. When using the metadata file location to create an external table, the database queries the most recent snapshot referenced by the metadata file. Subsequent updates to the Iceberg table, that create new snapshots and new metadata files, will not be visible to the database.

Query an Iceberg table that uses Hadoop Catalog on OCI

In this example, we query the Iceberg table `icebergTablePy`, created using OCI Data Flow, where Spark uses the HadoopCatalog implementation for the Iceberg catalog. HadoopCatalog uses a `warehouse` directory and puts the Iceberg metadata in a `$database_name/$table_name` subfolder under this directory. It also uses a `version-hint.text` file that contains the version number for the most recent metadata file version. See Iceberg Support on OCI Data Flow for the example on Github.

The sample table `db.icebergTablePy` was created using a `warehouse` folder, named `iceberg`, in OCI bucket `my-iceberg-bucket`. The storage layout on OCI for table `icebergTablePy` is shown below:



Create an external table for table `db.icebergTablePy` as follows:

```
BEGIN
DBMS_CLOUD.CREATE_EXTERNAL_TABLE (
    table_name       => 'iceberg_parquet_time_dim3',
    credential_name  => 'OCI_CRED',
    file_uri_list    => '',
    format           =>'{"access_protocol":{"protocol_type":"iceberg",
        "protocol_config":{"iceberg_catalog_type": "hadoop",
        "iceberg_warehouse":"https://objectstorage.uk-
cardiff-1.oraclecloud.com/n/my-tenancy/b/my-iceberg-bucket/o/iceberg",
        "iceberg_table_path": "db.icebergTablePy"}}}');
END;
/
```

Query an Iceberg table on OCI using the location of the root metadata file

We can query the Iceberg table described in the previous section by directly using the URL for the metadata file, as follows:

```
BEGIN
DBMS_CLOUD.CREATE_EXTERNAL_TABLE (
    table_name       => 'iceberg_parquet_time_dim4',
    credential_name  => 'OCI_CRED',
    file_uri_list    => 'https://objectstorage.uk-
cardiff-1.oraclecloud.com/n/my-tenancy/b/my-iceberg-bucket/o/iceberg/db/
icebergTablePy/metadata/v2.metadata.json',
    format           =>'{"access_protocol":{"protocol_type":"iceberg"}}'
    );
  END;
/
```

In this example, we use the `file_uri_list` parameter to specify the URI for the metadata file using the native OCI URI format. When using the metadata file URI, the external table always queries the latest snapshot stored in the specific file. Subsequent updates that generate new snapshots and new metadata files are not accessible to the query.

For more information on the native OCI URI format, see Cloud Object Storage URI Formats.

Validate External Data

To validate any external table, you can use the procedure `DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE`.

To validate a partitioned external table, see Validate External Partitioned Data. This procedure includes a parameter that lets you specify a specific partition to validate.

To validate a hybrid partitioned table, see Validate Hybrid Partitioned Data. This procedure includes a parameter that lets you specify a specific partition to validate.

Before validating an external table you need to create the external table. To create an external table use the procedure for your table type, either `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`. For example:

```
BEGIN
  DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE (
    table_name => 'CHANNELS_EXT' );
END;
/
```

This procedure scans your source files and validates them using the format options specified when you create the external table.

The validate operation, by default, scans all the rows in your source files and stops when a row is rejected. If you want to validate only a subset of the rows, use the `rowcount` parameter. When the `rowcount` parameter is set the validate operation scans rows and stops either when a row is rejected or when the specified number of rows are validated without errors.

For example, the following validate operation scans 100 rows and stops when a row is rejected or when 100 rows are validated without errors:

```
BEGIN
  DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE (
     table_name => 'CHANNELS_EXT',
     rowcount => 100 );
END;
/
```

If you do not want the validate to stop when a row is rejected and you want to see all rejected rows, set the `stop_on_error` parameter to `FALSE`. In this case `VALIDATE_EXTERNAL_TABLE` scans all rows and reports all rejected rows.

If you want to validate only a subset of rows use the `rowcount` parameter. When `rowcount` is set and `stop_on_error` is set to `FALSE`, the validate operation scans rows and stops either when the specified number of rows are rejected or when the specified number of rows are

validated without errors. For example, the following example scans 100 rows and stops when 100 rows are rejected or when 100 rows are validated without errors:

```
BEGIN
  DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE (
     table_name => 'CHANNELS_EXT',
     rowcount => 100
     stop_on_error => FALSE );
END;
/
```

See VALIDATE_EXTERNAL_TABLE Procedure for detailed information about `DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE` parameters.

See View Logs for Data Validation to see the results of validate operations in the tables `dba_load_operations` and `user_load_operations`.

Validate External Partitioned Data

To validate an external partitioned table, you can use the procedure `DBMS_CLOUD.VALIDATE_EXTERNAL_PART_TABLE`. This procedure includes a parameter that lets you specify a specific partition to validate.

Before validating an external partitioned table you need to create the external partitioned table. To create an external partitioned table use the procedure `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` (see Query External Partitioned Data (with Partitioning Clause) for more details):

```
BEGIN
  DBMS_CLOUD.VALIDATE_EXTERNAL_PART_TABLE (
    table_name => 'PET1',
    partition_name => 'P1');
END;
/
```

This procedure scans your source files for partition P1 and validates them using the format options specified when you create the external partitioned table.

The validation of a partitioned table by default validates all the partitions sequentially until `rowcount` is reached. If you specify a `partition_name` then only a specific partition is validated.

The validate operation, by default, scans all the rows in your source files and stops when a row is rejected. If you want to validate only a subset of the rows, use the `rowcount` parameter. When the `rowcount` parameter is set the validate operation scans rows and stops either when a row is rejected or when the specified `rowcount` number of rows are validated without errors.

For example, the following validate operation scans 100 rows and stops when a row is rejected or when 100 rows are validated without errors:

```
BEGIN
  DBMS_CLOUD.VALIDATE_EXTERNAL_PART_TABLE (
     table_name => 'PET1',
     rowcount => 100 );
END;
/
```

If you do not want the validate to stop when a row is rejected and you want to see all rejected rows, set the `stop_on_error` parameter to `FALSE`. In this case `DBMS_CLOUD.VALIDATE_EXTERNAL_PART_TABLE` scans all rows and reports all rejected rows.

If you want to validate only a subset of rows use the `rowcount` parameter. When `rowcount` is set and `stop_on_error` is set to `FALSE`, the validate operation scans rows and stops either when the specified number of rows are rejected or when the specified number of rows are validated without errors. For example, the following example scans 100 rows and stops when 100 rows are rejected or when 100 rows are validated without errors:

```
BEGIN
  DBMS_CLOUD.VALIDATE_EXTERNAL_PART_TABLE (
     table_name => 'PET1',
     rowcount => 100
     stop_on_error => FALSE );
END;
/
```

See VALIDATE_EXTERNAL_PART_TABLE Procedure for more information on `DBMS_CLOUD.VALIDATE_EXTERNAL_PART_TABLE`.

See View Logs for Data Validation to see the results of validate operations in the tables `dba_load_operations` and `user_load_operations`.

Validate Hybrid Partitioned Data

To validate a hybrid partitioned table, you can use the procedure `DBMS_CLOUD.VALIDATE_HYBRID_PART_TABLE`. This procedure includes a parameter that lets you specify a specific partition to validate.

Before validating a hybrid partitioned table you need to create the table. To create a hybrid partitioned table use the procedure `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE` (see Query Hybrid Partitioned Data for more details):

```
BEGIN
 DBMS_CLOUD.VALIDATE_HYBRID_PART_TABLE(
    table_name => 'HPT1',
    partition_name => 'P1');
END;
/
```

This procedure scans your source files for partition P1 and validates them using the format options specified when you create the hybrid partitioned table.

The validation of a hybrid partitioned table by default validates all the external partitions sequentially until `rowcount` is reached. If you specify a `partition_name` then only that specific partition is validated.

The validate operation, by default, scans all the rows in your source files and stops when a row is rejected. If you want to validate only a subset of the rows, use the `rowcount` parameter. When the `rowcount` parameter is set the validate operation scans rows and stops either when a row is rejected or when the specified `rowcount` number of rows are validated without errors.

For example, the following validate operation scans 100 rows and stops when a row is rejected or when 100 rows are validated without errors:

```
BEGIN
  DBMS_CLOUD.VALIDATE_HYBRID_PART_TABLE (
     table_name => 'HPT1',
     rowcount=>100 );
END;
/
```

If you do not want the validate to stop when a row is rejected and you want to see all rejected rows, set the `stop_on_error` parameter to `FALSE`. In this case `DBMS_CLOUD.VALIDATE_HYBRID_PART_TABLE` scans all rows and reports all rejected rows.

If you want to validate only a subset of rows use the `rowcount` parameter. When `rowcount` is set and `stop_on_error` is set to `FALSE`, the validate operation scans rows and stops either when the specified number of rows are rejected or when the specified number of rows are validated without errors. For example, the following example scans 100 rows and stops when 100 rows are rejected or when 100 rows are validated without errors:

```
BEGIN
  DBMS_CLOUD.VALIDATE_HYBRID_PART_TABLE (
     table_name => 'HPT1',
     rowcount => 100
     stop_on_error => FALSE );
END;
/
```

For detailed information about `DBMS_CLOUD.VALIDATE_HYBRID_PART_TABLE` parameters see VALIDATE_HYBRID_PART_TABLE Procedure.

See View Logs for Data Validation to see the results of validate operations in the tables `dba_load_operations` and `user_load_operations`.

View Logs for Data Validation

To validate an external table, use the procedures `DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE`, `DBMS_CLOUD.VALIDATE_EXTERNAL_PART_TABLE`, and `DBMS_CLOUD.VALIDATE_HYBRID_PART_TABLE`.

After you validate your source files you can see the result of the validate operation by querying a load operations table:

- `dba_load_operations`: shows all validate operations.
- `user_load_operations`: shows the validate operations in your schema.

You can use these files to view load validation information. For example use this select operation to query `user_load_operations`:

```
SELECT table_name, owner_name, type, status, start_time, update_time, logfile_table,
badfile_table
FROM user_load_operations
WHERE type = 'VALIDATE';

TABLE_NAME    OWNER_NAME  TYPE       STATUS     START_TIME     UPDATE_TIME
LOGFILE_TABLE    BADFILE_TABLE
```

```
------------ ----------  ---------- ---------- -------------- -------------
---------------  ---------------
CHANNELS_EXT  SH          VALIDATE   COMPLETED  13-NOV-17...   13-NOV-17...
VALIDATE$21_LOG  VALIDATE$21_BAD
```

Using this SQL statement with the `WHERE` clause on the `TYPE` column displays all of the load operations with type `VALIDATE`.

The `LOGFILE_TABLE` column shows the name of the table you can query to look at the log of a validate operation. For example, the following query shows the log for this validate operation:

```
SELECT * FROM VALIDATE$21_LOG;
```

The column `BADFILE_TABLE` shows the name of the table you can query to look at the rows where there were errors during validation. For example, the following query shows the rejected records for the above validate operation:

```
SELECT * FROM VALIDATE$21_BAD;
```

Depending on the errors shown in the log and the rows shown in the `BADFILE_TABLE`, you can correct the error by dropping the external table using the `DROP TABLE` command and recreating it by specifying the correct format options in `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`, `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` or `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE`.

> **Note:**
>
> The `LOGFILE_TABLE` and `BADFILE_TABLE` tables are stored for two days for each validate operation and then removed automatically.

Use Database Links with Autonomous Database

You can create database links from an Autonomous Database to another Autonomous Database, to other Oracle databases, or to non-Oracle databases. In addition, you can create database links from other databases to an Autonomous Database.

- Create Database Links from Autonomous Database to Another Autonomous Database
  You can create database links from Autonomous Database to another Autonomous Database that is on a private endpoint or on a public endpoint (publicly accessible).
- Create Database Links to an Oracle Database that is not an Autonomous Database
  You can create database links from an Autonomous Database to an Oracle database that is on a private endpoint or on a public endpoint (publicly accessible).
- Create Database Links to Non-Oracle Databases
  You can create database links from Autonomous Database to Non-Oracle databases.
- Create Database Links from Other Databases to Autonomous Database
  You can create database links to an Autonomous Database from an Oracle database that is not an Autonomous Database. For example from an on-premise Oracle database to an Autonomous Database
- Drop Database Links
  After you create a database link you can drop the database link.

Create Database Links from Autonomous Database to Another Autonomous Database

You can create database links from Autonomous Database to another Autonomous Database that is on a private endpoint or on a public endpoint (publicly accessible).

> **Note:**
>
> See Create Database Links to an Oracle Database that is not an Autonomous Database if the target for your database link is an Oracle Database that is not an Autonomous Database.

- Create Database Links from Autonomous Database to a Publicly Accessible Autonomous Database
  You can create database links from an Autonomous Database to a target Autonomous Database that is on a public endpoint.
- Create Database Links from Autonomous Database to an Autonomous Database on a Private Endpoint
  You can create database links from an Autonomous Database to a target Autonomous Database that is on a private endpoint.
- Database Link Notes with a Target that is an Autonomous Database
  Provides notes for creating database links to a target that is another Autonomous Database.

Create Database Links from Autonomous Database to a Publicly Accessible Autonomous Database

You can create database links from an Autonomous Database to a target Autonomous Database that is on a public endpoint.

See How to Create a Database Link from Your Autonomous Database to a Database Cloud Service Instance for more information.

- Create Database Links from Autonomous Database to a Publicly Accessible Autonomous Database without a Wallet (TLS)
  You can create database links from an Autonomous Database to a publicly accessible Autonomous Database without a wallet (TLS).
- Create Database Links from Autonomous Database to a Publicly Accessible Autonomous Database with a Wallet (mTLS)
  You can to create database links from an Autonomous Database to a publicly accessible Autonomous Database with a wallet (mTLS).

Create Database Links from Autonomous Database to a Publicly Accessible Autonomous Database without a Wallet (TLS)

You can create database links from an Autonomous Database to a publicly accessible Autonomous Database without a wallet (TLS).

To create database links to a public target, the target database must be accessible. Some databases, including Autonomous Databases, may limit access (for example, using Access Control Lists). Make sure you enable your target database to allow access from your source

database for the database link to work. If you limit access with Access Control Lists (ACLs), you can find the outbound IP address of your source Autonomous Database and allow that IP address to connect to your target database. When the target database is another Autonomous Database, you can add the outbound IP address of the source database to the ACL of the target database.

See Obtain Tenancy Details for information on finding the outbound IP address.

To create a database link to a target Autonomous Database without a wallet (TLS):

1. If you have not already done so, enable TLS connections on your Autonomous Database instance.

   See Update your Autonomous Database Instance to Allow both TLS and mTLS Authentication for details.

2. On the Autonomous Database instance where you are creating the database link, create credentials to access the target Autonomous Database. The `username` and `password` you specify with `DBMS_CLOUD.CREATE_CREDENTIAL` are the credentials for the target database (you use these credentials to create the database link).

   > **Note:**
   >
   > Supplying the `credential_name` parameter is required.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
           credential_name => 'DB_LINK_CRED',
           username => 'NICK',
           password => 'password' );
   END;
   /
   ```

   The characters in the `username` parameter must be all uppercase letters.

   > **Note:**
   >
   > You can use a vault secret credential for the target database credential in a database link. See Use Vault Secret Credentials for more information.

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name.

3. Create a database link to the target Autonomous Database instance using `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

   For example:

   ```
   BEGIN
       DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
           db_link_name => 'SALESLINK',
           hostname => 'adb.eu-frankfurt-1.oraclecloud.com',
   ```

```
        port => '1521',
        service_name => 'example_medium.adb.example.oraclecloud.com',
        credential_name => 'DB_LINK_CRED',
        directory_name => NULL);
END;
/
```

To create a database link with `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` to a target Autonomous Database on a public endpoint using a secure TCP connection without a wallet, the `directory_name` parameter must be `NULL`.

The `ssl_server_cert_dn` parameter can either be omitted or if included, specify a `NULL` value.

Users other than ADMIN require privileges to run `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

4. Use the database link you created to access data on the target database.

   For example:

   ```
   SELECT * FROM employees@SALESLINK;
   ```

For the credentials you create in Step 1, the target database credentials, if the password of the target user changes you can update the credential that contains the target user's credentials as follows:

```
BEGIN
    DBMS_CLOUD.UPDATE_CREDENTIAL (
          credential_name => 'DB_LINK_CRED',
          attribute => 'PASSWORD',
          value => 'password' );
END;
/
```

Where *password* is the new password.

After this operation, the existing database links that use this credential continue to work without having to drop and recreate the database links.

> **Note:**
>
> You can create links to Big Data Service using `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`. See Query Big Data Service Hadoop (HDFS) Data from Autonomous Database for more information.

For additional information, see:

- CREATE_DATABASE_LINK Procedure
- UPDATE_CREDENTIAL Procedure

Create Database Links from Autonomous Database to a Publicly Accessible Autonomous Database with a Wallet (mTLS)

You can to create database links from an Autonomous Database to a publicly accessible Autonomous Database with a wallet (mTLS).

To create database links to a public target, the target database must be accessible. Some databases, including Autonomous Databases, may limit access (for example, using Access Control Lists). Make sure you enable your target database to allow access from your source database for the database link to work. If you limit access with Access Control Lists (ACLs), you can find the outbound IP address of your source Autonomous Database and allow that IP address to connect to your target database. For example, if the target database is another Autonomous Database, you can add the outbound IP address of the source database to the ACL of the target database.

See Obtain Tenancy Details for information on finding the outbound IP address.

To create database links to a target Autonomous Database with a wallet (mTLS):

1. Copy your target database wallet, `cwallet.sso`, containing the certificates for the target database to Object Store.

   Note the following for the wallet file:

   - The wallet file, along with the Database user ID and password provide access to data in the target Oracle Database. Store wallet files in a secure location. Share wallet files only with authorized users.
   - Do not rename the wallet file. The wallet file in Object Storage must be named `cwallet.sso`.

2. Create credentials to access your Object Store where you store the wallet file `cwallet.sso`. See CREATE_CREDENTIAL Procedure for information about the username and password parameters for different object storage services.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

3. Create a directory on Autonomous Database for the wallet file `cwallet.sso`.

   For example:

   ```
   CREATE DIRECTORY dblink_wallet_dir AS 'directory_path_of_your_choice';
   ```

   See Create Directory in Autonomous Database for information on creating directories.

4. Use `DBMS_CLOUD.GET_OBJECT` to upload the target database wallet to the directory you created in the previous step, *DBLINK_WALLET_DIR*.

   For example:

   ```
   BEGIN
       DBMS_CLOUD.GET_OBJECT(
           credential_name => 'DEF_CRED_NAME',
           object_uri => 'https://objectstorage.us-
   phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/cwallet.sso',
           directory_name => 'DBLINK_WALLET_DIR');
   ```

```
END;
/
```

In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

> **Note:**
>
> The `credential_name` you use in this step is the credentials for the Object Store. In the next step you create the credentials to access the target database.

Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

5. On the Autonomous Database instance, create credentials to access the target database. The `username` and `password` you specify with `DBMS_CLOUD.CREATE_CREDENTIAL` are the credentials for the target database that you use to create the database link.

> **Note:**
>
> Supplying the `credential_name` parameter is required.

For example:

```
BEGIN
    DBMS_CLOUD.CREATE_CREDENTIAL(
        credential_name => 'DB_LINK_CRED',
        username => 'NICK',
        password => 'password');
END;
/
```

The characters in the `username` parameter must be all uppercase letters.

> **Note:**
>
> You can use a vault secret credential for the target database credential in a database link. See Use Vault Secret Credentials for more information.

This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name.

6. Create the database link to the target database using `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

For example:

```
BEGIN
    DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
        db_link_name => 'SALESLINK',
        hostname => 'adb.eu-frankfurt-1.oraclecloud.com',
        port => '1522',
        service_name => 'example_medium.adb.example.oraclecloud.com',
        ssl_server_cert_dn => 'CN=adb.example.oraclecloud.com,OU=Oracle
BMCS FRANKFURT,O=Oracle Corporation,L=Redwood City,ST=California,C=US',
        credential_name => 'DB_LINK_CRED',
        directory_name => 'DBLINK_WALLET_DIR');
END;
/
```

Users other than ADMIN require privileges to run `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

If the wallet file in the directory specified with `directory_name` is not `cwallet.sso`, the procedure reports an error such as: `ORA-28759: failure to open file`.

7. Use the database link you created to access data on the target database.

   For example:

   ```
   SELECT * FROM employees@SALESLINK;
   ```

For the credentials you create in Step 5, the target database credentials, if the password of the target user changes you can update the credential that contains the target user's credentials as follows:

```
BEGIN
    DBMS_CLOUD.UPDATE_CREDENTIAL (
        credential_name => 'DB_LINK_CRED',
        attribute => 'PASSWORD',
        value => 'password' );
END;
/
```

Where *password* is the new password.

After this operation, the existing database links that use this credential continue to work without having to drop and recreate the database links.

> **Note:**
>
> You can create links to Big Data Service using `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`. See Query Big Data Service Hadoop (HDFS) Data from Autonomous Database for more information.

For additional information, see:

- CREATE_DATABASE_LINK Procedure

- GET_OBJECT Procedure and Function
- UPDATE_CREDENTIAL Procedure

Create Database Links from Autonomous Database to an Autonomous Database on a Private Endpoint

You can create database links from an Autonomous Database to a target Autonomous Database that is on a private endpoint.

Depending on configuration of the target Autonomous Database, you have these options:

- Create Database Links to a Target Autonomous Database on a Private Endpoint without a Wallet (TLS)
- Create Database Links to a Target Autonomous Database on a Private Endpoint with a Wallet (mTLS)

See How to Create a Database Link from Your Autonomous Database to a Database Cloud Service Instance for more information.

Topics

- Prerequisites for Database Links from Autonomous Database to a Target Autonomous Database on a Private Endpoint
  Lists the prerequisites to create database links to a target Autonomous Database that is on a private endpoint.
- Create Database Links to a Target Autonomous Database on a Private Endpoint without a Wallet (TLS)
  You can create database links from an Autonomous Database to a target Autonomous Database that is on a private endpoint and connect without a wallet (TLS).
- Create Database Links to a Target Autonomous Database on a Private Endpoint with a Wallet (mTLS)
  You can create database links from an Autonomous Database to a target Autonomous Database that is on a private endpoint (mTLS).

Prerequisites for Database Links from Autonomous Database to a Target Autonomous Database on a Private Endpoint

Lists the prerequisites to create database links to a target Autonomous Database that is on a private endpoint.

To create a database link to a target Autonomous Database on a private endpoint:

- The target database must be accessible from the source database's Oracle Cloud Infrastructure VCN. For example, you can connect to the target database when:
  - The target database is on a private endpoint.
  - Both the source database and the target database are in the same Oracle Cloud Infrastructure VCN.
  - The source database and the target database are in different Oracle Cloud Infrastructure VCNs that are paired.
  - The target database is connected to the source database's Oracle Cloud Infrastructure VCN using FastConnect or VPN.

- For a target on a private endpoint, `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` supports specifying a single hostname with the `hostname` parameter. On a private endpoint, using an IP address, SCAN IP, or a SCAN hostname is not supported (when the target is on a public endpoint, `CREATE_DATABASE_LINK` supports using an IP address, a SCAN IP, or a SCAN hostname).
- `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` does not support a value of `localhost` for the `hostname` parameter.
- The following ingress and egress rules must be defined for the private endpoint:
  - Define an egress rule in the source database's subnet security list or network security group such that the traffic over TCP is allowed to the target database's IP address and port number.
  - Define an ingress rule in the target database's subnet security list or network security group such that the traffic over TCP is allowed from the source database IP address to the destination port.

  See Configure Network Access with Private Endpoints for information on configuring private endpoints with ingress and egress rules.

> **Note:**
>
> When your Autonomous Database instance is configured with a private endpoint, set the `ROUTE_OUTBOUND_CONNECTIONS` database property to '`PRIVATE_ENDPOINT`' to specify that all outgoing database links are subject to the Autonomous Database instance private endpoint VCN's egress rules. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

Create Database Links to a Target Autonomous Database on a Private Endpoint without a Wallet (TLS)

You can create database links from an Autonomous Database to a target Autonomous Database that is on a private endpoint and connect without a wallet (TLS).

Perform the prerequisite steps, as required. See Prerequisites for Database Links from Autonomous Database to a Target Autonomous Database on a Private Endpoint for details.

To create a database link to a target Autonomous Database on a private endpoint without a wallet:

1. If you have not already done so, enable TLS connections on your Autonomous Database instance.

   See Update your Autonomous Database Instance to Allow both TLS and mTLS Authentication for details.

2. Create credentials to access the target Autonomous Database instance. The `username` and `password` you specify with `DBMS_CLOUD.CREATE_CREDENTIAL` are the credentials for the target database used within the database link, (where the target database is accessed through the VCN).

   For example:

```
BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
         credential_name => 'PRIVATE_ENDPOINT_CRED',
```

```
        username => 'NICK',
        password => 'password'
        );
END;
/
```

The characters in the `username` parameter must be all uppercase letters.

> **Note:**
>
> You can use a vault secret credential for the target database credential in a database link. See Use Vault Secret Credentials for more information.

This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name.

3. Create the database link to the target database using `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

   For example:

```
BEGIN
     DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
         db_link_name => 'PRIVATE_ENDPOINT_LINK',
         hostname => 'exampleHostname',
         port => '1521',
         service_name => 'example_high.adb.oraclecloud.com',
         credential_name => 'PRIVATE_ENDPOINT_CRED',
         directory_name => NULL,
         private_target => TRUE);
END;
/
```

   For a target on a private endpoint, `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` supports specifying a single hostname with the `hostname` parameter. On a private endpoint, using an IP address, SCAN IP, or a SCAN hostname is not supported (when the target is on a public endpoint, `CREATE_DATABASE_LINK` supports using an IP address, a SCAN IP, or a SCAN hostname).

   Users other than ADMIN require privileges to run `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

   As shown in the example, to create a database link with `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` to a target database on a private endpoint without a wallet, all of the following are required:

   - The `directory_name` parameter must be `NULL`.
   - The `ssl_server_cert_dn` parameter can either be omitted or if included, specify a `NULL` value.
   - The `private_target` parameter must be `TRUE`.

> **Note:**
>
> If you set `ROUTE_OUTBOUND_CONNECTIONS` to `PRIVATE_ENDPOINT`, setting the `private_target` parameter to `TRUE` is not required in this API. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

4. Use the database link you created to access data in the target database.

   For example:

   ```
   SELECT * FROM employees@PRIVATE_ENDPOINT_LINK;
   ```

> **Note:**
>
> For the credentials you create in Step 1, the Oracle Database credentials, if the password of the target user changes you can update the credential that contains the target user's credentials as follows:
>
> ```
> BEGIN
>      DBMS_CLOUD.UPDATE_CREDENTIAL (
>          credential_name => 'DB_LINK_CRED',
>          attribute => 'PASSWORD',
>          value => 'password');
> END;
> /
> ```
>
> Where *password* is the new password.
>
> After this operation, the existing database links that use this credential continue to work without having to drop and recreate the database links.

See CREATE_DATABASE_LINK Procedure for additional information.

Create Database Links to a Target Autonomous Database on a Private Endpoint with a Wallet (mTLS)

You can create database links from an Autonomous Database to a target Autonomous Database that is on a private endpoint (mTLS).

Perform the prerequisite steps, as required. See Prerequisites for Database Links from Autonomous Database to a Target Autonomous Database on a Private Endpoint for details.

To create a database link to a target Autonomous Database on a private endpoint, with a wallet:

1. Copy your target database wallet, `cwallet.sso`, containing the certificates for the target database to Object Store.

> **Note:**
>
> The wallet file, along with the Database user ID and password provide access to data in the target Oracle database. Store wallet files in a secure location. Share wallet files only with authorized users.

2. Create credentials to access your Object Store where you store the `cwallet.sso`. See CREATE_CREDENTIAL Procedure for information about the username and password parameters for different object storage services.
3. Create a directory on Autonomous Database for the wallet file `cwallet.sso`.

   For example:

   ```
   CREATE DIRECTORY wallet_dir AS 'directory_path_of_your_choice';
   ```

   See Create Directory in Autonomous Database for information on creating directories.
4. Use `DBMS_CLOUD.GET_OBJECT` to upload the target database wallet to the directory you created in the previous step, *WALLET_DIR*.

   For example:

   ```
   BEGIN
       DBMS_CLOUD.GET_OBJECT(
           credential_name => 'DEF_CRED_NAME',
           object_uri => 'https://objectstorage.us-
   phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/cwallet.sso',
           directory_name => 'WALLET_DIR');
   END;
   /
   ```

   In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

   > **Note:**
   >
   > The `credential_name` you use in this step is the credentials for the Object Store. In the next step you create the credentials to access the target database.

5. On Autonomous Database create credentials to access the target database. The `username` and `password` you specify with `DBMS_CLOUD.CREATE_CREDENTIAL` are the credentials for the target database used within the database link, (where the target database is accessed through the VCN).

   > **Note:**
   >
   > Supplying the `credential_name` parameter is required.

For example:

```
BEGIN
    DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'DB_LINK_CRED',
       username => 'NICK',
       password => 'password');
END;
/
```

The characters in the `username` parameter must be all uppercase letters.

> **Note:**
>
> You can use a vault secret credential for the target database credential in a database link. See Use Vault Secret Credentials for more information.

This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name.

6. Create the database link to the target database using `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

   For example:

```
BEGIN
    DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
        db_link_name => 'PEDBLINK1',
        hostname => 'example1.adb.ap-osaka-1.oraclecloud.com',
        port => '1522',
        service_name => 'example_high.adb.oraclecloud.com',
        ssl_server_cert_dn => 'ssl_server_cert_dn',
        credential_name => 'DB_LINK_CRED',
        directory_name => 'WALLET_DIR',
        private_target => TRUE);
END;
/
```

> **Note:**
>
> If you set `ROUTE_OUTBOUND_CONNECTIONS` to `PRIVATE_ENDPOINT`, setting the `private_target` parameter to `TRUE` is not required in this API. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

For a target on a private endpoint, `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` supports specifying a single hostname with the `hostname` parameter. On a private endpoint, using an IP address, SCAN IP, or a SCAN hostname is not supported (when the target is on a public endpoint, `CREATE_DATABASE_LINK` supports using an IP address, a SCAN IP, or a SCAN hostname).

`DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` does not support a value of `localhost` for the `hostname` parameter.

Users other than ADMIN require privileges to run `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

7. Use the database link you created to access data in the target database.

   For example:

   ```
   SELECT * FROM employees@PEDBLINK1;
   ```

> **Note:**
>
> For the credentials you create in Step 5, the Oracle Database credentials, if the password of the target user changes you can update the credential that contains the target user's credentials as follows:
>
> ```
> BEGIN
>     DBMS_CLOUD.UPDATE_CREDENTIAL (
>         credential_name => 'DB_LINK_CRED',
>         attribute => 'PASSWORD',
>         value => 'password');
> END;
> /
> ```
>
> Where *password* is the new password.
>
> After this operation, the existing database links that use this credential continue to work without having to drop and recreate the database links.

See CREATE_DATABASE_LINK Procedure for additional information.

Database Link Notes with a Target that is an Autonomous Database

Provides notes for creating database links to a target that is another Autonomous Database.

Notes for database links to another Autonomous Database:

- Only one wallet file is valid per directory for use with database links. You can only upload one `cwallet.sso` at a time to the directory you choose for wallet files (for example *DBLINK_WALLET_DIR*). This means with a `cwallet.sso` in *DBLINK_WALLET_DIR* you can only create database links to the databases for which the wallet in that directory is valid. To use multiple `cwallet.sso` files with database links you need to create additional directories and put each `cwallet.sso` in a different directory. When you create database links with `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`, specify the directory that contains the wallet with the `directory_name` parameter.

  See Create Directory in Autonomous Database for information on creating directories.
- To list the database links, use the `ALL_DB_LINKS` view. See ALL_DB_LINKS for more information.
- The wallet file, along with the Database user ID and password provide access to data in the target Autonomous Database. Store wallet files in a secure location. Share wallet files only with authorized users.

Create Database Links to an Oracle Database that is not an Autonomous Database

You can create database links from an Autonomous Database to an Oracle database that is on a private endpoint or on a public endpoint (publicly accessible).

> **Note:**
>
> See Create Database Links from Autonomous Database to Another Autonomous Database if the target for your database link is another Autonomous Database.

- Create Database Links from Autonomous Database to a Publicly Accessible Oracle Database with a Wallet (mTLS)
  You can create database links from an Autonomous Database to a target Oracle database that is on a public endpoint.
- Create Database Links from Autonomous Database to an Oracle Database on a Private Endpoint
  You can create database links from an Autonomous Database to a target Oracle Database that is on a private endpoint.
- Database Link Notes with a Target Oracle Database
  Provides notes for creating database links to a target Oracle database (when the target is not an Autonomous Database)

Create Database Links from Autonomous Database to a Publicly Accessible Oracle Database with a Wallet (mTLS)

You can create database links from an Autonomous Database to a target Oracle database that is on a public endpoint.

To use database links with Autonomous Database the target database must be configured to use TCP/IP with SSL (TCPS) authentication. Autonomous Databases use TCP/IP with SSL (TCPS) authentication by default, so you do not need to do any additional configuration in your target database to link to another Autonomous Database. Other Oracle databases must be configured to use TCP/IP with SSL (TCPS) authentication. See Configuring Secure Sockets Layer Authentication for more information.

To create database links to a public target, the target Oracle Database must be accessible. Some databases may limit access (for example, using Access Control Lists). Make sure you enable your target database to allow access from your source database for the database link to work. If you limit access with Access Control Lists (ACLs), you can find the outbound IP address of your source Autonomous Database and allow that IP address to connect to your target database.

See How to Create a Database Link from Your Autonomous Database to a Database Cloud Service Instance for more information.

To create database links to a target Oracle database with a wallet (mTLS):

1. Copy your target database wallet, `cwallet.sso`, containing the certificates for the target database to Object Store.

   Note the following for the wallet file:

- The wallet file, along with the Database user ID and password provide access to data in the target Oracle Database. Store wallet files in a secure location. Share wallet files only with authorized users.
- Do not rename the wallet file. The wallet file in Object Storage must be named `cwallet.sso`.

2. Create credentials to access your Object Store where you store the wallet file `cwallet.sso`. See CREATE_CREDENTIAL Procedure for information about the username and password parameters for different object storage services.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

3. Create a directory on Autonomous Database for the wallet file `cwallet.sso`.

   For example:

   ```
   CREATE DIRECTORY dblink_wallet_dir AS 'directory_path_of_your_choice';
   ```

   See Create Directory in Autonomous Database for information on creating directories.

4. Use `DBMS_CLOUD.GET_OBJECT` to upload the target database wallet to the directory you created in the previous step, *DBLINK_WALLET_DIR*.

   For example:

   ```
   BEGIN
       DBMS_CLOUD.GET_OBJECT(
           credential_name => 'DEF_CRED_NAME',
           object_uri => 'https://objectstorage.us-
   phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/cwallet.sso',
           directory_name => 'DBLINK_WALLET_DIR');
   END;
   /
   ```

   In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

   > **Note:**
   >
   > The `credential_name` you use in this step is the credentials for the Object Store. In the next step you create the credentials to access the target database.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

5. On the Autonomous Database instance, create credentials to access the target database. The `username` and `password` you specify with `DBMS_CLOUD.CREATE_CREDENTIAL` are the credentials for the target database that you use to create the database link.

> **Note:**
>
> Supplying the `credential_name` parameter is required.

For example:

```
BEGIN
    DBMS_CLOUD.CREATE_CREDENTIAL(
        credential_name => 'DB_LINK_CRED',
        username => 'NICK',
        password => 'password');
END;
/
```

The characters in the `username` parameter must be all uppercase letters.

> **Note:**
>
> You can use a vault secret credential for the target database credential in a database link. See Use Vault Secret Credentials for more information.

This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name.

6. Create the database link to the target database using `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

   For example:

```
BEGIN
    DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
        db_link_name => 'SALESLINK',
        hostname => 'adb.eu-frankfurt-1.oraclecloud.com',
        port => '1522',
        service_name => 'example_medium.adb.example.oraclecloud.com',
        ssl_server_cert_dn => 'CN=adb.example.oraclecloud.com,OU=Oracle
BMCS FRANKFURT,O=Oracle Corporation,L=Redwood City,ST=California,C=US',
        credential_name => 'DB_LINK_CRED',
        directory_name => 'DBLINK_WALLET_DIR');
END;
/
```

   Users other than ADMIN require privileges to run `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

   If the wallet file in the directory specified with `directory_name` is not `cwallet.sso`, the procedure reports an error such as: `ORA-28759: failure to open file`.

7. Use the database link you created to access data on the target database.

For example:

```
SELECT * FROM employees@SALESLINK;
```

For the credentials you create in Step 5, the target database credentials, if the password of the target user changes you can update the credential that contains the target user's credentials as follows:

```
BEGIN
    DBMS_CLOUD.UPDATE_CREDENTIAL (
        credential_name => 'DB_LINK_CRED',
        attribute => 'PASSWORD',
        value => 'password' );
END;
/
```

Where *password* is the new password.

After this operation, the existing database links that use this credential continue to work without having to drop and recreate the database links.

> **Note:**
>
> You can create links to Big Data Service using `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`. See Query Big Data Service Hadoop (HDFS) Data from Autonomous Database for more information.

For additional information, see:

- CREATE_DATABASE_LINK Procedure
- GET_OBJECT Procedure and Function
- UPDATE_CREDENTIAL Procedure

Create Database Links from Autonomous Database to an Oracle Database on a Private Endpoint

You can create database links from an Autonomous Database to a target Oracle Database that is on a private endpoint.

Depending on the type and the configuration of the target Oracle database:

- **Another Oracle Database, such as on-premises or a Database Cloud Service database, on a private endpoint that is configured for SSL (TCPS)**: In this case you can create the database link with a wallet, and the database link communicates with TCPS. See Create Database Links from Autonomous Database to Oracle Databases on a Private Endpoint with a Wallet (mTLS) for details:
- **Oracle Database, such as on-premises or a Database Cloud Service database, on a private endpoint that is configured for TCP**: In this case you create the database link without a wallet and the database link communicates with TCP. See Create Database Links to Oracle Databases on a Private Endpoint without a Wallet for details

See How to Create a Database Link from Your Autonomous Database to a Database Cloud Service Instance for more information.

Topics

- Prerequisites for Database Links from Autonomous Database to Oracle Databases on a Private Endpoint
  Lists the prerequisites to create database links from an Autonomous Database to a target Oracle database that is on a private endpoint.
- Create Database Links to Oracle Databases on a Private Endpoint without a Wallet
  Use `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` to create database links from an Autonomous Database to a target Oracle database that is on a private endpoint and connect without a wallet (TCP).
- Create Database Links from Autonomous Database to Oracle Databases on a Private Endpoint with a Wallet (mTLS)
  You can create database links from an Autonomous Database to a target Oracle database that is on a private endpoint.

Prerequisites for Database Links from Autonomous Database to Oracle Databases on a Private Endpoint

Lists the prerequisites to create database links from an Autonomous Database to a target Oracle database that is on a private endpoint.

To create a database link to a target Oracle database on a private endpoint:

- The target database must be accessible from the source database's Oracle Cloud Infrastructure VCN. For example, you can connect to the target database when:
  - The target database is on a private endpoint.
  - Both the source database and the target database are in the same Oracle Cloud Infrastructure VCN.
  - The source database and the target database are in different Oracle Cloud Infrastructure VCNs that are paired.
  - The target database is an on-premises database that is connected to the source database's Oracle Cloud Infrastructure VCN using FastConnect or VPN.
- There are two options to specify the target database, use the `hostname` parameter or the `rac_hostnames` parameter:
  - For a target on a private endpoint, `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` supports specifying a single hostname with the `hostname` parameter. On a private endpoint, using an IP address, SCAN IP, or a SCAN hostname is not supported (when the target is on a public endpoint, `CREATE_DATABASE_LINK` supports using an IP address, a SCAN IP, or a SCAN hostname).
  - When the target is an Oracle RAC database, use the `rac_hostnames` parameter to specify one or more hostnames with `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`. This allows you to take advantage of the high availability capabilities of Oracle RAC. Using an IP address, a SCAN IP, or a SCAN hostname in the `rac_hostnames` value is not supported.

    When you specify a list of host names in the `rac_hostnames` parameter, `CREATE_DATABASE_LINK` uses all of the specified host names as addresses in the connect string. If one of the specified hosts is not available on the target Oracle RAC database, Autonomous Database automatically attempts to connect using another host name from the list.

- The following ingress and egress rules must be defined for the private endpoint:
  - Define an egress rule in the source database's subnet security list or network security group such that the traffic over TCP is allowed to the target database's IP address and port number.
  - Define an ingress rule in the target database's subnet security list or network security group such that the traffic over TCP is allowed from the source database IP address to the destination port.

  See Configure Network Access with Private Endpoints for information on configuring private endpoints with ingress and egress rules.

> **Note:**
>
> When your Autonomous Database instance is configured with a private endpoint, set the `ROUTE_OUTBOUND_CONNECTIONS` database parameter to '`PRIVATE_ENDPOINT`' to specify that all outgoing database links are subject to the Autonomous Database instance private endpoint VCN's egress rules. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

Create Database Links to Oracle Databases on a Private Endpoint without a Wallet

Use `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` to create database links from an Autonomous Database to a target Oracle database that is on a private endpoint and connect without a wallet (TCP).

> **Note:**
>
> This option is for target Oracle databases that are on a private endpoint and do not have SSL/TCPS configured.

Perform the prerequisite steps, as required. See Prerequisites for Database Links from Autonomous Database to a Target Autonomous Database on a Private Endpoint for details.

To create a database link to a target database on a private endpoint using a secure TCP connection without a wallet:

1. On Autonomous Database create credentials to access the target database. The `username` and `password` you specify with `DBMS_CLOUD.CREATE_CREDENTIAL` are the credentials for the target database used within the database link, (where the target database is accessed through the VCN).

   For example:

```
BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
         credential_name => 'PRIVATE_ENDPOINT_CRED',
         username => 'NICK',
         password => 'password'
         );
END;
/
```

The characters in the `username` parameter must be all uppercase letters.

> **Note:**
>
> You can use a vault secret credential for the target database credential in a database link. See Use Vault Secret Credentials for more information.

This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name.

2. Create the database link to the target database using `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

   For example:

```
BEGIN
     DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
         db_link_name => 'PRIVATE_ENDPOINT_LINK',
         hostname => 'exampleHostname',
         port => '1522',
         service_name => 'exampleServiceName',
         ssl_server_cert_dn => NULL,
         credential_name => 'PRIVATE_ENDPOINT_CRED',
         directory_name => NULL,
         private_target => TRUE);
END;
/
```

   For a target on a private endpoint, `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` supports specifying a single hostname with the `hostname` parameter. On a private endpoint, using an IP address, SCAN IP, or a SCAN hostname is not supported (when the target is on a public endpoint, `CREATE_DATABASE_LINK` supports using an IP address, a SCAN IP, or a SCAN hostname).

   When the target is an Oracle RAC database, use the `rac_hostnames` parameter to specify one or more hostnames with `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`. This allows you to take advantage of the high availability capabilities of Oracle RAC. Using an IP address, a SCAN IP, or a SCAN hostname in the `rac_hostnames` value is not supported.

   For example, with a target Oracle RAC database use the `rac_hostnames` parameter:

```
BEGIN
     DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
         db_link_name => 'PRIVATE_ENDPOINT_LINK',
         rac_hostnames => '["sales1-svr1.example.adb.us-
ashburn-1.oraclecloud.com",
                           "sales1-svr2.example.adb.us-
ashburn-1.oraclecloud.com",
                           "sales1-svr3.example.adb.us-
ashburn-1.oraclecloud.com"]',
         port => '1522',
         service_name => 'exampleServiceName',
         ssl_server_cert_dn => NULL,
         credential_name => 'PRIVATE_ENDPOINT_CRED',
         directory_name => NULL,
```

```
        private_target => TRUE);
END;
/
```

`DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` does not support a value of `localhost` for the `hostname` or in the `rac_hostnames` parameter.

Users other than ADMIN require privileges to run `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

As shown in the example, to create a database link with `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` to a target database on a private endpoint using a secure TCP connection without a wallet, all of the following are required:

- The `directory_name` parameter must be `NULL`.
- The `ssl_server_cert_dn` parameter must be `NULL`.
- The `private_target` parameter must be `TRUE`.

> **Note:**
>
> If you set `ROUTE_OUTBOUND_CONNECTIONS` to `PRIVATE_ENDPOINT`, setting the `private_target` parameter to `TRUE` is not required in this API. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

3. Use the database link you created to access data in the target database.

   For example:

```
SELECT * FROM employees@PRIVATE_ENDPOINT_LINK;
```

> **Note:**
>
> For the credentials you create in Step 1, the Oracle Database credentials, if the password of the target user changes you can update the credential that contains the target user's credentials as follows:
>
> ```
> BEGIN
>      DBMS_CLOUD.UPDATE_CREDENTIAL (
>          credential_name => 'DB_LINK_CRED',
>          attribute => 'PASSWORD',
>          value => 'password');
> END;
> /
> ```
>
> Where *password* is the new password.
>
> After this operation, the existing database links that use this credential continue to work without having to drop and recreate the database links.

See CREATE_DATABASE_LINK Procedure for additional information.

Create Database Links from Autonomous Database to Oracle Databases on a Private Endpoint with a Wallet (mTLS)

You can create database links from an Autonomous Database to a target Oracle database that is on a private endpoint.

> **Note:**
>
> This option is for target Oracle databases that have SSL/TCPS configured and that are on a private endpoint.

If the target Oracle database does not have SSL/TCPS configured, you have two options:

- You can configure the target Oracle database to use TCP/IP with SSL (TCPS) authentication. See Configuring Transport Layer Security Authentication for information on configuring SSL/TCPS.
- You can connect to the target Oracle database with TCP. See Create Database Links to Oracle Databases on a Private Endpoint without a Wallet for details.

Perform the prerequisite steps, as required. See Prerequisites for Database Links from Autonomous Database to a Target Autonomous Database on a Private Endpoint for details.

To create a database link to a target Oracle database on a private endpoint using TCP/IP with SSL (TCPS) authentication:

1. Copy your target database wallet, `cwallet.sso`, containing the certificates for the target database to Object Store.

   > **Note:**
   >
   > The wallet file, along with the Database user ID and password provide access to data in the target Oracle database. Store wallet files in a secure location. Share wallet files only with authorized users.

2. Create credentials to access your Object Store where you store the `cwallet.sso`. See CREATE_CREDENTIAL Procedure for information about the username and password parameters for different object storage services.
3. Create a directory on Autonomous Database for the wallet file `cwallet.sso`.

   For example:

   ```
   CREATE DIRECTORY wallet_dir AS 'directory_path_of_your_choice';
   ```

   See Create Directory in Autonomous Database for information on creating directories.
4. Use `DBMS_CLOUD.GET_OBJECT` to upload the target database wallet to the directory you created in the previous step, *WALLET_DIR*.

For example:

```
BEGIN
    DBMS_CLOUD.GET_OBJECT(
        credential_name => 'DEF_CRED_NAME',
        object_uri => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/cwallet.sso',
        directory_name => 'WALLET_DIR');
END;
/
```

In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

> **Note:**
>
> The `credential_name` you use in this step is the credentials for the Object Store. In the next step you create the credentials to access the target database.

5. On Autonomous Database create credentials to access the target database. The `username` and `password` you specify with `DBMS_CLOUD.CREATE_CREDENTIAL` are the credentials for the target database used within the database link, (where the target database is accessed through the VCN).

> **Note:**
>
> Supplying the `credential_name` parameter is required.

For example:

```
BEGIN
   DBMS_CLOUD.CREATE_CREDENTIAL(
      credential_name => 'DB_LINK_CRED',
      username => 'NICK',
      password => 'password');
END;
/
```

The characters in the `username` parameter must be all uppercase letters.

> **Note:**
>
> You can use a vault secret credential for the target database credential in a database link. See Use Vault Secret Credentials for more information.

This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name.

6. Create the database link to the target database using `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

   For example:

```
BEGIN
    DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
        db_link_name => 'PEDBLINK1',
        hostname => 'example1.adb.ap-osaka-1.oraclecloud.com',
        port => '1522',
        service_name => 'example_high.adb.oraclecloud.com',
        ssl_server_cert_dn => 'ssl_server_cert_dn',
        credential_name => 'DB_LINK_CRED',
        directory_name => 'WALLET_DIR',
        private_target => TRUE);
END;
/
```

> **Note:**
>
> If you set `ROUTE_OUTBOUND_CONNECTIONS` to `PRIVATE_ENDPOINT`, setting the `private_target` parameter to `TRUE` is not required in this API. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

For a target on a private endpoint, `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` supports specifying a single hostname with the `hostname` parameter. On a private endpoint, using an IP address, SCAN IP, or a SCAN hostname is not supported (when the target is on a public endpoint, `CREATE_DATABASE_LINK` supports using an IP address, a SCAN IP, or a SCAN hostname).

When the target is an Oracle RAC database, use the `rac_hostnames` parameter to specify one or more hostnames with `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`. This allows you to take advantage of the high availability capabilities of Oracle RAC. Using an IP address, a SCAN IP, or a SCAN hostname in the `rac_hostnames` value is not supported.

For example, with a target Oracle RAC database use the `rac_hostnames` parameter:

```
BEGIN
    DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
        db_link_name => 'PEDBLINK1',
        rac_hostnames => '["sales1-svr1.example.adb.us-
ashburn-1.oraclecloud.com",
                          "sales1-svr2.example.adb.us-
ashburn-1.oraclecloud.com",
                          "sales1-svr3.example.adb.us-
ashburn-1.oraclecloud.com"]',
        port => '1522',
        service_name => 'example_high.adb.oraclecloud.com',
        ssl_server_cert_dn => 'ssl_server_cert_dn',
        credential_name => 'DB_LINK_CRED',
        directory_name => 'WALLET_DIR',
        private_target => TRUE);
END;
/
```

`DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` does not support a value of `localhost` for the `hostname` or in the `rac_hostnames` parameter.

Users other than ADMIN require privileges to run `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

7. Use the database link you created to access data in the target database.

   For example:

```
SELECT * FROM employees@PEDBLINK1;
```

> **Note:**
>
> For the credentials you create in Step 5, the Oracle Database credentials, if the password of the target user changes you can update the credential that contains the target user's credentials as follows:
>
> ```
> BEGIN
>     DBMS_CLOUD.UPDATE_CREDENTIAL (
>         credential_name => 'DB_LINK_CRED',
>         attribute => 'PASSWORD',
>         value => 'password');
> END;
> /
> ```
>
> Where *password* is the new password.
>
> After this operation, the existing database links that use this credential continue to work without having to drop and recreate the database links.

See CREATE_DATABASE_LINK Procedure for additional information.

Database Link Notes with a Target Oracle Database

Provides notes for creating database links to a target Oracle database (when the target is not an Autonomous Database)

Notes for database links to other Oracle databases:

- If you are using database links between Autonomous Database and other Oracle Databases, you might need to apply Patch 33843368 on the Oracle Database that is not an Autonomous Database. This applies to cases where the Autonomous Database instance is either the source or the target of the database link.

  See My Oracle Support Knowledge Base: Patch Requirement For Database Links Between ADB-S And Other Oracle Databases(Doc ID 2874244.1) for further details.

- Only one wallet file is valid per directory for use with database links. You can only upload one `cwallet.sso` at a time to the directory you choose for wallet files (for example *DBLINK_WALLET_DIR*). This means with a `cwallet.sso` in *DBLINK_WALLET_DIR* you can only create database links to the databases for which the wallet in that directory is valid. To use multiple `cwallet.sso` files with database links you need to create additional directories and put each `cwallet.sso` in a different directory. When you create database

links with `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`, specify the directory that contains the wallet with the `directory_name` parameter.

See Create Directory in Autonomous Database for information on creating directories.

- Supported target Oracle database versions for database links to another Oracle Database are: 19c, 12.2.0, and 12.1.0.

> **Note:**
>
> For complete information on supported versions, see Client Server Interoperability Support Matrix for Different Oracle Versions (Doc ID 207303.1)

- Autonomous Database sets the `SEC_CASE_SENSITIVE_LOGON` parameter to `true` and this value cannot be changed. If your target database is not an Autonomous Database, then you must set `SEC_CASE_SENSITIVE_LOGON` parameter to `true` on the target database. If `SEC_CASE_SENSITIVE_LOGON` is set to `false` on the target database, then error `ORA-28040: No matching authentication protocol` is raised.
- To list the database links, use the `ALL_DB_LINKS` view. See ALL_DB_LINKS for more information.
- The wallet file, along with the Database user ID and password provide access to data in the target Oracle database. Store wallet files in a secure location. Share wallet files only with authorized users.
- When the Autonomous Database instance is on a private endpoint, there are two options to specify the target database: use either the `hostname` parameter or the `rac_hostnames` parameter:
  - For a target on a private endpoint, `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` supports specifying a single hostname with the `hostname` parameter. On a private endpoint, using an IP address, SCAN IP, or a SCAN hostname is not supported (when the target is on a public endpoint, `CREATE_DATABASE_LINK` supports using an IP address, a SCAN IP, or a SCAN hostname).
  - When the target is an Oracle RAC database, use the `rac_hostnames` parameter to specify one or more hostnames with `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`. This allows you to take advantage of the high availability capabilities of Oracle RAC. Using an IP address, a SCAN IP, or a SCAN hostname in the `rac_hostnames` value is not supported.

    When you specify a list of host names in the `rac_hostnames` parameter, `CREATE_DATABASE_LINK` uses all of the specified host names as addresses in the connect string. If one of the specified hosts is not available on the target Oracle RAC database, Autonomous Database automatically attempts to connect using another host name from the list.
  - `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` does not support a value of `localhost` for the `hostname` or in the `rac_hostnames` parameter.

Create Database Links to Non-Oracle Databases

You can create database links from Autonomous Database to Non-Oracle databases.

- Create Database Links to Non-Oracle Databases with Oracle-Managed Heterogeneous Connectivity
  Autonomous Database support for Oracle-managed heterogeneous connectivity makes it easy to create database links to non-Oracle databases.
- Create Database Links to Non-Oracle Databases with Customer-Managed Heterogeneous Connectivity
  Use `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` to create database links from an Autonomous Database to an Oracle Database Gateway to access Non-Oracle databases.

Create Database Links to Non-Oracle Databases with Oracle-Managed Heterogeneous Connectivity

Autonomous Database support for Oracle-managed heterogeneous connectivity makes it easy to create database links to non-Oracle databases.

When you use database links with Oracle-managed heterogeneous connectivity, Autonomous Database configures and sets up the connection to the non-Oracle database.

> **Note:**
>
> With Oracle-managed heterogeneous connectivity the support is for query-only connections on the remote database. That is, updates are not supported when using Oracle-managed heterogeneous connectivity.

The following is a prerequisite to use Oracle-managed heterogeneous connectivity with Autonomous Database:

- When the target database is on a public endpoint, database must be configured to allow incoming SSL/TLS connections with CA signed certificates.

  Oracle-managed heterogeneous connectivity supports connections to target database services on private endpoints (for example you can connect to Oracle MySQL Database Service when the service is on a private endpoint). When you connect to a non-oracle database with Oracle-managed heterogeneous connectivity on a private endpoint, the connection uses TCP protocol and it does not require SSL/TLS to be configured on the target database.

  See Create Database Links to Oracle MySQL on a Private Endpoint with Oracle-Managed Heterogeneous Connectivity for more information.

To create database links to a non-Oracle database using Oracle-managed heterogeneous connectivity, do the following:

1. On Autonomous Database create credentials to access the target database. The `username` and `password` you specify with `DBMS_CLOUD.CREATE_CREDENTIAL` are the credentials for the target database used within the database link.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'AWS_REDSHIFT_LINK_CRED',
       username => 'nick',
       password => 'password'
     );
   ```

```
END;
/
```

This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Supplying the `credential_name` parameter is required.

To access Google Analytics, Google BigQuery, Google Drive, Youtube, ServiceNow, or Microsoft SharePoint with OAuth2, the credential must include the `params` parameter with the value `gcp_oauth2`.

For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'SERVICENOW_OAUTH',
    params => JSON_OBJECT(
                'gcp_oauth2' value JSON_OBJECT(
                       'client_id' value 'CLIENT_ID',
                       'client_secret' value 'CLIENT_SECRET',
                       'refresh_token' value 'Refresh_Token')));
END;
/
```

See CREATE_CREDENTIAL Procedure for more information.

2. Create the database link to the target database using `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

   For example, to create a database link to AWS Redshift:

```
BEGIN
     DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
          db_link_name => 'AWSREDSHIFT_LINK',
          hostname => 'example.com',
          port => '5439',
          service_name => 'example_service_name',
          credential_name => 'AWS_REDSHIFT_LINK_CRED',
          gateway_params => JSON_OBJECT('db_type'  value 'awsredshift'),
          ssl_server_cert_dn => NULL);
END;
/
```

The `service_name` is the database name of the non-Oracle database.

The `gateway_params db_type` value that you supply must be one of the supported values:

db_type Value	Database Type
awsredshift	Amazon Redshift
azure	Microsoft SQL Server Azure SQL Azure Synapse Analytics
db2	IBM Db2
google_analytics	Google Analytics

db_type Value	Database Type
google_bigquery	Google BigQuery
google_drive	Google Drive
hive	Apache Hive
mongodb	MongoDB
mysql	MySQL
mysql_community	MySQL Community Edition
postgres	PostgreSQL
salesforce	Salesforce
servicenow	ServiceNow
sharepoint	Microsoft SharePoint
snowflake	Snowflake
youtube	YouTube

Autonomous Database automatically configures and handles the secure connection to a target database and your connections are end-to-end encrypted. Oracle-managed heterogeneous connectivity is preconfigured with a wallet that contains most of the common trusted root and intermediate SSL certificates. Thus, `NULL` must be provided as the value for the `ssl_server_cert_dn` parameter.

To ensure security when using database links with Oracle-managed heterogeneous connectivity, the connection port is restricted and must have SSL/TLS enabled. You specify the target database port with the `port` parameter.

See Oracle-Managed Heterogeneous Connectivity Database Types and Ports for the list of supported non-Oracle database types.

The `HETEROGENEOUS_CONNECTIVITY_INFO` view provides information on supported Oracle Heterogeneous Connectivity types and shows a PL/SQL code sample for each supported type. See Access Heterogeneous Connectivity Information and Samples for more information.

When you specify the `gateway_params` parameter, for some `db_type` values, additional `gateway_params` parameters are supported:

db_type Value	Supported Gateway Parameters with Specified db_type
azure	When the `db_type` value is `azure` there are two optional parameters, `auth_method` and `domain` to support NTLM/AD authentication. When NTLMv2 is configured, set `auth_method=10` and set `domain` to the Windows domain value. To use an Azure Active Directory password, set `auth_method=13`. Do not use `auth_method=13` to access an Oracle on-premises Database. See Microsoft SQL Server Authentication Method for more information.
google_analytics	When the `db_type` value is `google_analytics` the `hostname` is not used and you can provide value such as `example.com`. For `db_type google_analytics`, the credential must be an OAuth type credential using the `params` parameter with `gcp_oauth2` values specified (`client_id`, `client_secret`, and `refresh_token`). See CREATE_CREDENTIAL Procedure for more information.

db_type Value	Supported Gateway Parameters with Specified db_type
google_bigquery	When db_type is google_bigquery, the parameter project is valid. This parameter specifies the project name on google_bigquery and is required. When the db_type value is google_bigquery the hostname is not used and you can provide value such as example.com. For db_type google_bigquery, the credential must be an OAuth type credential using the params parameter with gcp_oauth2 values specified (client_id, client_secret, and refresh_token). See CREATE_CREDENTIAL Procedure for more information.

google_drive	For db_type google_drive, the credential you supply with DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK must be an OAuth type credential using the params parameter with gcp_oauth2 values specified (client_id, client_secret, and refresh_token).
	To connect to Google Drive and get data you must supply the gateway parameters directory_name and file_name. These parameters specify a model file (auto REST file) that maps the JSON response to the relational model. The model file specifies the endpoints, table mapping, and HTTP response code for processing the JSON response. See Model file syntax and Example Model file for more information.
	When you use gateway_params parameter with db_type google_drive you must supply the gateway_params parameter db_type with the value 'google_drive' and the directory_name and file_name parameters, along with OAuth type credentials:
	The directory_name parameter specifies the directory with the REST config file. You could create this directory as follows:
	`create or replace directory google_drive_dir as 'GOOGLE_DRIVE_DIR';`
	Obtain and download the auto REST file to the specified directory. For example:
	`exec DBMS_CLOUD.get_object('google_drive_dir_cred', 'https://objectstorage.<...>/ google_drive.rest','GOOGLE_DRIVE_DIR');`
	Set the file_name value to the name of the auto REST file you downloaded, "google_drive.rest". Then you can use the REST config file with OAuth 2.0.

hive	When db_type is hive, the parameter http_path is valid. This parameter specifies the HttpPath value, if required, to connect to the Hive instance.

db_type Value	Supported Gateway Parameters with Specified db_type
salesforce	When you use gateway_params parameter with db_type salesforce, you must supply the security_token option. The security token is a case-sensitive alphanumeric code. See Reset Your Security Token for more information. When you use gateway_params parameter with db_type salesforce, you must supply the correct hostname parameter. Salesforce provides two forms of URLs for the Salesforce service account: • *xxxx*.develop.lightning.force.com form of URL • *xxxx*my.salesforce.com form of URL as shown under the **View profile** tab. Oracle-Managed Heterogeneous Connectivity only supports the *xxxx*my.salesforce.com form of URL for the hostname parameter value.

servicenow	To connect to ServiceNow and get data you must supply the gateway parameters directory_name and file_name. These parameters specify a model file (REST config file) that maps the JSON response to the relational model. The model file specifies the endpoints, table mapping, and HTTP response code for processing the JSON response. See Model file syntax and Example Model file for more information.
	When you use gateway_params parameter with db_type servicenow there are two supported options:
	• Basic Authentication: you must supply the gateway_params parameter db_type with the value 'servicenow', and supply the directory_name and file_name parameters along with username/password type credentials.
	• OAuth 2.0 Authentication: you must supply the gateway_params parameter db_type with the value 'servicenow', and the directory_name, file_name, and token_uri parameters, along with OAuth type credentials.
	For OAuth 2.0 authentication with db_type servicenow, the credential must be an OAuth type credential using the params parameter with gcp_oauth2 values specified specified (client_id, client_secret, and refresh_token). See CREATE_CREDENTIAL Procedure for more information.
	The directory_name parameter specifies the directory with the ServiceNow REST config file. You could create this directory as follows:
	`create or replace directory servicenow_dir as 'SERVICENOW_DIR';`
	Obtain and download the ServiceNow REST config file to the specified directory. For example:
	`exec DBMS_CLOUD.get_object('servicenow_dir_cred', 'https://objectstorage.<...>/ servicenow.rest','SERVICENOW_DIR');`
	Set the file_name value to the name of the REST config file you downloaded, "servicenow.rest".
	Then you can use the ServiceNow REST config file with either basic authentication or OAuth2.0. See HETEROGENEOUS_CONNECTIVITY_INFO View for samples.

db_type Value	Supported Gateway Parameters with Specified db_type
sharepoint	When you use gateway_params parameter with db_type sharepoint, also specify values for auth_uri, scope, service_url, and token_uri. For db_type sharepoint, the credential you supply with DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK must be an OAuth type credential using the params parameter with gcp_oauth2 values specified (client_id, client_secret, and refresh_token). See CREATE_CREDENTIAL Procedure for more information. Generate the refresh_token for *tenant_id*.sharepoint.com/.default offline_access. See the following for more information: Determine the scope and OAuth 2.0 authentication. Set values for gateway_params for db_type sharepoint: • auth_uri: Set the auth_uri from Azure. See Obtain application client information and endpoints for more information. • scope: Set the scope. See Connection option descriptions for more information. • service_url: Set the service_url. See Service URL for more information. • token_uri: Get the token_uri from Azure. See Obtain application client information and endpoints for more information.
snowflake	When you use gateway_params parameter with db_type snowflake, use the Snowflake account identifier as the hostname parameter. In this case, the driver adds snowflakecomupting.com, so you do not pass this part of the hostname explicitly. To find your Snowflake account identifier, see Account Identifier Formats by Cloud Platform and Region. For example: for the Snowflake account: https://example-marketing_test_account.snowflakecomputing.com Set the hostname value to "example-marketing_test_account". When the db_type is SNOWFLAKE, the optional parameters: role, schema, and warehouse are valid. These values specify a different schema, role, or warehouse value, other than the default. For example: `gateway_params => JSON_OBJECT( 'db_type' value 'snowflake', 'role' value 'ADMIN', 'schema' value 'PUBLIC', 'warehouse' value 'TEST' )`

db_type Value	Supported Gateway Parameters with Specified db_type

youtube	To connect to YouTube and get data, you must supply the parameters directory_name and file_name. These parameters specify a model file (REST config file) that maps the JSON response to the relational model. The model file specifies the endpoints, table mapping, and HTTP response code for processing the JSON response. See Model file syntax and Example Model file for more information.
	For db_type youtube, the credential must be an OAuth type credential using the params parameter with gcp_oauth2 values specified (client_id, client_secret, and refresh_token). See CREATE_CREDENTIAL Procedure for more information.
	The directory_name parameter specifies the directory with the youtube REST config file.
	For example, to create this directory and add the required contents:
	a. Create a directory on the Autonomous Database instance:
	`create or replace directory youtube_dir as 'YOUTUBE_DIR';`
	b. Obtain and download the auto rest REST config file to the specified directory. For example:
	`exec DBMS_CLOUD.get_object('youtube_dir_cred', 'https://objectstorage.<...>/youtube.rest','YOUTUBE_DIR');`
	c. Set the file_name value to the name of the REST config file you downloaded, 'youtube.rest'.
	Then you can use the REST config file with OAuth2.0. See HETEROGENEOUS_CONNECTIVITY_INFO View for samples.

See CREATE_DATABASE_LINK Procedure for more information.

3. Use the database link to access data on the target database.

   For example:

```
SELECT count(*) FROM sales@AWSREDSHIFT_LINK
```

   The table name you specify when you use SELECT with Google BigQuery must be in quotes. For example:

```
SELECT count(*) FROM "sales"@GOOGLE_BIGQUERY_LINK
```

> **Note:**
>
> For the credentials you create in Step 1, the target database credentials, if the password of the target user changes you can update the credential that contains the target user's credentials as follows:
>
> ```
> BEGIN
>     DBMS_CLOUD.UPDATE_CREDENTIAL (
>        credential_name => 'AWS_REDSHIFT_LINK_CRED',
>        attribute => 'PASSWORD',
>        value => 'password');
> END;
> /
> ```
>
> Where *password* is the new password.
>
> After this operation, the existing database links that use this credential continue to work without having to drop and recreate the database links.

For additional information, see:

- Supported SQL Syntax and Functions in *Oracle Database Gateway for ODBC User's Guide*
- CREATE_DATABASE_LINK Procedure
- UPDATE_CREDENTIAL Procedure
- Create Database Links to Oracle MySQL on a Private Endpoint with Oracle-Managed Heterogeneous Connectivity
  Autonomous Database support for Oracle-managed heterogeneous connectivity makes it easy to create database links to Oracle MySQL Database Service on a private endpoint. When you use database links with Oracle-managed heterogeneous connectivity, Autonomous Database configures and sets up the connection to the Oracle MySQL Database Service.
- Oracle-Managed Heterogeneous Connectivity Database Types and Ports
  Shows the non-Oracle databases that you can connect to from Autonomous Database with Oracle-managed heterogeneous connectivity, and lists the supported port value for each database type. Also provides a link where you can see the supported database versions for each database type.
- Access Heterogeneous Connectivity Information and Samples
  Oracle Autonomous Database provides heterogeneous connectivity information for database links to non-Oracle databases.

Create Database Links to Oracle MySQL on a Private Endpoint with Oracle-Managed Heterogeneous Connectivity

Autonomous Database support for Oracle-managed heterogeneous connectivity makes it easy to create database links to Oracle MySQL Database Service on a private endpoint. When you use database links with Oracle-managed heterogeneous connectivity, Autonomous Database configures and sets up the connection to the Oracle MySQL Database Service.

The following are prerequisites to use Oracle-managed heterogeneous connectivity with Oracle MySQL Database Service on a private endpoint:

- Create a DNS name using private DNS Zone pointing to private IP of your Oracle MySQL Database Service in your VCN. See Private DNS.
- Create an Autonomous Database with a Private Endpoint on same subnet.
- Configure the VCN for the Oracle MySQL Database Service to allow incoming connections on port 3306.

To create database links to a Oracle MySQL Database Service on a private endpoint using Oracle-managed heterogeneous connectivity, do the following:

1. On Autonomous Database create credentials to access the Oracle MySQL Database Service. The `username` and `password` you specify with `DBMS_CLOUD.CREATE_CREDENTIAL` are the credentials for the Oracle MySQL Database Service used within the database link.

   > **Note:**
   >
   > Supplying the `credential_name` parameter is required.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'MYSQL_LINK_CRED',
       username => 'NICK',
       password => 'password'
     );
   END;
   /
   ```

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name.

2. Create the database link to the Oracle MySQL Database Service using `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

   For example, to create a database link:

   ```
   BEGIN
        DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
             db_link_name => 'MYSQL_LINK',
             hostname => 'mysql.example.com',
             port => '3306',
             service_name => 'mysql.example_service_name',
             ssl_server_cert_dn => NULL,
             credential_name => 'MYSQL_LINK_CRED',
             private_target => TRUE,
             gateway_params => JSON_OBJECT('db_type' value 'mysql'));
   END;
   /
   ```

3. Use the database link to access data on the target database.

   For example:

   ```
   SELECT count(*) FROM sales@MYSQL_LINK
   ```

> **Note:**
>
> For the credentials you create in Step 1, the target database credentials, if the password of the target user changes you can update the credential that contains the target user's credentials as follows:
>
> ```
> BEGIN
>     DBMS_CLOUD.UPDATE_CREDENTIAL (
>        credential_name => 'MYSQL_LINK_CRED',
>        attribute => 'PASSWORD',
>        value => 'password');
> END;
> /
> ```
>
> Where *password* is the new password.
>
> After this operation, the existing database links that use this credential continue to work without having to drop and recreate the database links.

For additional information, see:

- Supported SQL Syntax and Functions in *Oracle Database Gateway for ODBC User's Guide*
- CREATE_DATABASE_LINK Procedure
- UPDATE_CREDENTIAL Procedure

Oracle-Managed Heterogeneous Connectivity Database Types and Ports

Shows the non-Oracle databases that you can connect to from Autonomous Database with Oracle-managed heterogeneous connectivity, and lists the supported port value for each database type. Also provides a link where you can see the supported database versions for each database type.

> **Note:**
>
> Oracle uses Progress DataDirect connectors. The Database Support column provides links to the Progress website where you can find the supported database versions for each database type.

Database Type	db_type Value	Required Port	Database Support
Amazon Redshift	`awsredshift`	5439	Amazon Redshift Supported Versions
Azure SQL Microsoft SQL Server Azure Synapse Analytics	`azure`	1433	Azure SQL Supported Versions Azure Synapse Analytics Supported Versions
Apache Hive	`hive`	443	Hive Supported Versions
Google Analytics	`google_analytics`	443	Google Analytics Supported Versions

Database Type	db_type Value	Required Port	Database Support
Google BigQuery	`google_bigquery`	443	Google BigQuery Supported Versions
Google Drive	`google_drive`	443	Google Drive Supported Versions
Apache Hive	`hive`	443	Hive Supported Versions
IBM Db2 11.5.6 or greater IBM Db2 11.5.5 or less	`db2`	25000 50000	IBM Db2 Supported Versions
Microsoft SharePoint	`sharepoint`	443	Microsoft SharePoint Supported Versions
MongoDB	`mongodb`	27017	MongoDB Supported Versions
MySQL	`mysql`	3306	MySQL Supported Versions
MySQL Community Edition	`mysql_community`	3306	
PostgreSQL	`postgres`	5432	PostgreSQL Supported Versions
Salesforce	`salesforce`	19937	Salesforce Supported Versions
ServiceNow	`servicenow`	443	ServiceNow Supported Versions
Snowflake	`snowflake`	443	Snowflake Supported Versions
Youtube	`youtube`	443	Youtube Supported Versions

Access Heterogeneous Connectivity Information and Samples

Oracle Autonomous Database provides heterogeneous connectivity information for database links to non-Oracle databases.

The `HETEROGENEOUS_CONNECTIVITY_INFO` view provides information on supported Oracle Heterogeneous Connectivity types and shows a PL/SQL code sample for each supported type.

For example:

```
SELECT * FROM HETEROGENEOUS_CONNECTIVITY_INFO WHERE DATABASE_TYPE = 'hive';
```

- HETEROGENEOUS_CONNECTIVITY_INFO View
  The `HETEROGENEOUS_CONNECTIVITY_INFO` view lists connectivity information and examples for connecting with PL/SQL using database links and Oracle Managed Heterogeneous Connectivity.

HETEROGENEOUS_CONNECTIVITY_INFO View

The `HETEROGENEOUS_CONNECTIVITY_INFO` view lists connectivity information and examples for connecting with PL/SQL using database links and Oracle Managed Heterogeneous Connectivity.

Column	Datatype	Description
`DATABASE_TYPE`	`VARCHAR2(32)`	Database type value used with `gateway_params` parameter.

Column	Datatype	Description
REQUIRED_PORT	NUMBER	Supported port values for the database type.
DESCRIPTION	CLOB	Specifies a description for the DATABASE_TYPE.
OPTIONAL_PARAMETERS	VARCHAR2(1024)	Specifies the valid optional parameters for the DATABASE_TYPE.
SAMPLE_USAGE	CLOB	Shows sample PL/SQL usage for the DATABASE_TYPE.

Create Database Links to Non-Oracle Databases with Customer-Managed Heterogeneous Connectivity

Use DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK to create database links from an Autonomous Database to an Oracle Database Gateway to access Non-Oracle databases.

Topics

- Prerequisites to Create Database Links to Non-Oracle Databases with Customer-Managed Heterogeneous Connectivity
  Use DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK to create database links from an Autonomous Database to an Oracle Database Gateway to access Non-Oracle databases.
- Create Database Links with Customer-Managed Heterogeneous Connectivity to Publicly Accessible Non-Oracle Databases
  Use DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK to create database links from an Autonomous Database instance on a public endpoint to an Oracle Database Gateway to access Non-Oracle databases.
- Create Database Links with Customer-Managed Heterogeneous Connectivity to Non-Oracle Databases on a Private Endpoint
  You can create database links from an Autonomous Database to an Oracle Database Gateway to access Non-Oracle databases that are on a private endpoint.

Prerequisites to Create Database Links to Non-Oracle Databases with Customer-Managed Heterogeneous Connectivity

Use DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK to create database links from an Autonomous Database to an Oracle Database Gateway to access Non-Oracle databases.

An Oracle Database Gateway is a gateway that is designed for accessing a specific non-Oracle system. Using an Oracle Database Gateway, you can access data anywhere in a distributed database system without knowing either the location of the data or how it is stored. Using database links on Autonomous Database with Oracle Database Gateway supports heterogeneous environments and eliminates the need to customize your applications to access data from non-Oracle systems.

Before you create database links to a target gateway, do the following:

1. Configure the Oracle Database Gateway to access a non-Oracle database.
2. Configure Oracle Net Listener to handle incoming requests on the Oracle Database Gateway.
3. Create a self signed wallet on the Oracle Database Gateway.

For more information on Oracle Database Gateway, see Oracle Database Gateways in *Oracle Database Heterogeneous Connectivity User's Guide*. Also see the Installation and Configuration Guide and the Gateway User's Guide for the database you want to connect to. For example, for Oracle Database Gateway for SQL Server see:

- Installing and Configuring Oracle Database Gateway for SQL Server
- Introduction to the Oracle Database Gateway for SQL Server in *Oracle Database Gateway for SQL Server User's Guide*.
- Configure Oracle Net for the Gateway

Create Database Links with Customer-Managed Heterogeneous Connectivity to Publicly Accessible Non-Oracle Databases

Use `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` to create database links from an Autonomous Database instance on a public endpoint to an Oracle Database Gateway to access Non-Oracle databases.

To use database links from an Autonomous Database instance on a public endpoint, the target an Oracle Database Gateway must be configured to use TCP/IP with SSL (TCPS) authentication. See Configuring Secure Sockets Layer Authentication for more information.

To create database links from an Autonomous Database instance on a public endpoint to a target Oracle Database Gateway:

1. Copy the target Oracle Database Gateway self-signed wallet to Object Store. For example, copy `cwallet.sso`, containing the certificates for the Oracle Database Gateway, to Object Store.

   > **Note:**
   >
   > The wallet file, along with the Database user ID and password provide access to data available through the target Oracle Database Gateway. Store wallet files in a secure location. Share wallet files only with authorized users.

2. Create credentials to access the Object Store where you store the `cwallet.sso`. See CREATE_CREDENTIAL Procedure for information about the username and password parameters for different object storage services.
3. Create a directory on your Autonomous Database instance for the wallet file `cwallet.sso`.

   For example:

   ```
   CREATE DIRECTORY dblink_wallet_dir AS 'directory_path_of_your_choice';
   ```

   See Create Directory in Autonomous Database for information on creating directories.
4. Use `DBMS_CLOUD.GET_OBJECT` to upload the target gateway self signed wallet to the directory you created in the previous step, *DBLINK_WALLET_DIR*.

   For example:

   ```
   BEGIN
        DBMS_CLOUD.GET_OBJECT(
           credential_name => 'DEF_CRED_NAME',
           object_uri => 'https://objectstorage.us-
   phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/cwallet.sso',
   ```

```
        directory_name => 'DBLINK_WALLET_DIR');
END;
/
```

In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

The `credential_name` you use in this step is the credentials for the Object Store. In the next step you create the credentials to access the target gateway.

5. On your Autonomous Database instance create a credential to access the target database. The `username` and `password` you specify with `DBMS_CLOUD.CREATE_CREDENTIAL` are the credentials for the target database to be used when you create the database link, (where the target database is accessed through the Oracle Database Gateway).

   Supplying the `credential_name` parameter is required.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'DB_LINK_CRED',
       username => 'NICK',
       password => 'password'
     );
   END;
   /
   ```

   The characters in the `username` parameter must be all uppercase letters.

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name.

6. Create a database link to access the target gateway using `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

   For example:

   ```
   BEGIN
        DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
             db_link_name =>       'SALESLINK',
             hostname =>           'example.com',
             port =>               '1522',
             service_name =>       'example_service_name',
             ssl_server_cert_dn => 'ssl_server_cert_dn',
             credential_name =>    'DB_LINK_CRED',
             directory_name =>     'DBLINK_WALLET_DIR',
             gateway_link =>        TRUE);
   END;
   /
   ```

   Users other than ADMIN require privileges to run `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

7. Use the database link to access data on the target gateway.

For example:

```
SELECT * FROM employees@SALESLINK;
```

For the credentials you create in Step 5, the Oracle Database Gateway credentials, if the password of the target user changes you can update the credential that contains the target user's credentials as follows:

```
BEGIN
    DBMS_CLOUD.UPDATE_CREDENTIAL (
       credential_name => 'DB_LINK_CRED',
       attribute => 'PASSWORD',
       value => 'password');
END;
/
```

Where *password* is the new password.

After this operation, the existing database links that use this credential continue to work without having to drop and recreate the database links.

See the following for an example that shows you how to create a database link to an Oracle Database Gateway to access a Microsoft SQL Server database:

How to Access Non-Oracle Databases from Autonomous Database using Oracle Database Gateway

For additional information, see:

- CREATE_DATABASE_LINK Procedure
- GET_OBJECT Procedure and Function
- UPDATE_CREDENTIAL Procedure

Create Database Links with Customer-Managed Heterogeneous Connectivity to Non-Oracle Databases on a Private Endpoint

You can create database links from an Autonomous Database to an Oracle Database Gateway to access Non-Oracle databases that are on a private endpoint.

This section covers the steps for using database links to connect from Autonomous Database to a non-Oracle Database that through an Oracle Database Gateway, where the non-Oracle Database is on a private endpoint.

- Prerequisites to Create Database Links with Customer-Managed Heterogeneous Connectivity to Non-Oracle Databases on a Private Endpoint
  Lists the prerequisites to create database links from an Autonomous Database with Customer-Managed Heterogeneous Connectivity to Non-Oracle Databases that are on a Private Endpoint.
- Create Database Links with Customer-Managed Heterogeneous Connectivity to Non-Oracle Databases on a Private Endpoint (without a wallet)
  You can create database links from an Autonomous Database instance to an Oracle Database Gateway to access Non-Oracle databases that are on a private endpoint, connecting either with a wallet (TCPS) or without a wallet (TLS).

- Create Database Links with Customer-Managed Heterogeneous Connectivity to Non-Oracle Databases on a Private Endpoint (with a Wallet)
  You can create database links from an Autonomous Database instance to an Oracle Database Gateway to access Non-Oracle databases that are on a private endpoint, connecting either with a wallet (TCPS) or without a wallet (TLS).

Prerequisites to Create Database Links with Customer-Managed Heterogeneous Connectivity to Non-Oracle Databases on a Private Endpoint

Lists the prerequisites to create database links from an Autonomous Database with Customer-Managed Heterogeneous Connectivity to Non-Oracle Databases that are on a Private Endpoint.

To create a database link with Customer-Managed Heterogeneous Connectivity to Non-Oracle Databases that are on a Private Endpoint:

- The target database must be accessible from the source database's Oracle Cloud Infrastructure VCN. For example, you can connect to the target database when:
  - The target database is on a private endpoint.
  - Both the source database and the target database are in the same Oracle Cloud Infrastructure VCN.
  - The source database and the target database are in different Oracle Cloud Infrastructure VCNs that are paired.
  - For a target on a private endpoint, `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` supports specifying a single hostname with the `hostname` parameter. On a private endpoint, using an IP address, SCAN IP, or a SCAN hostname is not supported (when the target is on a public endpoint, `CREATE_DATABASE_LINK` supports using an IP address, a SCAN IP, or a SCAN hostname).
- The following ingress and egress rules must be defined for the private endpoint:
  - Define an egress rule in the source database's subnet security list or network security group such that the traffic over TCP is allowed to the target database's IP address and port number.
  - Define an ingress rule in the target database's subnet security list or network security group such that the traffic over TCP is allowed from the source database IP address to the destination port.

  See Configure Network Access with Private Endpoints for information on configuring private endpoints with ingress and egress rules.

> **Note:**
>
> When your Autonomous Database instance is configured with a private endpoint, set the `ROUTE_OUTBOUND_CONNECTIONS` database parameter to '`PRIVATE_ENDPOINT`' to specify that all outgoing database links are subject to the Autonomous Database instance private endpoint VCN's egress rules. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

Create Database Links with Customer-Managed Heterogeneous Connectivity to Non-Oracle Databases on a Private Endpoint (without a wallet)

You can create database links from an Autonomous Database instance to an Oracle Database Gateway to access Non-Oracle databases that are on a private endpoint, connecting either with a wallet (TCPS) or without a wallet (TLS).

To create database links from an Autonomous Database instance on a private endpoint to a target Oracle Database Gateway, connecting without a wallet (TLS):

1. On Autonomous Database create credentials to access the target database. The `username` and `password` you specify with `DBMS_CLOUD.CREATE_CREDENTIAL` are the credentials for the target database used within the database link, (where the target database is accessed through the Oracle Database Gateway).

   Supplying the `credential_name` parameter is required.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'DB_LINK_CRED',
       username => 'NICK',
       password => 'password'
     );
   END;
   /
   ```

   The characters in the `username` parameter must be all uppercase letters.

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name.

2. Create the database link to the target gateway using `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

   For example:

   ```
   BEGIN
        DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
             db_link_name =>       'SALESLINK',
             hostname =>           'example.com',
             port =>               '1522',
             service_name =>       'example_service_name',
             ssl_server_cert_dn => 'ssl_server_cert_dn',
             credential_name =>    'DB_LINK_CRED',
             directory_name =>      NULL,
             gateway_link =>        TRUE,
             private_target =>      TRUE,
             gateway_params =>      NULL);
   END;
   /
   ```

   The `ssl_server_cert_dn` parameter is optional if the connection is created as a TCP based database link (without a wallet).

   When the `directory_name` is `NULL`, the connection is created as a TCP based database link (without a wallet).

   The `private_target` parameter must be set to `TRUE` when the target non-Oracle Database is on a private endpoint (that is, the database link accesses a hostname that needs to be

resolved in a VCN DNS server). When `private_target` is `TRUE`, the `hostname` parameter must be a single hostname (on a private endpoint, using an IP address, a SCAN IP, or a SCAN hostname is not supported).

When `gateway_link` is `TRUE` and `gateway_params` is `NULL`, this specifies that the database link is to a customer-managed Oracle gateway.

Users other than ADMIN require privileges to run `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

3. Use the database link you created to access data on the target gateway.

   For example:

   ```
   SELECT * FROM employees@SALESLINK;
   ```

For the credentials you create in Step 1, the Oracle Database Gateway credentials, if the password of the target user changes you can update the credential that contains the target user's credentials as follows:

```
BEGIN
    DBMS_CLOUD.UPDATE_CREDENTIAL (
       credential_name => 'DB_LINK_CRED',
       attribute => 'PASSWORD',
       value => 'password');
END;
/
```

Where *password* is the new password.

After this operation, the existing database links that use this credential continue to work without having to drop and recreate the database links.

For additional information, see:

- CREATE_DATABASE_LINK Procedure
- UPDATE_CREDENTIAL Procedure

Create Database Links with Customer-Managed Heterogeneous Connectivity to Non-Oracle Databases on a Private Endpoint (with a Wallet)

You can create database links from an Autonomous Database instance to an Oracle Database Gateway to access Non-Oracle databases that are on a private endpoint, connecting either with a wallet (TCPS) or without a wallet (TLS).

To create database links from an Autonomous Database instance on a private endpoint to a target Oracle Database Gateway, and connect without a wallet:

1. Copy the target Oracle Database Gateway self-signed wallet to Object Store. For example, copy `cwallet.sso`, containing the certificates for the Oracle Database Gateway, to Object Store.

   The wallet file, along with the Database user ID and password provide access to data available through the target gateway. Store wallet files in a secure location. Share wallet files only with authorized users.

2. Create credentials to access the Object Store where you store the `cwallet.sso`. See CREATE_CREDENTIAL Procedure for information about the username and password parameters for different object storage services.

3. Create a directory on Autonomous Database for the wallet file `cwallet.sso`.

For example:

```
CREATE DIRECTORY dblink_wallet_dir AS 'directory_path_of_your_choice';
```

See Create Directory in Autonomous Database for information on creating directories.

4. Use `DBMS_CLOUD.GET_OBJECT` to upload the target gateway self signed wallet to the directory you created in the previous step, *DBLINK_WALLET_DIR*.

   For example:

```
BEGIN
     DBMS_CLOUD.GET_OBJECT(
        credential_name => 'DEF_CRED_NAME',
        object_uri => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/cwallet.sso',
        directory_name => 'DBLINK_WALLET_DIR');
END;
/
```

   In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

> **Note:**
>
> The `credential_name` you use in this step is the credentials for the Object Store. In the next step you create the credentials to access the target gateway.

5. On Autonomous Database create credentials to access the target database. The `username` and `password` you specify with `DBMS_CLOUD.CREATE_CREDENTIAL` are the credentials for the target database used within the database link, (where the target database is accessed through the Oracle Database Gateway).

   Supplying the `credential_name` parameter is required.

   For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DB_LINK_CRED',
    username => 'NICK',
    password => 'password'
  );
END;
/
```

   The characters in the `username` parameter must be all uppercase letters.

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name.

6. Create the database link to the target gateway using `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

For example:

```
BEGIN
     DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
          db_link_name =>       'SALESLINK',
          hostname =>           'example.com',
          port =>               '1522',
          service_name =>       'example_service_name',
          ssl_server_cert_dn => 'ssl_server_cert_dn',
          credential_name =>    'DB_LINK_CRED',
          directory_name =>     'DBLINK_WALLET_DIR',
          gateway_link =>        TRUE,
          private_target =>      TRUE,
          gateway_params =>      NULL);
END;
/
```

If `directory_name` is not `NULL`, a TCPS-based database link is created.

The `private_target` parameter must be set to `TRUE` when the target non-Oracle Database is on a private endpoint (that is, the database link accesses a hostname that needs to be resolved in a VCN DNS server). When `private_target` is `TRUE`, the `hostname` parameter must be a single hostname (on a private endpoint, using an IP address, a SCAN IP, or a SCAN hostname is not supported).

When `gateway_link` is `TRUE` and `gateway_params` is `NULL`, this specifies that the database link is to a customer-managed Oracle gateway.

Users other than ADMIN require privileges to run `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK`.

7. Use the database link you created to access data on the target gateway.

   For example:

   ```
   SELECT * FROM employees@SALESLINK;
   ```

For the credentials you create in Step 5, the Oracle Database Gateway credentials, if the password of the target user changes you can update the credential that contains the target user's credentials as follows:

```
BEGIN
    DBMS_CLOUD.UPDATE_CREDENTIAL (
       credential_name => 'DB_LINK_CRED',
       attribute => 'PASSWORD',
       value => 'password');
END;
/
```

Where *password* is the new password.

After this operation, the existing database links that use this credential continue to work without having to drop and recreate the database links.

For additional information, see:

- CREATE_DATABASE_LINK Procedure

- GET_OBJECT Procedure and Function
- UPDATE_CREDENTIAL Procedure

Create Database Links from Other Databases to Autonomous Database

You can create database links to an Autonomous Database from an Oracle database that is not an Autonomous Database. For example from an on-premise Oracle database to an Autonomous Database

- Supported Oracle Database versions for database links where the source is an Oracle Database and the target is an Autonomous Database instance are: 19.2 (or later), 21 (base release or later).
- If you are using database links between Autonomous Database and other Oracle Databases, you might need to apply Patch 33843368 on the Oracle Database that is not an Autonomous Database. This applies to cases where the Autonomous Database instance is either the source or the target of the database link.

  See My Oracle Support Knowledge Base: Patch Requirement For Database Links Between ADB-S And Other Oracle Databases(Doc ID 2874244.1) for further details.

To create database links to an Autonomous Database do the following:

1. Download your Autonomous Database wallet. See Download Client Credentials (Wallets) for more information.
2. Upload the wallet to the database instance where you want to create the link to the Autonomous Database.
3. Unzip the Autonomous Database wallet:

> **Note:**
>
> The wallet file, along with the Database user ID and password provide access to data in your Autonomous Database. Store wallet files in a secure location. Share wallet files only with authorized users.

```
[oracle@sys1 ~]$ cd/u01/targetwallet
[oracle@sys1 targetwallet]$ unzip Wallet_name1.zip
        Archive: Wallet_name1.zip
          inflating: cwallet.sso
          inflating: tnsnames.ora
          inflating: truststore.jks
          inflating: ojdbc.properties
          inflating: sqlnet.ora
          inflating: ewallet.p12
          inflating: keystore.jks
```

4. Set `GLOBAL_NAMES` to `FALSE`.

```
SQL> ALTER SYSTEM SET GLOBAL_NAMES = FALSE;

System altered.

SQL> SHOW PARAMETER GLOBAL_NAMES
NAME                      TYPE        VALUE
```

```
----------------------   ----------- -----------
global_names             boolean     FALSE
```

Set `GLOBAL_NAMES` to `FALSE` to use a database link name without checking that the name is different than the remote database name.When `GLOBAL_NAMES`, is set to `TRUE`, the database requires the database link to have the same name as the database to which it connects. See GLOBAL_NAMES for more information.

5. Create the database link to the target Autonomous Database. Note that the `security` path includes `my_wallet_directory`; the path where you unzip the Autonomous Database wallet.

```
CREATE DATABASE LINK ADBLINK
   CONNECT TO NAME1 IDENTIFIED BY ************
   USING
'(description=(retry_count=20)(retry_delay=3)
     (address=(protocol=tcps)(port=1522)
(host=example1.oraclecloud.com))
     (connect_data=(service_name=example2_high.adb.oraclecloud.com))
     (security=(my_wallet_directory=/u01/targetwallet)
        (ssl_server_dn_match=true)))';

Database link created.
```

6. Use the database link you created to access data on the target database (your Autonomous Database instance in this case):

   For example:

```
SELECT * FROM employees@ADBLINK;
```

To list the database links, use the `ALL_DB_LINKS` view. See ALL_DB_LINKS for more information.

For additional information, see:

- See CREATE DATABASE LINK for details on the procedure.
- See Create Database Links from Autonomous Database to a Publicly Accessible Autonomous Database with a Wallet (mTLS)

Drop Database Links

After you create a database link you can drop the database link.

1. Drop a database link to a target database using `DBMS_CLOUD_ADMIN.DROP_DATABASE_LINK`.

   For example:

```
BEGIN
     DBMS_CLOUD_ADMIN.DROP_DATABASE_LINK(
          db_link_name => 'SALESLINK' );
END;
/
```

2. Include the `public_link` parameter when the database link is a public database link.

For example:

```
BEGIN
     DBMS_CLOUD_ADMIN.DROP_DATABASE_LINK(
          db_link_name => 'SALESLINK'
          public_link => TRUE);
END;
/
```

To run `DBMS_CLOUD_ADMIN.DROP_DATABASE_LINK` with `public_link` set to `TRUE`, you must have the `DROP PUBLIC DATABASE LINK` system privilege.

See DROP_DATABASE_LINK Procedure for detailed information about the procedure.

Call Web Services from Autonomous Database

Describes options for calling Web Services from Autonomous Database.

There are a number of options for calling Web Services from Autonomous Database, including the following:

- Use `DBMS_CLOUD` REST APIs: The `DBMS_CLOUD.SEND_REQUEST` function begins an HTTP request, gets the response, and ends the response. This function provides a workflow for sending a cloud REST API request with arguments and the function returns a response code and payload. See SEND_REQUEST Function and Procedure for more information.
- Use Oracle APEX: You can interact with both SOAP and RESTful style web services from APEX in your Autonomous Database instance. See Use Web Services with Oracle APEX for more information.
- Use `UTL_HTTP` to submit a request to a public site: See Submit an HTTP Request to a Public Host for more information.
- Use `UTL_HTTP` to submit a request to a private site: See Submit an HTTP Request to a Private Host for more information.

  When your Autonomous Database instance is on a private endpoint you can use a customer-managed wallet with procedures in `UTL_HTTP`, `DBMS_LDAP`, `UTL_SMTP`, or `UTL_TCP`. See Make External Calls Using a Customer-Managed Wallet for more information.

See PL/SQL Package Notes for Autonomous Database for information on restrictions for `UTL_HTTP` on Autonomous Database.

Topics

- Submit an HTTP Request to a Public Host
  Provides details for using `UTL_HTTP` to submit an HTTP request on a public host.
- Submit an HTTP Request to a Private Host
  Describes the steps to use `UTL_HTTP` to submit an HTTP request on a private host.
- Submit an HTTP Request to Private Site with a Proxy
  When your Autonomous Database instance is on a private endpoint you can use a proxy to submit HTTP requests with `UTL_HTTP`.
- Use Credential Objects to Set HTTP Authentication
  Describes how to pass a credential objects to `UTL_HTTP.SET_CREDENTIAL` procedure.

- Notes for Submitting HTTP Requests with Oracle APEX or Database Actions
  When you use Oracle APEX SQL Commands or Database Actions SQL worksheet to run multiple sequential SQL commands, the commands can run in different database sessions that do not save the state of a previous statement. This behavior differs from desktop SQL clients such as SQL*Plus and SQL Developer that maintain a persistent connection to the database.

Submit an HTTP Request to a Public Host

Provides details for using `UTL_HTTP` to submit an HTTP request on a public host.

For example, to submit an HTTP request for a public host www.example.com, create an Access Control List for the host:

```
BEGIN
   DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
         host => 'www.example.com',
         ace =>  xs$ace_type( privilege_list => xs$name_list('http'),
                              principal_name => 'ADMIN',
                              principal_type => xs_acl.ptype_db));
END;
```

Then submit the HTTP request:

```
SELECT UTL_HTTP.REQUEST(url => 'https://www.example.com/') FROM dual;
```

> **Note:**
>
> If your Autonomous Database instance is on a private endpoint and you want your `UTL_HTTP` calls to public hosts to be subject to your private endpoint VCN's egress rules, set the `ROUTE_OUTBOUND_CONNECTIONS` database property to `PRIVATE_ENDPOINT`. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

See PL/SQL Package Notes for Autonomous Database for information on restrictions for `UTL_HTTP` on Autonomous Database.

Submit an HTTP Request to a Private Host

Describes the steps to use `UTL_HTTP` to submit an HTTP request on a private host.

To submit a request to a target host on a private endpoint, the target host must be accessible from the source database's Oracle Cloud Infrastructure VCN. For example, you can connect to the target host when:

- Both the source database and the target host are in the same Oracle Cloud Infrastructure VCN.
- The source database and the target host are in different Oracle Cloud Infrastructure VCNs that are paired.
- The target host is an on-premises network that is connected to the source database's Oracle Cloud Infrastructure VCN using FastConnect or VPN.

You can also make UTL_HTTP calls with a customer-managed wallet when your Autonomous Database is on a private endpoint. See Make External Calls Using a Customer-Managed Wallet for more information.

To make a UTL_HTTP request to a target on a private endpoint:

1. Create an Access Control List for the host.

   For example:

```
BEGIN
   DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
         host => 'www.example.com',
         ace => xs$ace_type( privilege_list => xs$name_list('http'),
                             principal_name => 'ADMIN',
                             principal_type => xs_acl.ptype_db),
                             private_target => TRUE);
END;
/
```

   As shown in this example, when you create an Access Control List for the host specify the private_target parameter with the value TRUE.

> **Note:**
>
> If you set ROUTE_OUTBOUND_CONNECTIONS to PRIVATE_ENDPOINT, setting the private_target parameter to TRUE is not required in this API. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

2. Submit the HTTP request:

```
SELECT UTL_HTTP.REQUEST(
               url => 'https://www.example.com/',
               https_host => 'www.example.com')
            FROM dual;
```

See PL/SQL Package Notes for Autonomous Database for information on restrictions for UTL_HTTP on Autonomous Database.

Submit an HTTP Request to Private Site with a Proxy

When your Autonomous Database instance is on a private endpoint you can use a proxy to submit HTTP requests with UTL_HTTP.

When your Autonomous Database instance is on a private endpoint, to use UTL_HTTP with a target proxy the target proxy must be accessible from the source database's Oracle Cloud Infrastructure VCN.

For example, you can connect using a proxy when:

- Both the source database and the proxy server are in the same Oracle Cloud Infrastructure VCN.
- The source database and the proxy server are in different Oracle Cloud Infrastructure VCNs that are paired.

- The proxy server is an on-premises network that is connected to the source database's Oracle Cloud Infrastructure VCN using FastConnect or VPN.

You can also make `UTL_HTTP` calls using a customer-managed wallet. See Make External Calls Using a Customer-Managed Wallet for more information.

To use a proxy server with `UTL_HTTP`:

1. Set the HTTP_PROXY ACL on the proxy server.

   For example:

```
BEGIN
  DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
       host =>'www-proxy-example.com',
       ace  => xs$ace_type(privilege_list => xs$name_list('HTTP_PROXY'),
                           principal_name => 'APPUSER1',
                           principal_type => xs_acl.ptype_db),
                           private_target => TRUE);
END;
/
```

   As shown in this example, when you create an Access Control List for the proxy server specify the `private_target` parameter with the value `TRUE`.

> **Note:**
>
> If you set `ROUTE_OUTBOUND_CONNECTIONS` to `PRIVATE_ENDPOINT`, setting the `private_target` parameter to `TRUE` is not required in this API. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

2. Set the HTTP ACL on the remote Web Server.

   For example:

```
BEGIN
  DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
        host =>'example.com',
        ace => xs$ace_type( privilege_list => xs$name_list('HTTP'),
                            principal_name => 'APPUSER1',
                            principal_type => xs_acl.ptype_db)
END;
/
```

3. Set the wallet and the proxy for `UTL_HTTP`.

   For example:

```
BEGIN
   UTL_HTTP.SET_WALLET('');
   UTL_HTTP.SET_PROXY('www-proxy-example:80');
END;
/
```

4. Submit an HTTP request:

```
SELECT UTL_HTTP.REQUEST(
                    url         => 'https://www.example.com/',
                    https_host  => 'www.example.com')
               FROM dual;
```

Notes for setting a proxy server with `UTL_HTTP.SET_PROXY`:

- `DBMS_CLOUD` requests do not honor the proxy server you set with `UTL_HTTP.SET_PROXY`. This includes `DBMS_CLOUD.SEND_REQUEST` and all object storage access for `DBMS_CLOUD` external tables that you define with `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`, `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE`, or `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE`.
- `APEX_WEB_SERVICE` requests do not honor the proxy server you set with `UTL_HTTP.SET_PROXY`.

See PL/SQL Package Notes for Autonomous Database for information on restrictions for `UTL_HTTP` on Autonomous Database.

Use Credential Objects to Set HTTP Authentication

Describes how to pass a credential objects to `UTL_HTTP.SET_CREDENTIAL` procedure.

The `UTL_HTTP.SET_CREDENTIAL` procedure sets HTTP authentication information in the HTTP request header. The Web server needs this information to authorize the request.

The `UTL_HTTP.SET_CREDENTIAL` procedure enables you to pass credential objects to set HTTP authentication. Credential objects are schema objects, hence they can be accessed only by privileged users and enable you to configure schema-level privileges to access control the credentials. Passing credential objects is an appropriate and secure way to store and manage username/password/keys to be used for authentication.

The `UTL_HTTP.SET_CREDENTIAL` procedure is a secure and convenient alternative to `UTL_HTTP.SET_AUTHENTICATION` procedure.

Example

```
...
UTL_HTTP.SET_AUTHENTICATION (l_http_request, 'web_app_user', 'xxxxxxxxxxxx');
...
```

As shown in the example above, when you invoke `SET_AUTHENTICATION` procedure, you must pass the username/password in clear text as part of PL/SQL formal parameters. You might need to embed the username/password into various PL/SQL automation or cron scripts. Passing clear text passwords is a compliance issue that is addressed in `UTL_HTTP.SET_CREDENTIAL` procedure.

See SET_AUTHENTICATION Procedure and SET_AUTHENTICATION_FROM_WALLET Procedure for more information.

`UTL_HTTP.SET_CREDENTIAL` Syntax

```
UTL_HTTP.SET_CREDENTIAL (
    r          IN OUT NOCOPY req,
    credential IN VARCHAR2,
```

```
    scheme      IN VARCHAR2 DEFAULT 'Basic',
    for_proxy   IN BOOLEAN  DEFAULT FALSE);
```

Example to pass a credential object in the `SET_CREDENTIAL` procedure:

- Create a credential object:

  ```
  BEGIN DBMS_CLOUD.CREATE_CREDENTIAL (
      credential_name => 'HTTP_CRED',
      username        => 'web_app_user',
      password        => '<password>' );
  END;
  ```

  This creates a credential object which creates a stored username/password pair.

  See CREATE_CREDENTIAL Procedure for more information.

  See Specifying Scheduler Job Credentials for more information.

- Invoke `UTL_HTTP.SET_CREDENTIAL` procedure:

  ```
  DECLARE
        l_http_request  UTL_HTTP.REQ;
      BEGIN
        l_http_request := UTL_HTTP.BEGIN_REQUEST('https://www.example.com/v1/
  dwcsdev/NAME/dwcs_small_xt1.csv');
        UTL_HTTP.SET_CREDENTIAL (l_http_request, 'HTTP_CRED','BASIC');
        ...
  END;
  ```

  This example first creates a request by invoking the `BEGIN_REQUEST` procedure and sets HTTP authentication information in the HTTP request header by invoking the `SET_CREDENTIAL` procedure. The Web server needs this information to authorize the request. The value `l_http_request` is the HTTP request, `HTTP_CRED` is the credentials name and `BASIC` is the HTTP authentication scheme.

See UTL_HTTP for more information.

See PL/SQL Package Notes for Autonomous Database for information on restrictions for `UTL_HTTP` on Autonomous Database.

Notes for Submitting HTTP Requests with Oracle APEX or Database Actions

When you use Oracle APEX SQL Commands or Database Actions SQL worksheet to run multiple sequential SQL commands, the commands can run in different database sessions that do not save the state of a previous statement. This behavior differs from desktop SQL clients such as SQL*Plus and SQL Developer that maintain a persistent connection to the database.

Oracle APEX SQL Commands and Database Actions SQL worksheet submissions to an Autonomous Database instance are stateless. This means running individual SQL and PL/SQL statements may save state in database memory, for example when you submit a command to use a wallet, but the state might be cleared before you run the next statement.

View the following table for the steps to keep the database memory state between statement executions for SQL commands you submit to Autonomous Database.

SQL Command Tool	Submit Statements as a Block
Database Actions SQL Worksheet	• Select all statements and click **Run Statement** • Select nothing and click **Run as SQL script**
Oracle APEX SQL Commands	APEX SQL Commands only supports running individual statements. When you want to run multiple statements you must wrap the statements in a single PL/SQL anonymous block. To run the block with APEX SQL Commands, click **Run**.

For example, use the following code block to run a `utl_http.request()` command that makes use of a Customer-Managed Wallet:

```
SELECT utl_http.request(url => 'https://api.example.com/', wallet_path =>
'file:path_to_wallet', wallet_password => 'password' ) FROM DUAL";
```

Compare this to running with two consecutive statements that might fail if the `utl_http.set_wallet()` command and the `utl_http.request()` statement run individually, rather than as a single code block:

```
EXEC utl_http.set_wallet('file:WALLET_DIR/wallet.sso', 'password');
SELECT utl_http.request('https://api.example.com/') FROM DUAL;
```

Move Files

Describes how to create directories and copy files to Autonomous Database.

- Creating and Managing Directories on Autonomous Database
- Copy Files Between Object Store and a Directory in Autonomous Database
  Use the procedure `DBMS_CLOUD.PUT_OBJECT` to copy a file from a directory to Object Store. Use the procedure `DBMS_CLOUD.GET_OBJECT` to copy a file from Object Store to a directory.
- Bulk Operations for Files in the Cloud
  The PL/SQL package `DBMS_CLOUD` offers parallel execution support for bulk file upload, download, copy, and transfer activities, which streamlines the user experience and delivers optimal performance for bulk file operations.

Creating and Managing Directories on Autonomous Database

Autonomous Database includes a predefined `data_pump_dir` directory in the database where you can place files. You can also create directories, drop directories, and attach network file systems.

For example you can use the `data_pump_dir` directory for Data Pump import and export operations. To create additional directories use the database `CREATE DIRECTORY` command. Use the database `DROP DIRECTORY` command to drop directories, and use `DBMS_CLOUD.LIST_FILES` to list the contents of a directory. You can attach a Network File System to a directory location in your Autonomous Database with the `DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM` procedure.

- Create Directory in Autonomous Database
  To create directories use the database `CREATE DIRECTORY` command. Using `CREATE DIRECTORY` you specify the path as a relative path for the new directory.

- Drop Directory in Autonomous Database
  Use the database `DROP DIRECTORY` command to drop a directory object.
- List Contents of Directory in Autonomous Database
  Use the function `DBMS_CLOUD.LIST_FILES` to list the contents of a directory.
- Access Network File System from Autonomous Database
  You can attach a Network File System to a directory location in your Autonomous Database.
- Load Data from Directories in Autonomous Database
  As an alternative to an object store location URI, you can specify a directory with `DBMS_CLOUD` procedures to load or unload data from files in a local directory, including directories created on attached network file systems.

Create Directory in Autonomous Database

To create directories use the database `CREATE DIRECTORY` command. Using `CREATE DIRECTORY` you specify the path as a relative path for the new directory.

`CREATE DIRECTORY` creates the database directory object and also creates the file system directory if it does not already exist. If the file system directory exists then `CREATE DIRECTORY` only creates the database directory object. For example, the following command creates the database directory named `staging` and creates the file system directory `stage`:

```
CREATE DIRECTORY staging AS 'stage';
```

You can also create subdirectories. For example, the following command creates the database directory object `sales_staging` and the file system directory `stage/sales`:

```
CREATE DIRECTORY sales_staging AS 'stage/sales';
```

When you create subdirectories you do not have to create the initial file system directory. For example, in the previous example if the directory `stage` does not exist then the `CREATE DIRECTORY` command creates both directories `stage` and `stage/sales`.

To add a directory, you must have the `CREATE ANY DIRECTORY` system privilege. The ADMIN user is granted the `CREATE ANY DIRECTORY` system privilege. The ADMIN user can grant `CREATE ANY DIRECTORY` system privilege to other users.

See CREATE DIRECTORY for more information.

Notes:

- `CREATE DIRECTORY` creates the database directory object in the database and also creates the file system directory. For example the directory path could be:

  ```
  /u03/dbfs/7C149E35BB1000A45FD/data/stage
  ```

- You can create a directory in the root file system to see all the files with the following commands:

  ```
  CREATE OR REPLACE DIRECTORY ROOT_DIR AS '';
  ```

  After you create the `ROOT_DIR` directory, use the following command to list all files:

  ```
  SELECT * FROM DBMS_CLOUD.list_files('ROOT_DIR');
  ```

  To run `DBMS_CLOUD.LIST_FILES` with a user other than ADMIN you need to grant read privileges on the directory to that user. See LIST_FILES Function for more information.
- Space in the file system allocated for the directories you create and their contents is part of your storage allocation. See Database Dashboard Overview to view the total space allocated.

Drop Directory in Autonomous Database

Use the database `DROP DIRECTORY` command to drop a directory object.

For example, the following command drops the database directory object `staging`:

```
DROP DIRECTORY staging;
```

The `DROP DIRECTORY` command does not delete files in the directory. If you want to delete the directory and the files in the directory, first use the procedure `DBMS_CLOUD.DELETE_FILE` to delete the files. See DELETE_FILE Procedure for more information.

To drop a directory, you must have the `DROP ANY DIRECTORY` system privilege. The ADMIN user is granted the `DROP ANY DIRECTORY` system privilege. The ADMIN user can grant `DROP ANY DIRECTORY` system privilege to other users.

See DROP DIRECTORY for more information.

Notes:

- You are not allowed to drop the predefined directories: `data_pump_dir` or `sql_tcb_dir`
- If you just want to drop the directory and you do not remove the files in the directory, after you drop the directory you can view all the files in the file system, including any files that were in the directory you dropped, as follows:

  ```
  CREATE OR REPLACE DIRECTORY ROOT_DIR AS '';
  ```

  Then list the contents of `ROOT_DIR` with the following command:

  ```
  SELECT * FROM DBMS_CLOUD.list_files('ROOT_DIR');
  ```

  To run `DBMS_CLOUD.LIST_FILES` with a user other than ADMIN you need to grant read privileges on the directory to that user. See LIST_FILES Function for more information.
- The `DROP DIRECTORY` command does not remove the underlying file system directory. Autonomous Database manages the underlying file system directory; users do not remove the file system directory.

List Contents of Directory in Autonomous Database

Use the function `DBMS_CLOUD.LIST_FILES` to list the contents of a directory.

For example, to list the contents of the `staging` directory, run the following query:

```
SELECT * FROM DBMS_CLOUD.LIST_FILES('STAGING');
```

To run `DBMS_CLOUD.LIST_FILES` with a user other than ADMIN you need to grant read privileges on the directory to that user. See LIST_FILES Function for more information.

Access Network File System from Autonomous Database

You can attach a Network File System to a directory location in your Autonomous Database.

This allows you to load data from Oracle Cloud Infrastructure File Storage in your Virtual Cloud Network (VCN) or from any other Network File System in on-premises data centers. Depending on the version of the Network File System you want to access, both NFSv3 and NFSv4 are supported.

Supporting Network File System allows you to do the following:

- Connect to an Autonomous Database instance from a legacy application and use the file system to load and unload data.
- Analyze data from different sources in an Autonomous Database.
- Secure access to data in an Autonomous Database from the file systems in an on-premises data center or Private Virtual Cloud Networks (VCNs).

Topics

- Attach Network File System to Autonomous Database
  Use `DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM` to attach a file system to a directory in your Autonomous Database.
- Detach Network File System from Autonomous Database
  Use the `DBMS_CLOUD_ADMIN.DETACH_FILE_SYSTEM` procedure to detach a file system from a directory in your Autonomous Database.
- Example: Set Up an NFSv4 Server on Oracle Cloud Compute
  Provides an example for setting up an NSFv4 server for use with Autonomous Database.
- DBA_CLOUD_FILE_SYSTEMS View
  The `DBA_CLOUD_FILE_SYSTEMS` view lists information about the network file system attached to a directory location in the database.

Attach Network File System to Autonomous Database

Use `DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM` to attach a file system to a directory in your Autonomous Database.

With an attached file system you can load data from any of the following:

- Oracle Cloud Infrastructure File Storage in your Virtual Cloud Network (VCN).

  See How to Attach a File System to your Autonomous Database for details on setting up Oracle Cloud Infrastructure File Storage with Autonomous Database.
- From a Network File System in an on-premises data center. Depending on the version of the Network File System you want to access, both NFSv3 and NFSv4 are supported.

  See Example: Set Up an NFSv4 Server on Oracle Cloud Compute for a configuration example with an NFSv4 Network File System.

> **Note:**
>
> The `DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM` procedure can only attach a private File Storage Service when the Autonomous Database instance is on a private endpoint.

To access data in an Autonomous Database from the file systems in an on-premises data center you must set up FastConnect or a Site-to-Site VPN to connect to the on-premises data center. See FastConnect and Site-to-Site VPN for more information.

1. Create a directory or use an existing directory to attach a Network File System in your Autonomous Database. You must have `WRITE` privilege on the directory object on your Autonomous Database instance to attach a file system to a directory location in the database.

   For example, the following command creates the database directory named `FSS_DIR` and creates the file system directory `fss`:

   ```
   CREATE DIRECTORY FSS_DIR AS 'fss';
   ```

   See Create Directory in Autonomous Database for more information.

2. Run `DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM` to attach a file system to a directory in your Autonomous Database. To run this procedure, you must be logged in as the ADMIN user or have `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.

   - By default `DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM` uses NFSv3:

     ```
     BEGIN
        DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM(
         file_system_name      => 'FSS',
         file_system_location  => 'myhost.sub000445.myvcn.oraclevcn.com:/
     results',
         directory_name        => 'FSS_DIR',
         description           => 'Source NFS for sales data'
     );
     END;
     /
     ```

     Optionally you can use the `params` parameter and specify the `nfs_version` with value 3 to specify NFSv3.

   - To use NFSv4, include the `params` parameter with `DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM` and specify the `nfs_version` with value 4 to specify NFSv4:

     ```
     BEGIN
        DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM(
         file_system_name      => 'FSS',
         file_system_location  => 'myhost.sub000445.myvcn.oraclevcn.com:/
     results',
         directory_name        => 'FSS_DIR',
         description           => 'Source NFS for sales data',
         params                => JSON_OBJECT('nfs_version' value 4)
     );
     END;
     /
     ```

   These examples attach the network file system specified in the `file_system_name` parameter to the Autonomous Database.

   The `file_system_location` parameter specifies the location of the file system. The value you supply with `file_system_location` consists of a Fully Qualified Domain Name (FQDN) and a file path in the form: *FQDN*:*file_path*.

   For example:

   - FQDN: `myhost.sub000445.myvcn.oraclevcn.com`

     For Oracle Cloud Infrastructure File Storage set the FQDN in **Show Advanced Options** when you create a file system. See Creating File Systems for more information.

   - File Path: `/results`

   The `directory_name` parameter specifies the directory name in the Autonomous Database where you want to attach the file system. This is the directory you created in Step 1, or another directory you previously created.

   The `description` parameter specifies the description for the task.

The `params` parameter is a JSON value that specifies an additional attribute `nfs_version`, whose value can be either 3 or 4 (NFSv3 or NFSv4).

After you attach a file system you can query the `DBA_CLOUD_FILE_SYSTEMS` view to retrieve information about the attached file system.

For example:

```
SELECT file_system_name, file_system_location,
directory_path
   FROM dba_cloud_file_systems
   WHERE file_system_name = 'FSS';
```

This query returns details for the `FSS` file system name.

See DBA_CLOUD_FILE_SYSTEMS View for more information.

With an attached file system you can read and write to files on an attached file system using any PL/SQL API that accepts a directory name. For example, you can use any of the following methods to work with an attached NFS directory:

- The `UTL_FILE` package.
- Data Pump Export and Import utilities.
- The `DBMS_CLOUD` APIs that work with directories such as `DBMS_CLOUD.LIST_FILES` and `DBMS_CLOUD.PUT_OBJECT`.

Example showing a write a file on an attached file system using `UTL_FILE`:

```
DECLARE
  l_file         UTL_FILE.FILE_TYPE;
  l_location     VARCHAR2(100) := 'FSS_DIR';
  l_filename     VARCHAR2(100) := 'test.csv';
BEGIN
  -- Open the file.
  l_file := UTL_FILE.FOPEN(l_location, l_filename, 'w');

  UTL_FILE.PUT(l_file, 'Scott, male, 1000');

  -- Close the file.
  UTL_FILE.FCLOSE(l_file);
END;
/
```

Example showing a read a file on an attached file system using `UTL_FILE`:

```
DECLARE
  l_file         UTL_FILE.FILE_TYPE;
  l_location     VARCHAR2(100) := 'FSS_DIR';
  l_filename     VARCHAR2(100) := 'test.csv';
  l_text         VARCHAR2(32767);
BEGIN
  -- Open the file.
  l_file := UTL_FILE.FOPEN(l_location, l_filename, 'r');

  UTL_FILE.GET_LINE(l_file, l_text, 32767);
```

```
  -- Close the file.
  UTL_FILE.FCLOSE(l_file);
END;
/
```

Example showing list files on an attached file system using `DBMS_CLOUD.LIST_FILES`:

```
SELECT object_name FROM DBMS_CLOUD.LIST_FILES('FSS_DIR');
```

Notes for using `DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM`:

- Oracle Cloud Infrastructure File Storage uses NFSv3 to share. See Overview of File Storage for more information.
- If you attach to non-Oracle Cloud Infrastructure File Storage systems, the procedure supports NFSv3 and NFSv4.
- If you have an attached NFS server that uses NFSv3 and the NFS version is updated to NFSv4 in the NFS server, you must run `DBMS_CLOUD_ADMIN.DETACH_FILE_SYSTEM` and then `DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM` (using the `params` parameter with `nfs_version` set to 4). This attaches NFS with the matching protocol so that Autonomous Database can access the NFSv4 Server. Without detaching and then reattaching, the NFS server will be inaccessible and you may see an error such as: `"Protocol not supported"`.

See the following for more information:

- UTL_FILE
- LIST_FILES Function
- OCI File Storage Service
- ATTACH_FILE_SYSTEM Procedure

Detach Network File System from Autonomous Database

Use the `DBMS_CLOUD_ADMIN.DETACH_FILE_SYSTEM` procedure to detach a file system from a directory in your Autonomous Database.

> **Note:**
>
> The `DBMS_CLOUD_ADMIN.DETACH_FILE_SYSTEM` procedure can only detach a private File Storage Service from databases that are on private endpoints.

You must have `WRITE` privilege on the directory object to detach a file system from a directory location.

Run `DBMS_CLOUD_ADMIN.DETACH_FILE_SYSTEM` procedure to detach a file system from a directory location in your Autonomous Database. To run this procedure, you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.

For example:

```
BEGIN
   DBMS_CLOUD_ADMIN.DETACH_FILE_SYSTEM (
    file_system_name      => 'FSS'
```

```
  );
END;
/
```

This example detaches the network file system specified in the `file_system_name` parameter from the Autonomous Database. You must provide a value for this parameter.

The information about this file system is removed from the `DBA_CLOUD_FILE_SYSTEMS` view.

See the following for more information:

- OCI File Storage Service
- Creating and Managing Directories on Autonomous Database
- DETACH_FILE_SYSTEM Procedure

Example: Set Up an NFSv4 Server on Oracle Cloud Compute

Provides an example for setting up an NSFv4 server for use with Autonomous Database.

1. Set up a private endpoint for the Autonomous Database instance.

   See Configuring Network Access with Private Endpoints for more information.

   The following ingress and egress rules need to be set for your VCN's security list so that Autonomous Database and the NFSv4 server can talk to each other

   - Stateful ingress from ALL ports in source CIDR block to TCP port 2049.
   - Stateful egress from TCP ALL ports to port 2049 in destination CIDR block.

2. Set up the NFS server on an Oracle Cloud VM with Oracle Linux 8 in the private subnet, which can connect to the Autonomous Database instance.

```
Compute Instance Image: Oracle-Linux-8.8-2023.09.26-0
$ sudo su -
$ mkdir /exports
$ chown nobody /exports
$ chgrp nobody /exports

# If the VM is using Linux 7, "data" directory may need having the
privilege 777 so that ADB can have access to NFS.
uname -a
chmod 777 /exports/data

$ mkdir /exports/data
$ chown nobody /exports/data
$ chgrp nobody /exports/data

# Either the private IP or the private FQDN can be used in "/etc/exports".
# Both can be found in the information of the autonomous database on the
OCI console.
$ tee -a /etc/exports <<'EOF'
  /exports/data *(rw,insecure)
  /exports/data example.adb.us-phoenix-1.oraclecloud.com(rw,insecure)
  EOF

$ systemctl start nfs-server
$ systemctl enable nfs-server
```

```
# Configure the firewall to allow NFS connections. "public" is the default
zone on Oracle Cloud VM.
$ firewall-cmd --get-default-zone
public

$ firewall-cmd --zone=public --add-service=nfs
$ firewall-cmd --permanent --zone=public --add-service=nfs

$ firewall-cmd --reload

# Display a list of the exported file systems.
$ showmount -e

# Displays all of the current clients and all of the file systems that the
clients have mounted.
$ showmount -a
```

3. Mount the NFSv4 file system with `DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM`.

   See Attach Network File System to Autonomous Database for more information.

DBA_CLOUD_FILE_SYSTEMS View

The `DBA_CLOUD_FILE_SYSTEMS` view lists information about the network file system attached to a directory location in the database.

Column	Data Type	Description
`FILE_SYSTEM_NAME`	`VARCHAR2(128)`	File system name
`FILE_SYSTEM_LOCATION`	`VARCHAR2(4000)`	File system location
`DIRECTORY_NAME`	`VARCHAR2(128)`	Attached directory name
`DIRECTORY_PATH`	`VARCHAR2(4000)`	Attached directory path
`NFS_VERSION`	`NUMBER`	The NFS version.
`DESCRIPTION`	`VARCHAR2(4000)`	The value provided for the description parameter when you run `DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM`
`CREATION_TIME`	`TIMESTAMP(6) WITH TIME ZONE`	Creation timestamp
`UPDATE_TIME`	`TIMESTAMP(6) WITH TIME ZONE`	Update timestamp

Load Data from Directories in Autonomous Database

As an alternative to an object store location URI, you can specify a directory with `DBMS_CLOUD` procedures to load or unload data from files in a local directory, including directories created on attached network file systems.

The following procedures support specifying files in a directory with the `file_uri_list` parameter:

- `DBMS_CLOUD.COPY_COLLECTION`
- `DBMS_CLOUD.COPY_DATA`

- DBMS_CLOUD.CREATE_EXTERNAL_TABLE

The following procedures support specifying files in a directory with the partitioning_clause parameter:

- DBMS_CLOUD.CREATE_EXTERNAL_TABLE
- DBMS_CLOUD.CREATE_HYBRID_PART_TABLE

You can specify one directory and one or more file names or use a comma separated list of directories and file names. The format to specify a directory is:'*MY_DIR*:*filename.ext*'. By default the directory name *MY_DIR* is a database object and is case-insensitive. The file name is case sensitive.

When you use the file_uri_list parameter to specify a directory you do not need to include the credential_name parameter, but you need READ object privileges on the directory.

For example, with a call to DBMS_CLOUD.COPY_DATA, use the file_uri_list parameter to specify files in a directory:

```
BEGIN
  DBMS_CLOUD.COPY_DATA(
     table_name => 'HRDATA1',
     file_uri_list => 'HR_DIR:test.csv',
     format => JSON_OBJECT('type' value 'csv')  );
END;
/
```

This example copies the data from test.csv in the local directory HR_DIR to the table HRDATA1.

Regular expressions are not supported when specifying the file names in a directory. You can only use wildcards to specify file names in a directory. The character "*" can be used as the wildcard for multiple characters, and the character "?" can be used as the wildcard for a single character. For example:'*MY_DIR*:*" or '*MY_DIR*:test?'

To specify multiple directories, use a comma separated list of directories: For example:'*MY_DIR1*:*, *MY_DIR2*:test?'

Use double quotes to specify a case-sensitive directory name. For example:'"*my_dir1*":*, "*my_dir2*":Test?'

To include a quote character, use two quotes. For example:'*MY_DIR*:''filename.ext'. This specifies the filename starts with a quote (').

See Attach Network File System to Autonomous Database for information on attaching network file systems.

Notes for Using Directories with DBMS_CLOUD Procedures

Note the following when you use DBMS_CLOUD procedures and specify a directory with the file_uri_list parameter:

- Compression options for files such as GZIP are not supported for directory files. See the compression format option in DBMS_CLOUD Package Format Options for more information.
- Special characters such as colon (:), single quote('), and comma(,) are not supported in the directory name.

Copy Files Between Object Store and a Directory in Autonomous Database

Use the procedure `DBMS_CLOUD.PUT_OBJECT` to copy a file from a directory to Object Store. Use the procedure `DBMS_CLOUD.GET_OBJECT` to copy a file from Object Store to a directory.

For example, to copy a file from Object Store to the `stage` directory, run the following command:

```
BEGIN
   DBMS_CLOUD.GET_OBJECT(
   credential_name => 'DEF_CRED_NAME',
   object_uri => 'https://objectstorage.usphoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/cwallet.sso',
   directory_name => 'STAGE');
END;
/
```

Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

To run `DBMS_CLOUD.GET_OBJECT` with a user other than ADMIN you need to grant write privileges on the directory to that user.

To run `DBMS_CLOUD.PUT_OBJECT` with a user other than ADMIN you need to grant read privileges on the directory to that user.

See GET_OBJECT Procedure and Function and PUT_OBJECT Procedure for more information.

Bulk Operations for Files in the Cloud

The PL/SQL package `DBMS_CLOUD` offers parallel execution support for bulk file upload, download, copy, and transfer activities, which streamlines the user experience and delivers optimal performance for bulk file operations.

The package `DBMS_CLOUD` supports loading and unloading files into the following cloud services:

- Oracle Cloud Infrastructure Object Storage
- Azure Blob Storage or Azure Data Lake Storage
- Amazon S3
- Amazon S3-Compatible, including Wasabi Hot Cloud Storage
- Google Cloud Storage

See DBMS_CLOUD Package File URI Formats for more information.

- About Bulk File Operations

- Bulk Copy Files in Cloud Object Storage
  Use the `DBMS_CLOUD.BULK_COPY` procedure to bulk copy files from one Object Store bucket or folder to another bucket or folder.
- Bulk Move Files Across Cloud Object Storage
- Bulk Download Files from Cloud Object Storage
  Use the `DBMS_CLOUD.BULK_DOWNLOAD` procedure to bulk download files from the Cloud Object Store location to the directory object in an Autonomous Database.
- Bulk Upload Files to Cloud Object Storage
  Use the `DBMS_CLOUD.BULK_UPLOAD` procedure to bulk upload files from a directory object in database to a Cloud Object Store location.
- Bulk Delete Files from Cloud Object Storage
  Use the `DBMS_CLOUD.BULK_DELETE` procedure to bulk delete files from Cloud Object Storage.
- Monitor and Troubleshoot Bulk File Loads
  All `DBMS_CLOUD` data load operations are logged in the `dba_load_operations` and `user_load_operations` views.
- Notes for Bulk File Operations
  Provides notes for the credentials you use with bulk file operations.

About Bulk File Operations

`DBMS_CLOUD` bulk file operations support downloading and uploading files on Autonomous Database, including support for copying, moving, and deleting files across Cloud Object Stores, either within the same Object Store provider or between Object Store providers. The bulk file operations support parallelism and provide optimal performance for uploading, downloading, copying, and moving files. Parallelism for bulk file operations is handled by specifying a priority for an operation. The supported priorities are:

- `HIGH`: Use the Autonomous Database instance CPU compute count to determine the number of parallel files processed.
- `MEDIUM`: Use the concurrency limit for Medium service to determine the parallelism.
- `LOW`: Process files in serial order (no concurrency).

Running bulk operations with a higher priority uses more database resources and operations complete faster when parallelism can speed up the operation. A lower priority consumes less database resources and operations take longer to complete when parallelism can speed up the operation. When bulk operations involve small numbers of files that contain little data, specifying higher priority may not change the performance.

To increase parallelism for bulk file operations use `HIGH` priority and increase the number of CPUs assigned to your Autonomous Database instance. The maximum supported concurrent file operations is limited to 64 for bulk file operations.

The default priority is `MEDIUM`, which specifies that bulk file operations use the concurrency limit defined for the MEDIUM consumer group. See Manage Concurrency and Priorities on Autonomous Database for more information.

See DBMS_CLOUD for Bulk File Management for details on using the `format` parameter to specify the `priority` with bulk file operations.

Bulk Copy Files in Cloud Object Storage

Use the `DBMS_CLOUD.BULK_COPY` procedure to bulk copy files from one Object Store bucket or folder to another bucket or folder.

1. Create a credential object to access the source location.

   The source credential name, as specified with the `source_credential_name` parameter is by default also used as the credential for the target location.

   See CREATE_CREDENTIAL Procedure for more information.

2. When the source and target are in distinct Object Stores or have different accounts with the same cloud provider, create a credential to access the target location and include the `target_credential_name` parameter.

3. Run `DBMS_CLOUD.BULK_COPY` procedure to copy files in parallel from one Object Store bucket or folder to another bucket or folder which can be across cloud provider, accounts, and buckets. To run the procedure, you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD`.

   ```
   BEGIN
   DBMS_CLOUD.BULK_COPY (
         source_credential_name => 'OCI_CRED',
         source_location_uri    => 'https://objectstorage.us-
   phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname1/o',
         target_location_uri    => 'https://objectstorage.us-
   phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname2/o',
         format                 => JSON_OBJECT ('logretention' value 7,
   'logprefix' value 'BULKOP')
   );
   END;
   /
   ```

   This example bulk copies files from one Oracle Cloud Infrastructure Object Storage bucket to another.

   See BULK_COPY Procedure for more information.

   See DBMS_CLOUD URI Formats for more information.

For example, use `DBMS_CLOUD.BULK_COPY` to copy files from Amazon S3 to Oracle Cloud Infrastructure Object Storage.

```
BEGIN
DBMS_CLOUD.BULK_COPY(
      source_credential_name  => 'AWS_CRED',
      source_location_uri     => 'https://bucketname.s3-us-
west-2.amazonaws.com/',
      target_credential_name  => 'OCI_CRED',
      target_location_uri     => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname2/o',
      format                  =>  JSON_OBJECT ('logretention' value 7,
'logprefix' value 'BULKOP')
);
END;
/
```

Bulk Move Files Across Cloud Object Storage

Use the `DBMS_CLOUD.BULK_MOVE` procedure to bulk move files from one Cloud Object Storage location to another.

The first step in moving files is copying them to the target location. After the files are successfully copied, they are deleted from the source location.

The files are renamed rather than copied if Object Store allows renaming operations between source and target locations.

1. Create a credential object to access the source location.

   The source credential name, as specified with the `source_credential_name` parameter is by default also used as the credential for the target location.

   See CREATE_CREDENTIAL Procedure for more information.

2. When the source and target are in distinct Object Stores or have different accounts with the same cloud provider, create a credential to access the target location and include the `target_credential_name` parameter.

3. Run `DBMS_CLOUD.BULK_MOVE` procedure to bulk move files from one Cloud Object Storage location to another. To run the procedure, you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD`.

   ```
   BEGIN
   DBMS_CLOUD.BULK_MOVE (
        source_credential_name => 'OCI_CRED',
        source_location_uri    => 'https://objectstorage.us-
   phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname1/o',
        target_location_uri    => 'https://objectstorage.us-
   phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname2/o',
        format                 => JSON_OBJECT ('logretention' value 7,
   'logprefix' value 'BULKMOVE')
   );
   END;
   /
   ```

   This example bulk moves files from one Oracle Cloud Infrastructure Object Storage location to another.

   See BULK_MOVE Procedure for more information.

   See DBMS_CLOUD URI Formats for more information.

For example, use `DBMS_CLOUD.BULK_MOVE` to move files from Amazon S3 to Oracle Cloud Infrastructure Object Storage.

```
BEGIN
DBMS_CLOUD.BULK_MOVE(
      source_credential_name  => 'AWS_CRED',
      source_location_uri     => 'https://bucketname.s3-us-
west-2.amazonaws.com/',
      target_credential_name  => 'OCI_CRED',
      target_location_uri     => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname2/o',
      format                  =>  JSON_OBJECT ('logretention' value 7,
```

```
'logprefix' value 'BULKOP')
);
END;
/
```

Bulk Download Files from Cloud Object Storage

Use the `DBMS_CLOUD.BULK_DOWNLOAD` procedure to bulk download files from the Cloud Object Store location to the directory object in an Autonomous Database.

1. Create a credential to access your Cloud Object Storage.

   The credential name, is specified with the `credential_name` parameter.

   See CREATE_CREDENTIAL Procedure for more information.

2. Run `DBMS_CLOUD.BULK_DOWNLOAD` procedure to download files into an Autonomous Database directory from Cloud Object Storage. To run the procedure, you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD`.

   ```
   BEGIN
   DBMS_CLOUD.BULK_DOWNLOAD (
        credential_name => 'OCI_CRED',
        location_uri    => 'https://objectstorage.us-
   phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o',
        directory_name  => 'BULK_TEST',
        format          => JSON_OBJECT ('logretention' value 7, 'logprefix'
   value 'BULKOP')
    );
   END;
   /
   ```

   This example bulk downloads files from the Oracle Cloud Infrastructure Object Store location URI to the directory object in an Autonomous Database.

   > **Note:**
   >
   > To write the files in the target directory object, you must have the `WRITE` privilege on the directory object.

   See BULK_DOWNLOAD Procedure for more information.

   See DBMS_CLOUD URI Formats for more information.

Bulk Upload Files to Cloud Object Storage

Use the `DBMS_CLOUD.BULK_UPLOAD` procedure to bulk upload files from a directory object in database to a Cloud Object Store location.

1. Create a credential to access your Cloud Object Storage.

   The credential name, is specified with the `credential_name` parameter.

   See CREATE_CREDENTIAL Procedure for more information.

2. Run `DBMS_CLOUD.BULK_UPLOAD` procedure to copy files into Cloud Object Storage from a database directory on your Autonomous Database instance. To run the procedure, you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD`.

```
BEGIN
DBMS_CLOUD.BULK_UPLOAD (
     credential_name => 'OCI_CRED',
     location_uri    => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o',
     directory_name  => 'BULK_TEST',
     format          => JSON_OBJECT ('logretention' value 5, 'logprefix'
value 'BULKUPLOAD')
 );
END;
/
```

This example bulk uploads files from a directory object, as specified with the `directory_name` parameter to the Oracle Cloud Infrastructure Object Store location URI.

> **Note:**
>
> To read the source files in the directory object, you must have the `READ` privilege on the directory object.

See BULK_UPLOAD Procedure for more information.

Bulk Delete Files from Cloud Object Storage

Use the `DBMS_CLOUD.BULK_DELETE` procedure to bulk delete files from Cloud Object Storage.

1. Create a credential to access your Cloud Object Storage.

   The credential name, is specified with the `credential_name` parameter.

   The `credential_name` can be `NULL` for public or Pre-authenticated or Pre-signed bucket URI.

   See CREATE_CREDENTIAL Procedure for more information.

2. Run `DBMS_CLOUD.BULK_DELETE` procedure to delete files from the Cloud Object Store. To run the procedure, you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD`.

```
BEGIN
DBMS_CLOUD.BULK_DELETE (
     credential_name => 'OCI_CRED',
     location_uri    => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o',
     format          => JSON_OBJECT ('logretention' value 5, 'logprefix'
value 'BULKDEL')
);
END;
/
```

This example bulk deletes files from the Oracle Cloud Infrastructure Object Store.

See BULK_DELETE Procedure for more information.

See DBMS_CLOUD URI Formats for more information.

Monitor and Troubleshoot Bulk File Loads

All `DBMS_CLOUD` data load operations are logged in the `dba_load_operations` and `user_load_operations` views.

You can use the following views to monitor and troubleshoot bulk file loads:

- `dba_load_operations`: shows all load operations.
- `user_load_operations`: shows the load operations in your schema.

Query these views to see information about ongoing and completed bulk file operations. For example, the following `SELECT` statement with a `WHERE` clause predicate on the `TYPE` shows `DOWNLOAD` operations:

```
SELECT owner_name, type, status, start_time, update_time, status_table, rows_loaded
   FROM user_load_operations WHERE type = 'DOWNLOAD';

OWNER_NAME TYPE     STATUS    START_TIME               UPDATE_TIME              STATUS_TABLE
ROWS_LOADED
---------- -------- --------- ------------------------ ------------------------ ------------
-----------
"ADMIN"    DOWNLOAD COMPLETED 2022-10-17T20:42:19.498Z 2022-10-17T20:42:21.054Z
DWN$2_STATUS           4
"ADMIN"    DOWNLOAD FAILED    2022-10-17T20:40:54.348Z 2022-10-17T20:40:55.679Z
DWN$1_STATUS
```

The `STATUS_TABLE` column shows the name of the table you can query to look at detailed logging information for the bulk download operation. For example:

```
DESCRIBE DWN$2_STATUS
Name          Null?    Type
------------- -------- ---------------------------
ID            NOT NULL NUMBER
NAME                   VARCHAR2(4000)
BYTES                  NUMBER
CHECKSUM               VARCHAR2(128)
LAST_MODIFIED          TIMESTAMP(6) WITH TIME ZONE
STATUS                 VARCHAR2(30)
ERROR_CODE             NUMBER
ERROR_MESSAGE          VARCHAR2(4000)
START_TIME             TIMESTAMP(6) WITH TIME ZONE
END_TIME               TIMESTAMP(6) WITH TIME ZONE
SID                    NUMBER
SERIAL#                NUMBER
ROWS_LOADED            NUMBER

SELECT id, name, bytes, status, error_message, start_time, end_time FROM
DWN$2_STATUS;
ID NAME       BYTES STATUS    ERROR_MESSAGE START_TIME
```

```
END_TIME
-- ---------- ----- --------- ------------- ------------------------
------------------------
 1 trees.txt     58 COMPLETED               2022-10-17T20:42:19.998Z
2022-10-17T20:42:20.421Z
 2 trees1.txt    58 COMPLETED               2022-10-17T20:42:20.425Z
2022-10-17T20:42:20.533Z
 3 trees2.txt    58 COMPLETED               2022-10-17T20:42:20.535Z
2022-10-17T20:42:20.894Z
 4 trees3.txt    58 COMPLETED               2022-10-17T20:42:20.896Z
2022-10-17T20:42:21.043Z
```

The status table shows each file name and its status for the bulk operation.

The relevant error number and message are recorded in the status table if an operation on a specific file fails.

For completed operations, the time needed for each operation can be calculated using the reported `START_TIME` and `END_TIME` time.

The file operation `STATUS` column can have one of the following values:

File Status	Description
`COMPLETED`	File operation completed successfully.
`FAILED`	File operation failed, a retry may be attempted for two times.
`PENDING`	The file operation has not yet started.
`RUNNING`	File operation is currently in progress.
`SKIPPED`	File operation skipped.

If any file operation fails after two retry attempts, then the bulk operation is marked as failed and an error is raised. For example:

```
ORA-20003: Operation failed, please query table DOWNLOAD$2_STATUS for error details
```

When you use a `DBMS_CLOUD` bulk file operation there are `format` parameter options that control status tables:

- `logretention`: Specifies an integer value that determines the duration in days that the status table is retained. The default value is 2 days.
- `logprefix`: Specifies a string value that determines the bulk operation status table's name prefix.

  Each bulk operation has its own default value for the `logprefix` option:

Procedure	Default Value for `logprefix`
`DBMS_CLOUD.BULK_COPY`	`COPYOBJ`
`DBMS_CLOUD.BULK_DELETE`	`DELETE`
`DBMS_CLOUD.BULK_DOWNLOAD`	`DOWNLOAD`
`DBMS_CLOUD.BULK_MOVE`	`MOVE`
`DBMS_CLOUD.BULK_UPLOAD`	`UPLOAD`

See DELETE_ALL_OPERATIONS Procedure for information on clearing the `user_load_operations` table.

Notes for Bulk File Operations

Provides notes for the credentials you use with bulk file operations.

- You may specify principals as the `credential_name`, `source_credential_name` or `target_credential_name` parameter to access cloud resources securely without storing user credentials. The supported principals are:
  - Oracle Cloud Infrastructure Resource Principals
  - AWS Amazon Resource Names (ARN)s
  - Azure service principal
  - Google Service Account

  See Configure Policies and Roles to Access Resources for more information.
- The `credential_name`, `source_credential_name` or the `target_credential_name` can be `NULL` for public, Pre-authenticated, or Pre-signed bucket URI.

Replicate Data

Describes using Oracle GoldenGate to replicate data between an Autonomous Database instance and to any target database or platform that Oracle GoldenGate supports.

- Use Oracle GoldenGate to Replicate Data to Autonomous Database
  You can replicate data to Autonomous Database using Oracle GoldenGate.

Use Oracle GoldenGate to Replicate Data to Autonomous Database

You can replicate data to Autonomous Database using Oracle GoldenGate.

Oracle GoldenGate Capture for Oracle Autonomous Database supports the following:

- Replication for different use cases: Report Offloading, Active-Active, Cloud to Cloud, and Cloud to on-premise.
- Inter-region and cross-region replication: Replicate data between different Oracle Cloud data centers around the world.
- Replicate between targets: Replicate from an Autonomous Database to any target database or platform that Oracle GoldenGate supports, including to other Oracle Autonomous Database environments.

The following are not available for Always Free Autonomous Databases:

- Supplemental logging
- Oracle GoldenGate Extract

See Supplemental Logging for more information.

See Configure Replicat to Apply to Autonomous Databases for more information.

Exporting Data from Autonomous Database to Object Store or to Other Oracle Databases

To export data from an Autonomous Database, use one of the following methods:

- Export to Cloud Object Store:
  - Use the procedure `DBMS_CLOUD.EXPORT_DATA` to export data to your Cloud Object Store as text, by specifying a query. This export method supports exporting data as CSV, JSON, Parquet, or XML files.

    This export method supports all the Cloud Object Stores supported by Autonomous Database.
  - Use Oracle Data Pump to export data to a directory on your database, and then move the data from the directory to Cloud Object Storage.
  - Use Oracle Data Pump to export the data to Cloud Object Storage directly. This method is only supported with Oracle Cloud Infrastructure Object Storage and Oracle Cloud Infrastructure Object Storage Classic.
  - Use the procedure `DBMS_CLOUD.EXPORT_DATA` to export data to your Cloud Object Store as Oracle Data Pump dump files, by specifying a query. This method is only supported with Oracle Cloud Infrastructure Object Storage and Oracle Cloud Infrastructure Object Storage Classic.
- Export to a directory:
  - Use the procedure `DBMS_CLOUD.EXPORT_DATA` to export data to a directory as text by specifying a query. This export method supports exporting data as CSV, JSON, Parquet, or XML files.
  - Use Oracle Data Pump to export data to a directory.
- Use Oracle GoldenGate Capture for Oracle Autonomous Database.

Topics

- Export Data to Object Store as Text
- Export Data to a Directory
- Move Data with Data Pump Export to an Autonomous Database Directory
  Oracle Data Pump offers very fast bulk data and metadata movement between Autonomous Database and other Oracle databases. To move data from an Autonomous Database to another Oracle database, use Oracle Data Pump to export to a directory on your Autonomous Database.
- Move Data with Data Pump Export to Object Store
  To move data from Autonomous Databases to other Oracle databases you can use Oracle Data Pump. Using this export method you use Oracle Data Pump to directly export data to your object store. This export method is supported with Oracle Cloud Infrastructure Object Storage and Oracle Cloud Infrastructure Object Storage Classic.
- Move Data to Object Store as Oracle Data Pump Files Using EXPORT_DATA
  You can export data to Oracle Data Pump dump files by specifying a query.

- Download Dump Files, Run Data Pump Import, and Clean Up Object Store
  If required, download the dump files from Cloud Object Store and use Oracle Data Pump Import to import the dump file set to the target database. Then perform any required clean up.
- Access Log Files for Data Pump Export
  The log files for Data Pump Export operations are stored in the directory you specify with the data pump `directory` parameter.
- Oracle GoldenGate Capture for Oracle Autonomous Database
  Using Oracle GoldenGate you can capture changes from an Oracle Autonomous Database and replicate to any target database or platform that Oracle GoldenGate supports, including another Oracle Autonomous Database.

Export Data to Object Store as Text

Use `DBMS_CLOUD.EXPORT_DATA` to export data as text from an Autonomous Database to cloud Object Store. The text format export options are CSV, JSON, Parquet, or XML.

See Move Data to Object Store as Oracle Data Pump Files Using EXPORT_DATA for details on exporting to cloud Object Store in Oracle Data Pump dump format.

- Export JSON Data to Cloud Object Storage
  Shows the steps to export table data from your Autonomous Database to Cloud Object Storage as JSON data by specifying a query.
- Export Data as CSV to Cloud Object Storage
  Shows the steps to export table data from your Autonomous Database to Cloud Object Storage as CSV data by specifying a query.
- Export Data as Parquet to Cloud Object Storage
  Shows the steps to export table data from your Autonomous Database to Cloud Object Storage as Parquet data by specifying a query.
- Export Data as XML to Cloud Object Storage
  Shows the steps to export table data from your Autonomous Database to Cloud Object Storage as XML data by specifying a query.
- File Naming for Text Output (CSV, JSON, Parquet, or XML)
  Describes the output file naming using `DBMS_CLOUD.EXPORT_DATA` with CSV, JSON, Parquet, or XML text file output.

Export JSON Data to Cloud Object Storage

Shows the steps to export table data from your Autonomous Database to Cloud Object Storage as JSON data by specifying a query.

This export method supports all the Cloud Object Stores supported by Autonomous Database, and you can use an Oracle Cloud Infrastructure resource principal to access your Oracle Cloud Infrastructure Object Store, Amazon Resource Names (ARNs) to access AWS Simple Storage Service (S3), an Azure service principal to access Azure BLOB storage or Azure Data Lake Storage, or a Google service account to access Google Cloud Platform (GCP) resources.

1. Connect to your Autonomous Database instance.

   See Connect to Autonomous Database for more information.
2. Store your Cloud Object Storage credential using `DBMS_CLOUD.CREATE_CREDENTIAL`.

For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'user1@example.com',
    password => 'password'
  );
END;
/
```

The values you provide for `username` and `password` depend on the Cloud Object Storage service you are using.

Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

3. Run `DBMS_CLOUD.EXPORT_DATA` and specify the `format` parameter `type` with the value `json` to export the results as JSON files on Cloud Object Storage.

   To generate the JSON output files there are two options for the `file_uri_list` parameter:

   - Set the `file_uri_list` value to the URL for an existing bucket on your Cloud Object Storage.
   - Set the `file_uri_list` value to the URL for an existing bucket on your Cloud Object Storage and include a file name prefix to use when generating the file names for the exported JSON.

   If you do not include the file name prefix in the `file_uri_list`, `DBMS_CLOUD.EXPORT_DATA` supplies a file name prefix. See File Naming for Text Output (CSV, JSON, Parquet, or XML) for details.

   For example, the following shows `DBMS_CLOUD.EXPORT_DATA` with a file name prefix specified in `file_uri_list`:

```
BEGIN
  DBMS_CLOUD.EXPORT_DATA(
    credential_name => 'DEF_CRED_NAME',
    file_uri_list   => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/dept_export',
    query           => 'SELECT * FROM DEPT',
    format          => JSON_OBJECT('type' value 'json'));
END;
/
```

   In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

   When record delimiters include escape characters, such as \r\n or \t, enclose the record delimiters in double quotes. For example, to use the record delimiter `\r\n`, enclose the value in double quotes:`"\r\n"`.

```
BEGIN
  DBMS_CLOUD.EXPORT_DATA(
    credential_name => 'DEF_CRED_NAME',
```

```
    file_uri_list   => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/dept_export',
    query => 'SELECT * FROM DEPT',
    format => JSON_OBJECT('type' value 'json', 'recorddelimiter' value
'"\r\n"' format json));
END;
/
```

In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

For detailed information about the parameters, see EXPORT_DATA Procedure.

For detailed information about the available `format` parameters you can use with `DBMS_CLOUD.EXPORT_DATA`, see DBMS_CLOUD Package Format Options for EXPORT_DATA.

Notes for exporting with `DBMS_CLOUD.EXPORT_DATA`:

- The `query` parameter that you supply can be an advanced query, if required, such as a query that includes joins or subqueries.
- Specify the `format` parameter with the `compression` option to compress the output files.
- Specify the `format` parameter with the `encryption` option to encrypt data while exporting. See Encrypt Data While Exporting to Object Storage for more information.
- When you no longer need the files that you export, use the procedure `DBMS_CLOUD.DELETE_OBJECT` or use native Cloud Object Storage commands to delete the files.

Export Data as CSV to Cloud Object Storage

Shows the steps to export table data from your Autonomous Database to Cloud Object Storage as CSV data by specifying a query.

This export method supports all the Cloud Object Stores supported by Autonomous Database, and you can use an Oracle Cloud Infrastructure resource principal to access your Oracle Cloud Infrastructure Object Store, Amazon Resource Names (ARNs) to access AWS Simple Storage Service (S3), an Azure service principal to access Azure BLOB storage or Azure Data Lake Storage, or a Google service account to access Google Cloud Platform (GCP) resources.

1. Connect to your Autonomous Database instance.

   See Connect to Autonomous Database for more information.

2. Store your Cloud Object Storage credential using `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'user1@example.com',
    password => 'password'
  );
END;
/
```

The values you provide for `username` and `password` depend on the Cloud Object Storage service you are using.

Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

3. Run `DBMS_CLOUD.EXPORT_DATA` and specify the `format` parameter `type` with the value `csv` to export the results as CSV files on Cloud Object Storage.

   To generate the CSV output files there are two options for the `file_uri_list` parameter:

   - Set the `file_uri_list` value to the URL for an existing bucket on your Cloud Object Storage.
   - Set the `file_uri_list` value to the URL for an existing bucket on your Cloud Object Storage and include a file name prefix to use when generating the file names for the exported CSV files.

   If you do not include the file name prefix in the `file_uri_list`, `DBMS_CLOUD.EXPORT_DATA` supplies a file name prefix. See File Naming for Text Output (CSV, JSON, Parquet, or XML) for details.

   For example, the following shows `DBMS_CLOUD.EXPORT_DATA` with a file name prefix specified in `file_uri_list`:

```
BEGIN
  DBMS_CLOUD.EXPORT_DATA(
    credential_name => 'DEF_CRED_NAME',
    file_uri_list   => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/dept_export',
    query           => 'SELECT * FROM DEPT',
    format          => JSON_OBJECT('type' value 'csv', 'delimiter' value
'|', 'compression' value 'gzip'));
END;
/
```

   In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

   For detailed information about the parameters, see EXPORT_DATA Procedure.

   For detailed information about the available `format` parameters you can use with `DBMS_CLOUD.EXPORT_DATA`, see DBMS_CLOUD Package Format Options for EXPORT_DATA.

Notes for exporting with `DBMS_CLOUD.EXPORT_DATA`:

- The `query` parameter that you supply can be an advanced query, if required, such as a query that includes joins or subqueries.
- Specify the `format` parameter with the `compression` option to compress the output files.
- Specify the `format` parameter with the `encryption` option to encrypt data while exporting. See Encrypt Data While Exporting to Object Storage for more information.
- When you no longer need the files that you export, use the procedure `DBMS_CLOUD.DELETE_OBJECT` or use native Cloud Object Storage commands to delete the files.

Export Data as Parquet to Cloud Object Storage

Shows the steps to export table data from your Autonomous Database to Cloud Object Storage as Parquet data by specifying a query.

This export method supports all the Cloud Object Stores supported by Autonomous Database, and you can use an Oracle Cloud Infrastructure resource principal to access your Oracle Cloud Infrastructure Object Store, Amazon Resource Names (ARNs) to access AWS Simple Storage Service (S3), an Azure service principal to access Azure BLOB storage or Azure Data Lake Storage, or a Google service account to access Google Cloud Platform (GCP) resources.

1. Connect to your Autonomous Database instance.

   See Connect to Autonomous Database for more information.

2. Store your Cloud Object Storage credential using `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'DEF_CRED_NAME',
       username => 'user1@example.com',
       password => 'password'
     );
   END;
   /
   ```

   The values you provide for `username` and `password` depend on the Cloud Object Storage service you are using.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

3. Run `DBMS_CLOUD.EXPORT_DATA` and specify the `format` parameter `type` with the value `parquet` to export the results as parquet files on Cloud Object Storage.

   To generate the parquet output files there are two options for the `file_uri_list` parameter:

   - Set the `file_uri_list` value to the URL for an existing bucket on your Cloud Object Storage.
   - Set the `file_uri_list` value to the URL for an existing bucket on your Cloud Object Storage and include a file name prefix to use when generating the file names for the exported parquet files.

   If you do not include the file name prefix in the `file_uri_list`, `DBMS_CLOUD.EXPORT_DATA` supplies a file name prefix. See File Naming for Text Output (CSV, JSON, Parquet, or XML) for details.

   For example, the following shows `DBMS_CLOUD.EXPORT_DATA` with a file name prefix specified in `file_uri_list`:

   ```
   BEGIN
     DBMS_CLOUD.EXPORT_DATA(
       credential_name => 'DEF_CRED_NAME',
       file_uri_list   => 'https://objectstorage.us-
   ```

```
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/dept_export',
    query           => 'SELECT * FROM DEPT',
    format          => JSON_OBJECT('type' value 'parquet', 'compression'
value 'snappy'));
END;
/
```

In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

For detailed information about the parameters, see EXPORT_DATA Procedure.

For detailed information about the available `format` parameters you can use with `DBMS_CLOUD.EXPORT_DATA`, see DBMS_CLOUD Package Format Options for EXPORT_DATA.

Notes for exporting with `DBMS_CLOUD.EXPORT_DATA`:

- The `query` parameter that you supply can be an advanced query, if required, such as a query that includes joins or subqueries.
- Specify the `format` parameter with the `compression` option to compress the output files. The default `compression` for `type parquet` is `snappy`.
- When you no longer need the files that you export, use the procedure `DBMS_CLOUD.DELETE_OBJECT` or use native Cloud Object Storage commands to delete the files.
- See DBMS_CLOUD Package Oracle Data Type to Parquet Mapping for details on Oracle Type to Parquet Type mapping.

  The following types are not supported or have limitations on their support for exporting Parquet with `DBMS_CLOUD.EXPORT_DATA`:

Oracle Type	Notes
`BFILE`	Not supported
`BLOB`	Not supported
`DATE`	Supported with the following limitation: `DATE` format supports only date, month and year. Hour, minute and seconds are not supported. See DBMS_CLOUD Package Oracle Data Type to Parquet Mapping for details on NLS format limitations for exporting `DATE` to Parquet.
`INTERVAL DAY TO SECOND`	Supported and is treated as string internally
`INTERVAL YEAR TO MONTH`	Supported and is treated as string internally
`LONG`	Not supported
`LONG RAW`	Not supported

Oracle Type	Notes
NUMBER	Supported with the following limitations: – Can have maximum precision of 38 and scale equal to less than precision. – If no precision and scale is provided for the column NUMBER type, by default precision of 38 and scale of 20 is used. – Negative scale is not supported for NUMBER types.
Object Types	Not supported
TIMESTAMP	Supported with the following limitations: – If there are multiple columns with different precision, highest precision will be taken. – TIMESTAMP WITH TIME ZONE Oracle datatype will use the timestamp only. See DBMS_CLOUD Package Oracle Data Type to Parquet Mapping for details on NLS format limitations for exporting TIMESTAMP to Parquet.

Export Data as XML to Cloud Object Storage

Shows the steps to export table data from your Autonomous Database to Cloud Object Storage as XML data by specifying a query.

This export method supports all the Cloud Object Stores supported by Autonomous Database, and you can use an Oracle Cloud Infrastructure resource principal to access your Oracle Cloud Infrastructure Object Store, Amazon Resource Names (ARNs) to access AWS Simple Storage Service (S3), an Azure service principal to access Azure BLOB storage or Azure Data Lake Storage, or a Google service account to access Google Cloud Platform (GCP) resources.

1. Connect to your Autonomous Database instance.

   See Connect to Autonomous Database for more information.

2. Store your Cloud Object Storage credential using DBMS_CLOUD.CREATE_CREDENTIAL.

   For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'user1@example.com',
    password => 'password'
  );
END;
/
```

   The values you provide for username and password depend on the Cloud Object Storage service you are using.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

3. Run DBMS_CLOUD.EXPORT_DATA and specify the format parameter type with the value xml to export the results as XML files on Cloud Object Storage.

To generate the XML output files there are two options for the `file_uri_list` parameter:

- Set the `file_uri_list` value to the URL for an existing bucket on your Cloud Object Storage.
- Set the `file_uri_list` value to the URL for an existing bucket on your Cloud Object Storage and include a file name prefix to use when generating the file names for the exported JSON.

If you do not include the file name prefix in the `file_uri_list`, `DBMS_CLOUD.EXPORT_DATA` supplies a file name prefix. See File Naming for Text Output (CSV, JSON, Parquet, or XML) for details.

For example, the following shows `DBMS_CLOUD.EXPORT_DATA` with a file name prefix specified in `file_uri_list`:

```
BEGIN
  DBMS_CLOUD.EXPORT_DATA(
    credential_name => 'DEF_CRED_NAME',
    file_uri_list   => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/dept_export',
    query           => 'SELECT * FROM DEPT',
    format          => JSON_OBJECT('type' value 'xml', 'compression' value
'gzip'));
END;
/
```

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

For detailed information about the parameters, see EXPORT_DATA Procedure.

For detailed information about the available `format` parameters you can use with `DBMS_CLOUD.EXPORT_DATA`, see DBMS_CLOUD Package Format Options for EXPORT_DATA.

Notes for exporting with `DBMS_CLOUD.EXPORT_DATA`:

- The `query` parameter that you supply can be an advanced query, if required, such as a query that includes joins or subqueries.
- Specify the `format` parameter with the `compression` option to compress the output files.
- Specify the `format` parameter with the `encryption` option to encrypt data while exporting. See Encrypt Data While Exporting to Object Storage for more information.
- When you no longer need the files that you export, use the procedure `DBMS_CLOUD.DELETE_OBJECT` or use native Cloud Object Storage commands to delete the files.

File Naming for Text Output (CSV, JSON, Parquet, or XML)

Describes the output file naming using `DBMS_CLOUD.EXPORT_DATA` with CSV, JSON, Parquet, or XML text file output.

`DBMS_CLOUD.EXPORT_DATA` performs the query specified with the `query` parameter and sends the results to text files either in the Cloud Object Store bucket or to a directory. The output format depends on the `format` parameter `type` you specify (one of CSV, JSON, Parquet, or XML).

To speed up the procedure and to generate the output as fast as possible, DBMS_CLOUD.EXPORT_DATA divides its work. This means that, depending on system resources, when you run DBMS_CLOUD.EXPORT_DATA the procedure creates multiple output files either in the Cloud Object Store bucket or in the directory.

The format for each generated file is:

[*FileNamePrefix* | *client_info_module_action*]*_sequenceNum_timestamp.format_extension.*[*compression_extension*]

- *FileNamePrefix*: (optional) If a *FileNamePrefix* is supplied, DBMS_CLOUD.EXPORT_DATA uses the file name prefix to generate file names for the results. The *FileNamePrefix* is specified using the text supplied after the bucket or directory name in the file_uri_list parameter value.

  You cannot provide multiple values for the *FileNamePrefix* in the file_uri_list.

- *client_info_module_action*: If a file name prefix is not supplied with the file_uri_list parameter, DBMS_CLOUD.EXPORT_DATA uses the combination of *client_info*, application *module* and *action* as the file name prefix (when this information is available). The procedure obtains these names from the application information for the database session that runs the query. See DBMS_APPLICATION_INFO for information on *client_info*, *module* name, and *action* name.

  If a file name prefix is not supplied with the file_uri_list and the database session attributes are not available, DBMS_CLOUD.EXPORT_DATA uses the file name prefix "data".

- *sequenceNum*: The sequence number associated with the DBMS_CLOUD.EXPORT_DATA query. Depending on the query, the database service, and the number of ECPUs (OCPUs if your database uses OCPUs) there are one or more *sequenceNum*s. Also, depending on the size of the results, there are one or more output files for each *sequenceNum*.

  See Manage Concurrency and Priorities on Autonomous Database for information on database services.

- *timestamp*: Timestamp when the file is uploaded.

- *format_extension*: The default value depends on the format type value:

  - CSV format: .csv
  - JSON format: .json
  - PARQUET format .parquet
  - XML format: .xml

  For more information, see the description for format option fileextension in DBMS_CLOUD Package Format Options for EXPORT_DATA.

- *compression_extension*: When you include the format parameter with the compression option with the value gzip, this is "gz".

  When the format type is parquet, the compression value snappy is also supported and is the default.

For example, the file name prefix in the following DBMS_CLOUD.EXPORT_DATA procedure is specified in the file_uri_list parameter, as dept_export. The example generates the output to the provided Object Store in the specified format.

```
BEGIN
  DBMS_CLOUD.EXPORT_DATA(
    credential_name => 'DEF_CRED_NAME',
```

```
    file_uri_list   => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/dept_export',
    query           => 'SELECT * FROM DEPT',
    format          => JSON_OBJECT('type' value 'json'));
END;
/
```

When you specify a file name prefix the generated output files include the file name prefix, similar to the following:

```
dept_export_1_20210809T173033Z.json
dept_export_2_20210809T173034Z.json
dept_export_3_20210809T173041Z.json
dept_export_4_20210809T173035Z.json
```

The number of generated output files depends on the size of the results, the database service, and the number of ECPUs (OCPUs if your database uses OCPUs) in the Autonomous Database instance.

In the following example the `file_uri_list` parameter does not include a file name prefix and the `compression` parameter is supplied, with value `gzip`:

```
BEGIN
  DBMS_CLOUD.EXPORT_DATA(
    credential_name => 'DEF_CRED_NAME',
    file_uri_list   => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/',
    query           => 'SELECT * FROM DEPT',
    format          => json_object('type' value 'json', 'compression' value
'gzip'));
END;
/
```

When a file name prefix is not in the `file_uri_list` parameter, `DBMS_CLOUD.EXPORT_DATA` uses a file name prefix of the form: *client_info_module_action*. For this example the generated output files include the file name prefix that `DBMS_CLOUD.EXPORT_DATA` supplies and the files are compressed with gzip and the file extension `.gz` is added, as follows:

```
Client1_Module1_Action1_1_20210809T173033Z.json.gz
Client1_Module1_Action1_2_20210809T173034Z.json.gz
Client1_Module1_Action1_3_20210809T173041Z.json.gz
Client1_Module1_Action1_4_20210809T173035Z.json.gz
```

If the *client_info_module_action* session information is not available when you run `DBMS_CLOUD.EXPORT_DATA`, the file name prefix is set to `data`. For example:

```
data_1_20210809T173033Z.json.gz
data_2_20210809T173034Z.json.gz
data_3_20210809T173041Z.json.gz
data_4_20210809T173035Z.json.gz
```

For example, the file name prefix in the following `DBMS_CLOUD.EXPORT_DATA` procedure is specified in the `file_uri_list` parameter, as `dept_export`. The example generates the output to the provided directory in the specified format.

```
BEGIN
  DBMS_CLOUD.EXPORT_DATA(
    file_uri_list   => 'DATA_PUMP_DIR:sales.json',
    query           => 'SELECT * FROM SALES',
    format          => JSON_OBJECT('type' value 'json'));
END;
/
```

When you specify a file name prefix the generated output file include the file name prefix, similar to the following:

```
sales_1_20230705T124523275915Z.csv
```

Notes for file naming with `DBMS_CLOUD.EXPORT_DATA`:

- `DBMS_CLOUD.EXPORT_DATA` does not create buckets or directories.
- The number of files that `DBMS_CLOUD.EXPORT_DATA` generates is determined by the number of ECPUs (OCPUs if your database uses OCPUs), the database service, and the size of the result data.
- The following applies when providing a directory object name in the `file_uri_list` parameter:
  - The provided directory must exist and you must have `WRITE` access to the directory.
  - The directory name is case-sensitive when enclosed in double quotes.
  - The credential name parameter must not be provided.
- For CSV, JSON, or XML output, by default when a generated file contains 10MB of data a new output file is created. However, if you have less than 10MB of result data you may have multiple output files, depending on the database service and the number of ECPUs (OCPUs if your database uses OCPUs) for the Autonomous Database instance.

  The default output file chunk size is 10MB for CSV, JSON, or XML. You can change this value with the `format` parameter `maxfilesize` option. See DBMS_CLOUD Package Format Options for EXPORT_DATA for more information.
- For Parquet output, each generated file is less than 128MB and multiple output files may be generated. However, if you have less than 128MB of result data, you may have multiple output files depending on the database service and the number of ECPUs (OCPUs if your database uses OCPUs) for the Autonomous Database instance.

  The `format` parameter `maxfilesize` option does not apply for Parquet files.

Export Data to a Directory

Use `DBMS_CLOUD.EXPORT_DATA` to export files to a directory.

The directory you export files to can be in the Autonomous Database file system or an attached Network File System. See Creating and Managing Directories on Autonomous Database for more information.

- Export Data as CSV to a Directory
  Shows the steps to export table data from your Autonomous Database to a directory as CSV data by specifying a query.
- Export Data as JSON a Directory
  Shows the steps to export table data from your Autonomous Database to a directory as JSON data by specifying a query.
- Export Data as Parquet to a Directory
  Shows the steps to export table data from your Autonomous Database to a directory as Parquet data by specifying a query.
- Export Data as XML to a Directory
  Shows the steps to export table data from your Autonomous Database to Directory as XML data by specifying a query.
- Export Data to a Directory as Oracle Data Pump Files
  You can export data to a directory as Oracle Data Pump dump files by specifying a query.

Export Data as CSV to a Directory

Shows the steps to export table data from your Autonomous Database to a directory as CSV data by specifying a query.

1. Connect to your Autonomous Database instance.

   See Connect to Autonomous Database for more information.

2. Create a directory.

   For example:

   ```
   CREATE DIRECTORY export_dir AS 'export_dir';
   ```

   See Create Directory in Autonomous Database for more information.

3. Run `DBMS_CLOUD.EXPORT_DATA` and specify the `format` parameter `type` with the value `json` to export the results as CSV files to a directory. Do not include the `credential` parameter when sending output to a directory.

   ```
   BEGIN
    DBMS_CLOUD.EXPORT_DATA(
       file_uri_list => 'export_dir:sales.csv',
       format => JSON_OBJECT('type' value 'csv'),
       query => 'SELECT * FROM sales'
    );
   END;
   /
   ```

   When record delimiters include escape characters, such as \r\n or \t, enclose the record delimiters in double quotes. For example, to use the record delimiter `\r\n`, enclose the value in double quotes:`"\r\n"`.

   ```
   BEGIN
     DBMS_CLOUD.EXPORT_DATA(
       file_uri_list => 'export_dir:sales.csv',
       query => 'SELECT * FROM sales',
       format => JSON_OBJECT('type' value 'json', 'recorddelimiter' value
   ```

```
'"\r\n"' format json));
END;
/
```

The directory name is case-sensitive when the directory name is enclosed in double quotes. For example:

```
BEGIN
 DBMS_CLOUD.EXPORT_DATA(
    file_uri_list => '"export_dir":sales.csv',
    format => JSON_OBJECT('type' value 'csv'),
    query => 'SELECT * FROM sales'
 );
END;
/
```

For detailed information about the parameters, see EXPORT_DATA Procedure.

For detailed information about the available `format` parameters you can use with `DBMS_CLOUD.EXPORT_DATA`, see DBMS_CLOUD Package Format Options for EXPORT_DATA.

The parameters are:

- `file_uri_list`: is a comma delimited list of the export file(s). Use of wildcard and substitution characters is not supported in the `file_uri_list`.
- `format`: specifies the required `type` parameter. The valid values are `datapump`, `json`, `xml`, `csv` and `parquet` and it also optionally defines the options you can specify for the export with the `ORACLE_DATAPUMP` Access Driver.
- `query`: specifies a `SELECT` statement so that only the required data is exported. The query determines the contents of the dump file(s).

> **Note:**
>
> The `DBMS_CLOUD.EXPORT_DATA` procedure creates the dump file(s) that you specify in the `file_uri_list`. The procedure does not overwrite files. If a dump file in the `file_uri_list` exists, `DBMS_CLOUD.EXPORT_DATA` generates another file with a unique name. `DBMS_CLOUD.EXPORT_DATA` does not create directories.

For detailed information about the parameters, see EXPORT_DATA Procedure.

Notes for exporting with `DBMS_CLOUD.EXPORT_DATA`:

- The `query` parameter that you supply can be an advanced query, if required, such as a query that includes joins or subqueries.
- Specify the `format` parameter with the `compression` option to compress the output files.

Export Data as JSON a Directory

Shows the steps to export table data from your Autonomous Database to a directory as JSON data by specifying a query.

1. Connect to your Autonomous Database instance.

See Connect to Autonomous Database for more information.

2. Create a directory.

   For example:

   ```
   CREATE DIRECTORY export_dir AS 'export_dir';
   ```

   See Create Directory in Autonomous Database for more information.

3. Run `DBMS_CLOUD.EXPORT_DATA` and specify the `format` parameter `type` with the value `json` to export the results as JSON files to a directory. Do not include the `credential` parameter when sending output to a directory.

   ```
   BEGIN
    DBMS_CLOUD.EXPORT_DATA(
       file_uri_list => 'export_dir:sales.dmp',
       format => json_object('type' value 'json'),
       query => 'SELECT * FROM sales'
    );
   END;
   /
   ```

   When record delimiters include escape characters, such as \r\n or \t, enclose the record delimiters in double quotes. For example, to use the record delimiter `\r\n`, enclose the value in double quotes:`"\r\n"`.

   ```
   BEGIN
     DBMS_CLOUD.EXPORT_DATA(
       file_uri_list => 'export_dir:sales.dmp',
       query => 'SELECT * FROM sales',
       format => JSON_OBJECT('type' value 'json', 'recorddelimiter' value
   '"\r\n"' format json));
   END;
   /
   ```

   The directory name is case-sensitive when the directory name is enclosed in double quotes. For example:

   ```
   BEGIN
    DBMS_CLOUD.EXPORT_DATA(
       file_uri_list => '"export_dir":sales.dmp',
       format => json_object('type' value 'json'),
       query => 'SELECT * FROM sales'
    );
   END;
   /
   ```

   For detailed information about the parameters, see EXPORT_DATA Procedure.

   For detailed information about the available `format` parameters you can use with `DBMS_CLOUD.EXPORT_DATA`, see DBMS_CLOUD Package Format Options for EXPORT_DATA.

   The parameters are:

- `file_uri_list`: is a comma delimited list of the export file(s). Use of wildcard and substitution characters is not supported in the `file_uri_list`.
- `format`: specifies the required `type` parameter. The valid values are `datapump`, `json`, `xml`, `csv` and `parquet` and it also optionally defines the options you can specify for the export with the `ORACLE_DATAPUMP` Access Driver.
- `query`: specifies a `SELECT` statement so that only the required data is exported. The query determines the contents of the dump file(s).

> **Note:**
>
> The `DBMS_CLOUD.EXPORT_DATA` procedure creates the dump file(s) that you specify in the `file_uri_list`. The procedure does not overwrite files. If a dump file in the `file_uri_list` exists, `DBMS_CLOUD.EXPORT_DATA` generates another file with a unique name. `DBMS_CLOUD.EXPORT_DATA` does not create directories.

For detailed information about the parameters, see EXPORT_DATA Procedure.

Notes for exporting with `DBMS_CLOUD.EXPORT_DATA`:

- The `query` parameter that you supply can be an advanced query, if required, such as a query that includes joins or subqueries.
- Specify the `format` parameter with the `compression` option to compress the output files.

Export Data as Parquet to a Directory

Shows the steps to export table data from your Autonomous Database to a directory as Parquet data by specifying a query.

1. Connect to your Autonomous Database instance.

   See Connect to Autonomous Database for more information.

2. Create a directory.

   For example:

   ```
   CREATE DIRECTORY export_dir AS 'export_dir';
   ```

   See Create Directory in Autonomous Database for more information.

3. Run `DBMS_CLOUD.EXPORT_DATA` and specify the `format` parameter `type` with the value `json` to export the results as Parquet files to a directory. Do not include the `credential` parameter when sending output to a directory.

   ```
   BEGIN
    DBMS_CLOUD.EXPORT_DATA(
       file_uri_list => 'export_dir:sales.parquet',
       format => JSON_OBJECT('type' value 'parquet'),
       query => 'SELECT * FROM sales'
    );
   END;
   /
   ```

The directory name is case-sensitive when the directory name is enclosed in double quotes. For example:

```
BEGIN
 DBMS_CLOUD.EXPORT_DATA(
    file_uri_list => '"export_dir":sales.parquet',
    format => JSON_OBJECT('type' value 'parquet'),
    query => 'SELECT * FROM sales'
 );
END;
/
```

For detailed information about the parameters, see EXPORT_DATA Procedure.

For detailed information about the available `format` parameters you can use with `DBMS_CLOUD.EXPORT_DATA`, see DBMS_CLOUD Package Format Options for EXPORT_DATA.

The parameters are:

- `file_uri_list`: is a comma delimited list of the export file(s). Use of wildcard and substitution characters is not supported in the `file_uri_list`.
- `format`: specifies the required `type` parameter. The valid values are `datapump`, `json`, `xml`, `csv` and `parquet` and it also optionally defines the options you can specify for the export with the `ORACLE_DATAPUMP` Access Driver.
- `query`: specifies a `SELECT` statement so that only the required data is exported. The query determines the contents of the dump file(s).

For detailed information about the parameters, see EXPORT_DATA Procedure.

Notes for exporting with `DBMS_CLOUD.EXPORT_DATA`:

- The `query` parameter that you supply can be an advanced query, if required, such as a query that includes joins or subqueries.
- Specify the `format` parameter with the `compression` option to compress the output files.
- See DBMS_CLOUD Package Oracle Data Type to Parquet Mapping for details on Oracle Type to Parquet Type mapping.

  The following types are not supported or have limitations on their support for exporting Parquet with `DBMS_CLOUD.EXPORT_DATA`:

Oracle Type	Notes
`BFILE`	Not supported
`BLOB`	Not supported
`DATE`	Supported with the following limitation: `DATE` format supports only date, month and year. Hour, minute and seconds are not supported. See DBMS_CLOUD Package Oracle Data Type to Parquet Mapping for details on NLS format limitations for exporting `DATE` to Parquet.
`INTERVAL DAY TO SECOND`	Supported and is treated as string internally
`INTERVAL YEAR TO MONTH`	Supported and is treated as string internally
`LONG`	Not supported

Oracle Type	Notes
LONG RAW	Not supported
NUMBER	Supported with the following limitations: – Can have maximum precision of 38 and scale equal to less than precision. – If no precision and scale is provided for the column NUMBER type, by default precision of 38 and scale of 20 is used. – Negative scale is not supported for NUMBER types.
Object Types	Not supported
TIMESTAMP	Supported with the following limitations: – If there are multiple columns with different precision, highest precision will be taken. – TIMESTAMP WITH TIME ZONE Oracle datatype will use the timestamp only. See DBMS_CLOUD Package Oracle Data Type to Parquet Mapping for details on NLS format limitations for exporting TIMESTAMP to Parquet.

Export Data as XML to a Directory

Shows the steps to export table data from your Autonomous Database to Directory as XML data by specifying a query.

1. Connect to your Autonomous Database instance.

   See Connect to Autonomous Database for more information.

2. Create a directory.

   For example:

   ```
   CREATE DIRECTORY export_dir AS 'export_dir';
   ```

   See Create Directory in Autonomous Database for more information.

3. Run DBMS_CLOUD.EXPORT_DATA and specify the format parameter type with the value json to export the results as XML files to a directory. Do not include the credential parameter when sending output to a directory.

   ```
   BEGIN
    DBMS_CLOUD.EXPORT_DATA(
       file_uri_list => 'export_dir:sales.csv',
       format => JSON_OBJECT('type' value 'csv'),
       query => 'SELECT * FROM sales'
    );
   END;
   /
   ```

   The directory name is case-sensitive when the directory name is enclosed in double quotes. For example:

   ```
   BEGIN
    DBMS_CLOUD.EXPORT_DATA(
   ```

```
    file_uri_list => '"export_dir":sales.xml',
    format => JSON_OBJECT('type' value 'xml'),
    query => 'SELECT * FROM sales'
 );
END;
/
```

For detailed information about the parameters, see EXPORT_DATA Procedure.

For detailed information about the available `format` parameters you can use with `DBMS_CLOUD.EXPORT_DATA`, see DBMS_CLOUD Package Format Options for EXPORT_DATA.

The parameters are:

- `file_uri_list`: is a comma delimited list of the export file(s). Use of wildcard and substitution characters is not supported in the `file_uri_list`.
- `format`: specifies the required `type` parameter. The valid values are `datapump`, `json`, `xml`, `csv` and `parquet` and it also optionally defines the options you can specify for the export with the `ORACLE_DATAPUMP` Access Driver.
- `query`: specifies a `SELECT` statement so that only the required data is exported. The query determines the contents of the dump file(s).

> **Note:**
>
> The `DBMS_CLOUD.EXPORT_DATA` procedure creates the dump file(s) that you specify in the `file_uri_list`. The procedure does not overwrite files. If a dump file in the `file_uri_list` exists, `DBMS_CLOUD.EXPORT_DATA` generates another file with a unique name. `DBMS_CLOUD.EXPORT_DATA` does not create directories.

For detailed information about the parameters, see EXPORT_DATA Procedure.

Notes for exporting with `DBMS_CLOUD.EXPORT_DATA`:

- The `query` parameter that you supply can be an advanced query, if required, such as a query that includes joins or subqueries.
- Specify the `format` parameter with the `compression` option to compress the output files.

Export Data to a Directory as Oracle Data Pump Files

You can export data to a directory as Oracle Data Pump dump files by specifying a query.

With this export method you use the `DBMS_CLOUD.EXPORT_DATA` procedure to specify a query to select the data to export as a dump file to a directory.

1. Connect to your Autonomous Database instance.
2. Create a directory.

   For example:

   ```
   CREATE DIRECTORY export_dir AS 'export_dir';
   ```

   See Create Directory in Autonomous Database for more information.

3. Export data from Autonomous Database to your directory as Oracle Data Pump dump file(s) with `DBMS_CLOUD.EXPORT_DATA` and specify the `format` parameter `type` as `datapump`. For example:

```
BEGIN
 DBMS_CLOUD.EXPORT_DATA(
    file_uri_list => 'export_dir:sales.dmp',
    format => json_object('type' value 'datapump'),
    query => 'SELECT * FROM sales'
 );
END;
/
```

Example to export data as multiple Data Pump files to a directory:

```
BEGIN
 DBMS_CLOUD.EXPORT_DATA(
    file_uri_list => 'export_dir:sales1.dmp, export_dir:sales2.dmp',
    format => json_object('type' value 'datapump'),
    query => 'SELECT * FROM sales'
 );
END;
/
```

The parameters are:

- `file_uri_list`: is a comma delimited list of the export file(s). Use of wildcard and substitution characters is not supported in the `file_uri_list`.
- `format`: specifies the required `type` parameter. The valid values are `datapump`, `json`, `xml`, `csv` and `parquet` and it also optionally defines the options you can specify for the export with the `ORACLE_DATAPUMP` Access Driver.
- `query`: specifies a `SELECT` statement so that only the required data is exported. The query determines the contents of the dump file(s).

> **Note:**
>
> The `DBMS_CLOUD.EXPORT_DATA` procedure creates the dump file(s) that you specify in the `file_uri_list`.

For detailed information about the parameters, see EXPORT_DATA Procedure.

Notes for exporting data with `DBMS_CLOUD.EXPORT_DATA`:

- The provided directory must exist and you must be logged in as the `ADMIN` user or have `WRITE` access to the directory.
- The procedure does not overwrite files. If a dump file in the `file_uri_list` exists, `DBMS_CLOUD.EXPORT_DATA` reports an error such as:

```
ORA-31641: unable to create dump file  "/u02/exports/123.dmp"
ORA-27038: created file already exists
```

- `DBMS_CLOUD.EXPORT_DATA` does not create directories.
- The directory name is case-sensitive when enclosed in double quotes.
- The number of dump files that `DBMS_CLOUD.EXPORT_DATA` generates is determined when the procedure runs. The number of dump files that are generated depends on the number of file names you provide in the `file_uri_list` parameter, as well as on the number of ECPUs available to the instance, the service level, and the size of the data.

  For example, if you use a 2 ECPU Autonomous Database instance or the `low` service, then a single dump file is exported with no parallelism, even if you provide multiple file names. If you use an 8 ECPU Autonomous Database instance with the `medium` or `high` service, then the jobs can run in parallel and multiple dump files are exported if you provide multiple file names.
- The `query` parameter value that you supply can be an advanced query, if required, such as a query that includes joins or subqueries.
- The dump files you create with `DBMS_CLOUD.EXPORT_DATA` cannot be imported using Oracle Data Pump `impdp`. Depending on the database, you can use these files as follows:
  - On an Autonomous Database, you can use the dump files with the `DBMS_CLOUD` procedures that support the `format` parameter `type` with the value '`datapump`'. You can import the dump files using `DBMS_CLOUD.COPY_DATA` or you can call `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` to create an external table.
  - On any other Oracle Database, such as Oracle Database 19c on-premise, you can import the dump files created with the procedure `DBMS_CLOUD.EXPORT_DATA` using the `ORACLE_DATAPUMP` access driver. See Unloading and Loading Data with the ORACLE_DATAPUMP Access Driver for more information.

Move Data with Data Pump Export to an Autonomous Database Directory

Oracle Data Pump offers very fast bulk data and metadata movement between Autonomous Database and other Oracle databases. To move data from an Autonomous Database to another Oracle database, use Oracle Data Pump to export to a directory on your Autonomous Database.

This export method is supported on all of the object stores that Autonomous Database supports:

- Oracle Cloud Infrastructure Object Storage
- Oracle Cloud Infrastructure Object Storage Classic
- Azure Blob Storage or Azure Data Lake Storage
- Amazon S3

> **Note:**
>
> If you are using Oracle Cloud Infrastructure Object Storage or Oracle Cloud Infrastructure Object Storage Classic, then you can use the alternative methods to export directly to object store. See Move Data with Data Pump Export to Object Store and Move Data to Object Store as Oracle Data Pump Files Using EXPORT_DATA for more information.

This export method includes the following steps:

1. Use Data Pump Export to save a dump file set to a directory on your database.

   See Use Data Pump to Create a Dump File Set on Autonomous Database for details.

2. Move the dump file set from the directory on your database to your cloud object store.

   See Move Dump File Set from Autonomous Database to Your Cloud Object Store for details.

3. Download the dump file set from the cloud object store, run Data Pump Import, and then perform any required clean up such as removing the dump file set from your cloud object store.

   See Download Dump Files, Run Data Pump Import, and Clean Up Object Store for details.

Topics

- Use Data Pump to Create a Dump File Set on Autonomous Database
  Shows the steps to export data from your Autonomous Database to a directory with Oracle Data Pump.
- Move Dump File Set from Autonomous Database to Your Cloud Object Store
  To move the dump file set to your Cloud Object Store, upload the files from the database directory to your Cloud Object Store.

Use Data Pump to Create a Dump File Set on Autonomous Database

Shows the steps to export data from your Autonomous Database to a directory with Oracle Data Pump.

Oracle recommends using the latest Oracle Data Pump version for exporting data from Autonomous Database to other Oracle databases, as it contains enhancements and fixes for a better experience. Download the latest version of Oracle Instant Client and download the Tools Package, which includes Oracle Data Pump, for your platform. See the installation instructions on the platform install download page for the installation steps required after you download Oracle Instant Client and the Tools Package. See Oracle Instant Client Downloads for details.

> **Note:**
>
> Database Actions provides a link for Oracle Instant Client. To access this link from Database Actions, under **Downloads**, click **Download Oracle Instant Client**.

1. Run Data Pump Export with the *dumpfile* parameter set, the *filesize* parameter set to less than 5G, and the *directory* parameter set. For example, the following shows how to export a schema named `SALES` in an Autonomous Database named `DB2022ADB` with 16 ECPUs:

   ```
   expdp sales/password@db2022adb_high
   directory=data_pump_dir
   dumpfile=exp%L.dmp
   parallel=4
   encryption_pwd_prompt=yes
   filesize=1G
   logfile=export.log
   ```

   Notes for Data Pump parameters:

- If during the export with `expdp` you use the `encryption_pwd_prompt=yes` parameter then use `encryption_pwd_prompt=yes` with your import and input the same password at the `impdp` prompt to decrypt the dump files (remember the password you supply with export). The maximum length of the encryption password is 128 bytes.
- The `dumpfile` parameter supports the `%L` and `%l` wildcards in addition to the legacy `%U` and `%u` wildcards. For example, `dumpfile=export%L.dmp`. Use the `%L` or `%l` wildcard for exports from Oracle Database Release 12.2 and higher. This wildcard expands the dumpfile file name into a 3-digit to 10-digit, variable-width incrementing integer, starting at 100 and ending at 2147483646.

For the best export performance use the `HIGH` database service for your export connection and set the `parallel` parameter to one quarter the number of ECPUs (.25 x **ECPU count**). If you are using OCPU compute model, set the parallel parameter to the number of OCPUs (1 x **OCPU count**). For information on which database service name to connect to run Data Pump Export, see Manage Concurrency and Priorities on Autonomous Database.

After the export is finished you can see the generated dump files by running the following query:

```
SELECT * FROM DBMS_CLOUD.LIST_FILES('DATA_PUMP_DIR');
```

For example, the output from this query shows the generated dump files and the export log file:

```
OBJECT_NAME                    BYTES  CHECKSUM                    CREATED
LAST_MODIFIED
---------------------- ---------- --------- ------------------------------
--------------------
exp01.dmp                      12288               12-NOV-19 06.10.47.0 PM GMT       12-
NOV-19...
exp02.dmp                       8192               12-NOV-19 06.10.48.0 PM GMT       12-
NOV-19...
exp03.dmp                    1171456               12-NOV-19 06.10.48.0 PM GMT       12-
NOV-19...
exp04.dmp                     348160               12-NOV-19 06.10.48.0 PM GMT       12-
NOV-19...
export.log                      1663               12-NOV-19 06.10.50.0 PM GMT       12-
NOV-19...
```

2. Move the dump file set to your cloud object store. See Move Dump File Set from Autonomous Database to Your Cloud Object Store for details.

Notes:

- To perform a full export or to export objects that are owned by other users, you need the `DATAPUMP_CLOUD_EXP` role.
- The `DATA_PUMP_DIR` is the only predefined directory. You can specify a different directory as the *directory* argument if you previously created the directory and you have write privileges on the directory. See Create Directory in Autonomous Database for information on creating directories.
- The API you use to move the dump files to Cloud Object Storage has a maximum file transfer size, so make sure you use a *filesize* argument that is less than or equal to the maximum supported size for your Cloud Object Storage service. See PUT_OBJECT Procedure for the Cloud Object Storage Service file transfer size limits.
- For more information on Oracle Data Pump Export see *Oracle Database Utilities*.

Move Dump File Set from Autonomous Database to Your Cloud Object Store

To move the dump file set to your Cloud Object Store, upload the files from the database directory to your Cloud Object Store.

1. Connect to your database.
2. Store your object store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'DEF_CRED_NAME',
       username => 'adb_user@example.com',
       password => 'password'
     );
   END;
   /
   ```

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials you can then use the same credential name.

   See CREATE_CREDENTIAL Procedure for information about the `username` and `password` parameters for different object storage services.
3. Move the dump files from the database to your Cloud Object Store by calling `DBMS_CLOUD.PUT_OBJECT`.

   For example:

   ```
   BEGIN
      DBMS_CLOUD.PUT_OBJECT(credential_name => 'DEF_CRED_NAME',
        object_uri => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
   namespace-string/b/bucketname/o/exp01.dmp',
        directory_name => 'DATA_PUMP_DIR',
   ```

```
      file_name => 'exp01.dmp');
   DBMS_CLOUD.PUT_OBJECT(credential_name => 'DEF_CRED_NAME',
      object_uri => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/exp02.dmp',
      directory_name => 'DATA_PUMP_DIR',
      file_name => 'exp02.dmp');
   DBMS_CLOUD.PUT_OBJECT(credential_name => 'DEF_CRED_NAME',
      object_uri => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/exp03.dmp',
      directory_name => 'DATA_PUMP_DIR',
      file_name => 'exp03.dmp');
   DBMS_CLOUD.PUT_OBJECT(credential_name => 'DEF_CRED_NAME',
      object_uri => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/exp04.dmp',
      directory_name => 'DATA_PUMP_DIR',
      file_name => 'exp04.dmp');
END;
/
```

In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

See PUT_OBJECT Procedure for information on `PUT_OBJECT`.

4. Perform the required steps to use Oracle Data Pump import and clean up. See Download Dump Files, Run Data Pump Import, and Clean Up Object Store for more details.

Move Data with Data Pump Export to Object Store

To move data from Autonomous Databases to other Oracle databases you can use Oracle Data Pump. Using this export method you use Oracle Data Pump to directly export data to your object store. This export method is supported with Oracle Cloud Infrastructure Object Storage and Oracle Cloud Infrastructure Object Storage Classic.

See Move Data with Data Pump Export to an Autonomous Database Directory for details on the alternative export method to use with other supported cloud object stores.

To export data from Autonomous Databases to other Oracle databases, do the following:

1. Use Data Pump Export to export from your database to Object Storage.

   You have two options for setting the credential to access Object Store, depending on your Oracle Data Pump version:

   - Using Oracle Data Pump version 19.9 (or later) use the `CREDENTIAL` parameter to set the credential or set the database property `DEFAULT_CREDENTIAL`, both options are supported. See Use Oracle Data Pump to Export Data to Object Store Using CREDENTIAL Parameter (Version 19.9 or Later) for details.
   - For Oracle Data Pump versions before 19.9 the `CREDENTIAL` parameter is not supported. Use the `DEFAULT_CREDENTIAL` database property to set the credential to use to access your Object Store. See Use Oracle Data Pump to Export Data to Object Store Setting DEFAULT_CREDENTIAL Property for details.

2. Download the dump files from the Object Storage service, run Data Pump Import with the dump files, and perform any required clean up such as removing the dump file set from Object Storage.

See Download Dump Files, Run Data Pump Import, and Clean Up Object Store for details.

Topics

- Use Oracle Data Pump to Export Data to Object Store Using CREDENTIAL Parameter (Version 19.9 or Later)
  Shows the steps to export data from your Autonomous Database to Cloud Object Storage with Oracle Data Pump Export using the `credential` parameter.
- Use Oracle Data Pump to Export Data to Object Store Setting DEFAULT_CREDENTIAL Property
  Shows the steps to export data from your Autonomous Database to Cloud Object Storage with Oracle Data Pump Export and setting the `DEFAULT_CREDENTIAL` property.

Use Oracle Data Pump to Export Data to Object Store Using CREDENTIAL Parameter (Version 19.9 or Later)

Shows the steps to export data from your Autonomous Database to Cloud Object Storage with Oracle Data Pump Export using the `credential` parameter.

Oracle recommends using the latest Oracle Data Pump version for exporting data from Autonomous Database to other Oracle databases, as it contains enhancements and fixes for a better experience. Download the latest version of Oracle Instant Client and download the Tools Package, which includes Oracle Data Pump, for your platform from Oracle Instant Client Downloads. See the installation instructions on the platform install download page for the installation steps required after you download Oracle Instant Client and the Tools Package.

> **Note:**
>
> Database Actions provides a link for Oracle Instant Client. To access this link from Database Actions, under **Downloads**, click **Download Oracle Instant Client**.

If you are using Oracle Data Pump Version 19.9 or later, then you can use the `credential` parameter as shown in these steps. For instructions for using Oracle Data Pump Versions 19.8 and earlier, see Use Oracle Data Pump to Export Data to Object Store Setting DEFAULT_CREDENTIAL Property.

1. Connect to your database.
2. Store your Cloud Object Storage credential using `DBMS_CLOUD.CREATE_CREDENTIAL`. For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'user1@example.com',
    password => 'password'
  );
END;
/
```

The values you provide for `username` and `password` depend on the Cloud Object Storage service you are using.

If you are exporting to Oracle Cloud Infrastructure Object Storage, you can use the Oracle Cloud Infrastructure native URIs or Swift URIs, but the credentials must be auth tokens. See CREATE_CREDENTIAL Procedure for more information.

3. Run Data Pump Export with the `dumpfile` parameter set to the URL for an existing bucket on your Cloud Object Storage, ending with a file name or a file name with a substitution variable, such as `exp%L.dmp`, and with the `credential` parameter set to the name of the credential you created in the previous step. For example:

```
expdp admin/password@db2022adb_high \
     filesize=5GB \
     credential=def_cred_name \
     dumpfile=https://namespace-string.objectstorage.us-ashburn-1.oci.customer-
oci.com/n/namespace-string/b/bucketname/o/exp%L.dmp \
     parallel=16 \
     encryption_pwd_prompt=yes \
     logfile=export.log \
     directory=data_pump_dir
```

Notes for Oracle Data Pump parameters:

- If during the export with `expdp` you used the `encryption_pwd_prompt=yes` parameter then use `encryption_pwd_prompt=yes` and input the same password at the `impdp` prompt that you specified during the export.
- The `dumpfile` parameter supports the `%L` and `%l` wildcards in addition to the legacy `%U` and `%u` wildcards. For example, `dumpfile=export%L.dmp`. Use the `%L` or `%l` wildcard for exports from Oracle Database Release 12.2 and higher. This wildcard expands the dumpfile file name into a 3-digit to 10-digit, variable-width incrementing integer, starting at 100 and ending at 2147483646.
- If your bucket and data do not reside in a commercial (OC1) realm, where OCI Dedicated Endpoints are not supported, use the following URI format for `dumpfile`:

  `https://objectstorage.region.oraclecloud.com/n/namespace-string/b/bucketname/o/filename`

For `dumpfile`, this example uses the recommended URI format using OCI Dedicated Endpoints for commercial realm (OC1). The *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Object Storage Dedicated Endpoints, Regions and Availability Domains, and Understanding Object Storage Namespaces for more information.
The `credential` parameter cannot be an OCI resource principal, Azure service principal, Amazon Resource Name (ARN), or a Google service account. See Accessing Cloud Resources by Configuring Policies and Roles for more information on resource principal based authentication.

For the best export performance use the `HIGH` database service for your export connection and set the `parallel` parameter to one quarter the number of ECPUs (.25 x **ECPU count**). If you are using OCPU compute model, set the parallel parameter to the number of OCPUs (1 x **OCPU count**).

For information on which database service name to connect to run Data Pump Export, see Database Service Names for Autonomous Database.

For the dump file URL format for different Cloud Object Storage services, see DBMS_CLOUD URI Formats.

This example shows the recommended parameters for exporting from Autonomous Database. For these `expdp` parameters, note the following:

- The maximum `filesize` parameter value is 10000MB for Oracle Cloud Infrastructure Object Storage exports.
- The maximum `filesize` parameter value is 20GB for Oracle Cloud Infrastructure Object Storage Classic exports.
- If the specified `filesize` is too large, the export shows the error message:

  ```
  ORA-17500: ODM err:ODM HTTP Request Entity Too Large
  ```

- The `directory` parameter specifies the directory `data_pump_dir` for the specified log file, `export.log`. See Access Log Files for Data Pump Export for more information.

Oracle Data Pump divides each dump file part into smaller chunks for faster uploads. The Oracle Cloud Infrastructure Object Storage console shows multiple files for each dump file part that you export. The size of the actual dump files will be displayed as zero (0) and its related file chunks as 10mb or less. For example:

```
exp100.dmp
exp100.dmp_aaaaaa
exp101.dmp
exp101.dmp_aaaaaa
```

Downloading the zero byte dump file from the Oracle Cloud Infrastructure console or using the Oracle Cloud Infrastructure CLI will not give you the full dump files. To download the full dump files from the Object Store, use a tool that supports Swift such as curl, and provide your user login and Swift auth token. For example:

```
curl -O -v -X GET -u 'user1@example.com:auth_token' \
   https://swiftobjectstorage.us-ashburn-1.oraclecloud.com/v1/namespace-
string/bucketname/exp100.dmp
```

If you import a file with the `DBMS_CLOUD` procedures that support the `format` parameter `type` with the value '`datapump`', you only need to provide the primary file name. The procedures that support the '`datapump`' format type automatically discover and download the chunks.

When you use `DBMS_CLOUD.DELETE_OBJECT`, the procedure automatically discovers and deletes the chunks when the procedure deletes the primary file.

4. Perform the required steps to use Oracle Data Pump import and clean up.

   See Download Dump Files, Run Data Pump Import, and Clean Up Object Store for details.

> **Note:**
>
> To perform a full export or to export objects that are owned by other users, you need the `DATAPUMP_CLOUD_EXP` role.

For detailed information on Oracle Data Pump Export parameters see *Oracle Database Utilities*.

Use Oracle Data Pump to Export Data to Object Store Setting DEFAULT_CREDENTIAL Property

Shows the steps to export data from your Autonomous Database to Cloud Object Storage with Oracle Data Pump Export and setting the `DEFAULT_CREDENTIAL` property.

Oracle recommends using the latest Oracle Data Pump version for exporting data from Autonomous Database to other Oracle databases, as it contains enhancements and fixes for a better experience. Download the latest version of Oracle Instant Client and download the Tools Package, which includes Oracle Data Pump, for your platform from Oracle Instant Client Downloads. See the installation instructions on the platform install download page for the installation steps required after you download Oracle Instant Client and the Tools Package.

> **Note:**
>
> Database Actions provides a link for Oracle Instant Client. To access this link from Database Actions, under **Downloads**, click **Download Oracle Instant Client**.

1. Connect to your Autonomous Database.
2. Store your Cloud Object Storage credential using `DBMS_CLOUD.CREATE_CREDENTIAL`. For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'user1@example.com',
    password => 'password'
  );
END;
/
```

   The values you provide for `username` and `password` depend on the Cloud Object Storage service you are using.

   If you are exporting to Oracle Cloud Infrastructure Object Storage, you can use the Oracle Cloud Infrastructure native URIs or Swift URIs, but the credentials must be auth tokens. See CREATE_CREDENTIAL Procedure for more information.
3. As the ADMIN user, set the credential you defined in step 2 as the default credential for your database. For example:

```
ALTER DATABASE PROPERTY SET DEFAULT_CREDENTIAL = 'ADMIN.DEF_CRED_NAME'
```

   The `DEFAULT_CREDENTIAL` value cannot be an OCI resource principal, Azure service principal, Amazon Resource Name (ARN), or a Google service account. See Configure Policies and Roles to Access Resources for more information on resource principal based authentication.
4. Run Data Pump Export with the `dumpfile` parameter set to the URL for an existing bucket on your Cloud Object Storage (ending with a file name or a file name with a substitution variable, such as `exp%L.dmp`). For example:

```
expdp admin/password@db2022adb_high \
     filesize=5GB \
     dumpfile=default_credential:https://namespace-string.objectstorage.us-
ashburn-1.oci.customer-oci.com/n/namespace-string/b/bucketname/o/exp%L.dmp \
     parallel=16 \
     encryption_pwd_prompt=yes \
     logfile=export.log \
     directory=data_pump_dir
```

Notes for Data Pump parameters:

- If during the export with `expdp` you used the `encryption_pwd_prompt=yes` parameter then use `encryption_pwd_prompt=yes` and input the same password at the `impdp` prompt that you specified during the export.
- The `dumpfile` parameter supports the `%L` and `%l` wildcards in addition to the legacy `%U` and `%u` wildcards. For example, `dumpfile=export%L.dmp`. Use the `%L` or `%l` wildcard for exports from Oracle Database Release 12.2 and higher. This wildcard expands the dumpfile file name into a 3-digit to 10-digit, variable-width incrementing integer, starting at 100 and ending at 2147483646.
- If your bucket and data do not reside in a commercial (OC1) realm, where OCI Dedicated Endpoints are not supported, use the following URI format for `dumpfile`:

  ```
  https://objectstorage.region.oraclecloud.com/n/namespace-string/b/
  bucketname/o/filename
  ```

For `dumpfile`, this example uses the recommended URI format using OCI Dedicated Endpoints for commercial realm (OC1). The *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Object Storage Dedicated Endpoints, Regions and Availability Domains, and Understanding Object Storage Namespaces for more information.
The `default_credential` keyword in the `dumpfile` parameter is required.

For the best export performance use the `HIGH` database service for your export connection and set the `parallel` parameter to one quarter the number of ECPUs (.25 x **ECPU count**). If you are using OCPU compute model, set the parallel parameter to the number of OCPUs (1 x **OCPU count**).

For information on which database service name to connect to run Data Pump Export, see Manage Concurrency and Priorities on Autonomous Database.

For the dump file URL format for different Cloud Object Storage services, see DBMS_CLOUD URI Formats.

This example shows the recommended parameters for exporting from Autonomous Database. For these `expdp` parameters, note the following:

- The maximum `filesize` parameter value is 10000MB for Oracle Cloud Infrastructure Object Storage exports.
- The maximum `filesize` parameter value is 20GB for Oracle Cloud Infrastructure Object Storage Classic exports.
- If the specified `filesize` is too large, the export shows the error message:

  ```
  ORA-17500: ODM err:ODM HTTP Request Entity Too Large
  ```

- The `directory` parameter specifies the directory `data_pump_dir` for the specified log file, `export.log`. See Access Log Files for Data Pump Export for more information.

Oracle Data Pump divides each dump file part into smaller chunks for faster uploads. The Oracle Cloud Infrastructure Object Storage console shows multiple files for each dump file part that you export. The size of the actual dump files will be displayed as zero (0) and its related file chunks as 10mb or less. For example:

```
exp100.dmp
exp100.dmp_aaaaaa
exp101.dmp
exp101.dmp_aaaaaa
```

Downloading the zero byte dump file from the Oracle Cloud Infrastructure console or using the Oracle Cloud Infrastructure CLI will not give you the full dump files. To download the full dump files from the Object Store, use a tool that supports Swift such as curl, and provide your user login and Swift auth token. For example:

```
curl -O -v -X GET -u 'user1@example.com:auth_token' \
   https://swiftobjectstorage.us-ashburn-1.oraclecloud.com/v1/namespace-
string/bucketname/exp100.dmp
```

If you import a file with the `DBMS_CLOUD` procedures that support the `format` parameter `type` with the value '`datapump`', you only need to provide the primary file name. The procedures that support the '`datapump`' format type automatically discover and download the chunks.

When you use `DBMS_CLOUD.DELETE_OBJECT`, the procedure automatically discovers and deletes the chunks when the procedure deletes the primary file.

5. Perform the required steps to use Oracle Data Pump import and clean up.

   See Download Dump Files, Run Data Pump Import, and Clean Up Object Store for details.

> **Note:**
>
> To perform a full export or to export objects that are owned by other users, you need the `DATAPUMP_CLOUD_EXP` role.

For detailed information on Oracle Data Pump Export parameters see *Oracle Database Utilities*.

Move Data to Object Store as Oracle Data Pump Files Using EXPORT_DATA

You can export data to Oracle Data Pump dump files by specifying a query.

With this export method you use the `DBMS_CLOUD.EXPORT_DATA` procedure to specify a query to select the data to export, as follows:

1. Connect to your database.
2. Store your object store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
```

```
    credential_name => 'DEF_CRED_NAME',
    username => 'adb_user@example.com',
    password => 'password'
  );
END;
/
```

Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials you can then use the same credential name.

See CREATE_CREDENTIAL Procedure for information about the `username` and `password` parameters for different object storage services.

3. Export data from Autonomous Database to your Cloud Object Store as Oracle Data Pump dump file(s) by calling `DBMS_CLOUD.EXPORT_DATA` with the `format` parameter `type` set to value `datapump`. For example:

```
BEGIN
 DBMS_CLOUD.EXPORT_DATA(
    credential_name =>'DEF_CRED_NAME',
    file_uri_list =>'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/exp01.dmp',
    format => json_object('type' value 'datapump'),
    query => 'SELECT warehouse_id, quantity FROM inventories'
 );
END;
/
```

The parameters are:

- `credential_name`: is the name of the credential created in the previous step.
- `file_uri_list`: is a comma delimited list of the export file(s). Use of wildcard and substitution characters is not supported in the `file_uri_list`.
- `format`: specifies the required `type` parameter with the value `datapump`, and optionally defines the options you can specify for the export with the `ORACLE_DATAPUMP` Access Driver.
- `query`: specifies a `SELECT` statement so that only the required data is exported. The query determines the contents of the dump file(s).

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

> **Note:**
>
> The `DBMS_CLOUD.EXPORT_DATA` procedure creates the dump file(s) that you specify in the `file_uri_list`. The procedure does not overwrite files. If a dump file in the `file_uri_list` exists, `DBMS_CLOUD.EXPORT_DATA` reports an error. `DBMS_CLOUD.EXPORT_DATA` does not create buckets.

For detailed information about the parameters, see EXPORT_DATA Procedure.

4. Perform the required steps to use Oracle Data Pump import and clean up. See Download Dump Files, Run Data Pump Import, and Clean Up Object Store for more details.

Notes for exporting data with `DBMS_CLOUD.EXPORT_DATA`:

- The dump files you create with `DBMS_CLOUD.EXPORT_DATA` cannot be imported using Oracle Data Pump `impdp`. Depending on the database, you can use these files as follows:
  - On an Autonomous Database, you can use the dump files with the `DBMS_CLOUD` procedures that support the `format` parameter `type` with the value '`datapump`'. You can import the dump files using `DBMS_CLOUD.COPY_DATA` or you can call `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` to create an external table.
  - On any other Oracle Database, such as Oracle Database 19c on-premise, you can import the dump files created with the procedure `DBMS_CLOUD.EXPORT_DATA` using the `ORACLE_DATAPUMP` access driver. See Unloading and Loading Data with the ORACLE_DATAPUMP Access Driver for more information.
- The number of dump files that `DBMS_CLOUD.EXPORT_DATA` generates is determined when the procedure runs. The number of dump files that are generated depends on the number of file names you provide in the `file_uri_list` parameter, as well as on the number of ECPUs available to the instance, the service level, and the size of the data.

  For example, if you use a 2 ECPU Autonomous Database instance or the `low` service, then a single dump file is exported with no parallelism, even if you provide multiple file names. If you use an 8 ECPU Autonomous Database instance with the `medium` or `high` service, then the jobs can run in parallel and multiple dump files are exported if you provide multiple file names.
- The `query` parameter value that you supply can be an advanced query, if required, such as a query that includes joins or subqueries.

Download Dump Files, Run Data Pump Import, and Clean Up Object Store

If required, download the dump files from Cloud Object Store and use Oracle Data Pump Import to import the dump file set to the target database. Then perform any required clean up.

1. Download the dump files from Cloud Object Store.

> **Note:**
>
> This step is not needed if you are importing the data to an Autonomous Data Warehouse database, to an Autonomous Transaction Processing database, or to Autonomous JSON Database.

If you export directly to Object Store using Oracle Data Pump, as shown in Move Data with Data Pump Export to Object Store, then the dump files on Object Store show size 0. Oracle Data Pump divides each dump file part into smaller chunks for faster uploads. The Oracle Cloud Infrastructure Object Storage console shows multiple files for each dump file part that you export. The size of the actual dump files will be displayed as zero (0) and its related file chunks as 10mb or less. For example:

```
exp01.dmp
exp01.dmp_aaaaaa
exp02.dmp
exp02.dmp_aaaaaa
```

Downloading the zero byte dump file from the Oracle Cloud Infrastructure console or using the Oracle Cloud Infrastructure CLI will not give you the full dump files. To download the full dump files from the Object Store, use a tool that supports Swift such as curl, and provide your user login and Swift auth token.

```
curl -O -v -X GET -u 'user1@example.com:auth_token' \
   https://swiftobjectstorage.us-ashburn-1.oraclecloud.com/v1/namespace-
string/bucketname/exp01.dmp
```

The cURL command does not support wildcards or substitution characters in its URL. You need to use multiple cURL commands to download the dump file set from your Object Store. Alternatively, you can use a script that supports substitution characters to download all the dump files from your Object Store in a single command. See How To: Download all files from an export to object store job in Autonomous Database using cURL for an example.

2. Run Data Pump Import to import the dump file set to the target database.

   If you are importing the data to another Autonomous Database, see Import Data Using Oracle Data Pump on Autonomous Database.

3. Perform post import clean up tasks. If you are done importing the dump files to your target database then drop the bucket containing the data or remove the dump files from the Cloud Object Store bucket, and remove the dump files from the location where you downloaded the dump files to run Data Pump Import.

For detailed information on Oracle Data Pump Import parameters see *Oracle Database Utilities*.

Access Log Files for Data Pump Export

The log files for Data Pump Export operations are stored in the directory you specify with the data pump `directory` parameter.

To access the log file you need to move the log file to your Cloud Object Storage using the procedure `DBMS_CLOUD.PUT_OBJECT`. For example, the following PL/SQL block moves the file `export.log` to your Cloud Object Storage:

```
BEGIN
  DBMS_CLOUD.PUT_OBJECT(
    credential_name => 'DEF_CRED_NAME',
    object_uri => 'https://objectstorage.us-ashburn-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/import.log',
    directory_name  => 'DATA_PUMP_DIR',
```

```
    file_name => 'export.log');
END;
/
```

In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

See PUT_OBJECT Procedure for more information.

Oracle GoldenGate Capture for Oracle Autonomous Database

Using Oracle GoldenGate you can capture changes from an Oracle Autonomous Database and replicate to any target database or platform that Oracle GoldenGate supports, including another Oracle Autonomous Database.

Oracle GoldenGate Capture for Oracle Autonomous Database supports the following:

- Replication for different use cases: Report Offloading, Active-Active, Cloud to Cloud, and Cloud to on-premise.
- Inter-region and cross-region replication: Replicate data between different Oracle Cloud data centers around the world
- Replicate between targets: Replicate from an Autonomous Database to any target database or platform that Oracle GoldenGate supports, including to other Oracle Autonomous Database environments.

See Configure Replicat to Apply to an Autonomous Databases for more information.

Data Sharing

Share data on Autonomous Database

- Use Cloud Links for Read Only Data Access on Autonomous Database
  Cloud Links provide a cloud-based method to remotely access read only data on an Autonomous Database instance.
- Use Table Hyperlinks for Read Only Data Access on Autonomous Database
  You can generate and manage Table Hyperlinks to access data on Autonomous Database.
- Share Data with Data Studio
  The Data Share tool enables you to share data and metadata with other users. The Data Share page displays information about the different types of shares available in the Oracle Autonomous Database.

Use Cloud Links for Read Only Data Access on Autonomous Database

Cloud Links provide a cloud-based method to remotely access read only data on an Autonomous Database instance.

- About Cloud Links on Autonomous Database
  With Cloud Links a data owner registers a table or view for remote access for a selected audience as defined by the data owner, and the data is then accessible to those with access granted at registration time. No further actions are required to set up Cloud Links and whoever is supposed to see and access your data is able to discover and work with the data.

- Grant Cloud Links Access for Database Users
  The ADMIN user grants privileges to database users to register data sets. The ADMIN user also grants privileges to database users to access registered data sets.
- Register a Data Set
  Describes the options and steps to register a table or view you own as a registered data set to share with Cloud Links.
- Find Data Sets and Use Cloud Links
  A user who is granted access to read Cloud Links can search for data sets available to an Autonomous Database instance and can access and use registered data sets with their queries.
- Use Cloud Links Consumer Options
  You can set the service name mapping to use to access data from a consumer database and you can enable caching on a data set consumer for the results of a query or for a query fragment that accesses Cloud Link data.
- Monitor and View Cloud Links Information
  Autonomous Database provides views that allow you to monitor and audit Cloud Links.
- Define a Virtual Private Database Policy to Secure a Registered Data Set
- Notes for Cloud Links
  Provides notes and restrictions for using Cloud Links.

About Cloud Links on Autonomous Database

With Cloud Links a data owner registers a table or view for remote access for a selected audience as defined by the data owner, and the data is then accessible to those with access granted at registration time. No further actions are required to set up Cloud Links and whoever is supposed to see and access your data is able to discover and work with the data.

The Cloud Links implementation leverages Oracle Cloud Infrastructure access mechanisms to make data accessible within a specific scope. Scope indicates who can remotely access the data. Scope can be set to various levels, including to the region where the database resides, to individual tenancies, or to compartments. In addition, you can specify that authorization to access a data set is limited to one or more Autonomous Database instances.

By creating one or more cross-region refreshable clones from the source (data set owner's) Autonomous Database instance, you can use Cloud Links to share data across multiple regions.

Cloud Links greatly simplify the sharing of tables or views across Autonomous Database instances, as compared to traditional database linking mechanisms. With Cloud Links you can discover data without requiring a complex database link setup. Autonomous Database provides transparent access using SQL, and implements privilege enforcement with the Cloud Links scope and by granting authorization to individual Autonomous Database instances.

Cloud Links introduce the concepts of regional namespaces and names for data that is made remotely accessible. This is similar to existing Oracle tables where there is a table, for example `"EMP"` that resides in a namespace (schema), for example `"LWARD"`. There can only be one `LWARD.EMP` in your database. Cloud Links provide a similar namespace and name on a regional level, that is not tied to a single database but applies across many Autonomous Database instances as specified with the scope and optionally with database authorization.

For example, you can register a data set under the namespace `FOREST` and for security purposes or for naming convenience, you can provide a namespace and a name other than the original schema and object names. In this example, `TREE_DATA` is the visible name of the registered data set and this name is not required to be the name of the source table. In

addition to the namespace and the name, the `cloud$link` keyword indicates to the database that it must resolve the source as a Cloud Link.

To access a registered data set, include the namespace, the name, and the `cloud$link` keyword in the `FROM` clause of a `SELECT` statement:

```
SELECT county, species, height FROM FOREST.TREE_DATA@cloud$link;
```

Optionally, you can specify that access to a data set from one or more databases is offloaded to a refreshable clone. When a consumer Autonomous Database is listed in a data set's offload list, access to the data set is directed to the refreshable clone.

> **Note:**
>
> Cloud Links provide read only access to remote objects on an Autonomous Database instance. If you want to use database links with other Oracle databases or with non-Oracle databases, or if you want to use your remote data with DML operations, you need to use database links. See Use Database Links with Autonomous Database for more information.

Cloud Links support private and public synonyms. For example, you can define and use a synonym for `FOREST.TREE_DATA@cloud$link`:

```
CREATE SYNONYM S1 for FOREST.TREE_DATA@cloud$link;
SELECT county, species, height FROM S1;

CREATE PUBLIC SYNONYM S2 for FOREST.TREE_DATA@cloud$link;
SELECT * FROM S2;
```

See Overview of Synonyms for more information on synonyms.

Cloud Links Terms

There are several concepts and terms to use when you work with Cloud Links:

- **Registered Data Set (Data Set)**: Identifies a table or view that has been enabled for remote access in an Autonomous Database. A registered data set also indicates who is allowed to access the data set (its scope). Data set registration defines a namespace and a name for use with Cloud Links. After data set registration, these values together specify the Fully Qualified Name (FQN) for remote access, and allow Cloud Links to manage accessibility for the data set.
- **Data Set Owner**: Specify a data set owner to provide a point of contact for questions about a data set.
- **Scope**: Specifies who and from where a user is allowed to access a registered data set. See Data Set Scope, Access Control, and Authorization for more details on scope.
- **OCID** (Oracle Cloud Identifier): Identifies a specific Tenancy, Compartment, or Database. The scope for a registered data set can be expressed in terms of OCIDs. See Resource Identifiers for more information.
- **Data Registration**: With data registration a user makes a table or view available for remote access, subject to the access restrictions imposed by scope, and optionally subject to an additional authorization step. You can allow remote access with Cloud Links on a table or view stored in the database or on data stored in Object Store.

- **Data Discovery**: A registered data set can be discovered using textual queries from the database. The discovery only shows a data set if there is a privilege to access the data set. You can search to find a registered data set by name or by description.
- **Data Describe**: Allows a user to retrieve a description or metadata for a table or view made available as a registered data set.
- **Offload Targets**: Optionally, you can specify one or more offload targets. Offload targets are refreshable clones that provide Cloud Links data sets to specified Autonomous Database instances. By specifying an offload target, you can dedicate an Autonomous Database instance to provide data sets to separate production from development, for performance, to assure security, or for other reasons. See Use Refreshable Clones with Autonomous Database for more information.

Cloud Links Auditing

Any access to a registered data set using Cloud Links is logged for audit purposes. The log includes the tenancy, compartment, or database that accessed the data set, the amount of data accessed, and additional information. The V$CLOUD_LINK_ACCESS_STATS and GV$CLOUD_LINK_ACCESS_STATS Views show audit information.

Data Set Metadata and Audit Views

Each Autonomous Database instance provides views that expose data set metadata and that allow you to monitor and audit data usage. See Monitor and View Cloud Links Information for more information.

- Data Set Scope, Access Control, and Authorization

Data Set Scope, Access Control, and Authorization

Cloud Links leverage Oracle Cloud Infrastructure access mechanisms to make registered data sets accessible, and to protect your data with access restrictions.
Autonomous Database determines the accessibility of registered data sets as follows:

- The ADMIN user specifies a scope for a user that allows the user to register data sets based on the scope provided.
- A user who has been granted privileges to register data sets specifies a scope when they register a data set.
- Optionally, when you register a data set, a user who has been granted authorization privileges can specify that an authorization step required to access a data set. This authorization step is in addition to scope level access authorization.

Data Set Scope

The ADMIN sets a user's scope with `DBMS_CLOUD_LINK_ADMIN.GRANT_REGISTER` to be one of:

- `'MY$REGION'`
- `'MY$TENANCY'`
- `'MY$COMPARTMENT'`

A user's scope is hierarchical, meaning a user who is granted one of these scopes can allow access as follows:

- `MY$REGION`: A user can grant remote data access to other tenancies in the region of the Autonomous Database instance that is registering the data set. This is the least restrictive scope.

- `MY$TENANCY`: A user can grant remote data access to any resource, tenancy, compartment, or database in the tenancy of the Autonomous Database instance that is registering the data set. This scope is more restrictive than `MY$REGION` scope.
- `MY$COMPARTMENT`: A user can grant remote data access to any resource, compartment, or database in the compartment of the Autonomous Database instance that is registering the data set. This is the most restrictive scope you can set for a user with `GRANT_REGISTER`.

Next, the scope you set when you register a data set determines from where users are allowed to access the data set. The `DBMS_CLOUD_LINK.REGISTER scope` is a comma separated list consisting of one or more of the following:

- **Database OCID**: Access to the data set is allowed for the specific Autonomous Database instances identified by OCID.
- **Compartment OCID**: Access to the data set is allowed for databases in the compartments identified by compartment OCID.
- **Tenancy OCID**: Access to the data set is allowed for databases in the tenancies identified by tenancy OCID.
- **Region name**: Access to the data set is allowed for databases in the region identified by the named region. In all cases, Cloud Links access is limited to within a single region and is not cross-region.
- **`MY$COMPARTMENT`**: Access to the data set is allowed for databases in the same compartment as the data set owner.
- **`MY$TENANCY`**: Access to the data set is allowed for databases in the same tenancy as the data set owner.
- **`MY$REGION`**: Access to the data set is allowed for databases in the same region as the data set owner.

The scope values, `MY$REGION`, `MY$TENANCY`, and `MY$COMPARTMENT` are variables that act as convenience macros and resolve to OCIDs.

> **Note:**
>
> The scope you set when you register a data set is only honored when it matches or is more restrictive than the value set with `DBMS_CLOUD_LINK_ADMIN.GRANT_REGISTER`. For example, assume the ADMIN granted the scope `'MY$TENANCY'` with `GRANT_REGISTER`, and the user specifies `'MY$REGION'` when they register a data set with `DBMS_CLOUD_LINK.REGISTER`. In this case they would see an error such as the following:
>
> ```
> ORA-20001: Share privileges are not enabled for current user or it is
> enabled but not for scope MY$REGION
> ```

You can also use an additional access control mechanism based on a `SYS_CONTEXT` value. This mechanism uses the function `DBMS_CLOUD_LINK.GET_DATABASE_ID` that returns an identifier that is available as a `SYS_CONTEXT` value.

When you register a data set with `DBMS_CLOUD_LINK.REGISTER` you can use the `SYS_CONTEXT` value in Oracle Virtual Private Database (VPD) security policies to control database access to further restrict and control what specific data can be accessed by individual Autonomous Database instances.

See Define a Virtual Private Database Policy to Secure a Registered Data Set for more information on using VPD policies.

The database ID value is also available in the V$CLOUD_LINK_ACCESS_STATS and GV$CLOUD_LINK_ACCESS_STATS Views that tracks access statistics and audit information.

See Using Oracle Virtual Private Database to Control Data Access for more information.

Data Set Authorization

When you register a data set, if you have been granted authorization privileges you can specify that database OCID authorization is required to access a data set. To provide database OCID authorization for a data set, use the `DBMS_CLOUD_LINK.GRANT_AUTHORIZATION` procedure to specify the Autonomous Database instances that are authorized to access the data set. Before you run `DBMS_CLOUD_LINK.GRANT_AUTHORIZATION`, the ADMIN must authorize you to run this procedure with `DBMS_CLOUD_LINK_ADMIN.GRANT_AUTHORIZE`.

Data set registration with authorization required specifies database level access for a data set, in addition to the scope access control specified for the data set, as follows:

- Databases that are within the `SCOPE` specified and have been authorized with `DBMS_CLOUD_LINK.GRANT_AUTHORIZATION` can view rows from the data set.
- Any databases that are within the specified `SCOPE` but have not been authorized with `DBMS_CLOUD_LINK.GRANT_AUTHORIZATION` cannot view data set rows. In this case, consumers without authorization see the data set as empty.
- Databases that are not within the `SCOPE` specified see an error when attempting to access the data set.

Grant Cloud Links Access for Database Users

The ADMIN user grants privileges to database users to register data sets. The ADMIN user also grants privileges to database users to access registered data sets.

When the ADMIN user grants register privileges, they provide a scope that specifies the maximum scope that a user can provide when they register a data set (within the scope hierarchy). The valid `scope` values for use with `DBMS_CLOUD_LINK_ADMIN.GRANT_REGISTER` are:

- `'MY$REGION'`
- `'MY$TENANCY'`
- `'MY$COMPARTMENT'`

See Data Set Scope, Access Control, and Authorization for more information.

1. As the ADMIN user, allow a user to register data sets within a specified scope.

```
BEGIN
DBMS_CLOUD_LINK_ADMIN.GRANT_REGISTER(
   username => 'DB_USER1',
   scope    => 'MY$REGION');
END;
/
```

This specifies that the user `DB_USER1`, has privileges to register data sets within a specified scope, `'MY$REGION'`.

A user can query `SYS_CONTEXT('USERENV', 'CLOUD_LINK_REGISTER_ENABLED')` to check if they are enabled for registering data sets.

For example, the following query:

```
SELECT SYS_CONTEXT('USERENV', 'CLOUD_LINK_REGISTER_ENABLED') FROM DUAL;
```

Returns '`YES`' or '`NO`'.

See GRANT_REGISTER Procedure for more information.

2. As the ADMIN user, allow a user to access registered data sets.

```
EXEC DBMS_CLOUD_LINK_ADMIN.GRANT_READ('LWARD');
```

This allows `LWARD` to access registered data sets that are available to the Autonomous Database instance.

A user can query `SYS_CONTEXT('USERENV', 'CLOUD_LINK_READ_ENABLED')` to check if they are enabled for `READ` access to a data set.

For example, the following query:

```
SELECT SYS_CONTEXT('USERENV', 'CLOUD_LINK_READ_ENABLED') FROM DUAL;
```

Returns '`YES`' or '`NO`'.

See GRANT_READ Procedure for more information.

3. During data registration you can set the authorization required parameter to `TRUE`. When authorization required is `TRUE`, the ADMIN user must run `DBMS_CLOUD_LINK_ADMIN.GRANT_AUTHORIZE` to provide authorization to invoke `DBMS_CLOUD_LINK.GRANT_AUTHORIZATION`. Use `DBMS_CLOUD_LINK.GRANT_AUTHORIZATION` to specify authorization details.

   See Register a Data Set with Authorization Required for more information.

4. When the Autonomous Database instance has Database Vault enabled and the table or view to be registered with Cloud Links is protected by a realm, the owner of the table or view must be authorized to the realm as a realm-owner before registration.

```
BEGIN
     DBMS_MACADM.ADD_AUTH_TO_REALM(
           realm_name    => 'myrealm',
           grantee       => 'object_owner',
           auth_option   => dbms_macutl.g_realm_auth_owner);
END;
/
```

If the realm protecting the table or view is a mandatory realm, the Autonomous Database common schema named `C##DATA$SHARE`, which is used internally as the connecting schema, must be authorized to the realm as a realm-participant.

For example:

```
BEGIN
     DBMS_MACADM.ADD_AUTH_TO_REALM(
           realm_name    => 'myrealm',
```

```
            grantee      => 'C##DATA$SHARE',
            auth_option  => dbms_macutl.g_realm_auth_participant);
END;
/
```

See Use Oracle Database Vault with Autonomous Database for more information.

Notes for granting privileges to database users to register data sets:

- To register data sets or see and access remote data sets, you have to have granted the appropriate privilege for either registering with `DBMS_CLOUD_LINK_ADMIN.GRANT_REGISTER` or for reading data sets with `DBMS_CLOUD_LINK_ADMIN.GRANT_READ`.

  This is true for ADMIN user as well; however the ADMIN user can grant privileges to themself.
- The views `DBA_CLOUD_LINK_PRIVS` and `USER_CLOUD_LINK_PRIVS` provide information on user privileges. See Monitor and View Cloud Links Information for more information.
- A user can run the following query to check if they are enabled for authenticating registered, protected data sets:

  ```
  SELECT SYS_CONTEXT('USERENV', 'CLOUD_LINK_AUTH_ENABLED') FROM DUAL;
  ```

- Revoke Cloud Links Access for Database Users
  The ADMIN user can revoke registration to disallow a user from registering data sets for remote access. The ADMIN user can also revoke privileges or database users to access registered data sets.

Revoke Cloud Links Access for Database Users

The ADMIN user can revoke registration to disallow a user from registering data sets for remote access. The ADMIN user can also revoke privileges or database users to access registered data sets.

1. As the ADMIN user, revoke a user's privileges to register data sets.

   For example:

   ```
   EXEC DBMS_CLOUD_LINK_ADMIN.REVOKE_REGISTER('DB_USER1');
   ```

   This revokes the privileges to register data sets for the user, `DB_USER1`.

   Running `DBMS_CLOUD_LINK_ADMIN.REVOKE_REGISTER` does not affect data sets that are already registered. Use `DBMS_CLOUD_LINK.UNREGISTER` to remove access for a registered data set.

   See REVOKE_REGISTER Procedure and UNREGISTER Procedure for more information.
2. As the ADMIN user, revoke access to registered data sets.

   For example:

   ```
   EXEC DBMS_CLOUD_LINK_ADMIN.REVOKE_READ('LWARD');
   ```

   This revokes access to Cloud Links data sets for the user `LWARD`.

   See REVOKE_READ Procedure for more information.

Register a Data Set

Describes the options and steps to register a table or view you own as a registered data set to share with Cloud Links.

- Register or Unregister a Data Set
  You can register a table or view you own as a registered data set. When you want to remove or replace the data set, you need to unregister it. After you register a data set you can changes the values of a data set's attributes.
- Register or Unregister a Data Set in a Different Region
  You can use Cloud Links in multiple regions, where the source region contains the data set's source database and one or more remote regions contain refreshable clones of the source database.
- Register a Data Set with Authorization Required
  Optionally, when you register a data set, in addition to the scope you can specify that database level authorization is required to access a data set.
- Register a Data Set with Offload Targets for Data Set Access
  Optionally, when you register a data set you can offload access to the data set to one or more Autonomous Database instances that are refreshable clones.
- Update Data Set Registration Attributes
  After you register a data set you can update some data set attributes. You cannot update the schema name, schema object, namespace, or name attributes.

Register or Unregister a Data Set

You can register a table or view you own as a registered data set. When you want to remove or replace the data set, you need to unregister it. After you register a data set you can changes the values of a data set's attributes.

Data set registration defines a namespace and a name for use with Cloud Links. After data set registration, these values together specify the Fully Qualified Name (FQN) for remote access, and allow Cloud Links to manage accessibility for the data set.

To register a data set:

1. Obtain grant register privileges from the ADMIN user.

   See Grant Cloud Links Access for Database Users for more information.

2. Register one ore more data sets.

   For example, assuming there is a schema `CLOUDLINK` on your Autonomous Database instance, you can register two objects, `SALES_VIEW_AGG` and `SALES_ALL` and provide different `scope` parameters to determine how the objects are accessed.

```
BEGIN
   DBMS_CLOUD_LINK.REGISTER(
    schema_name    => 'CLOUDLINK',
    schema_object  => 'SALES_VIEW_AGG',
    namespace      => 'REGIONAL_SALES',
    name           => 'SALES_AGG',
    description    => 'Aggregated regional sales information.',
    scope          => 'MY$TENANCY',
    auth_required  =>  FALSE,
    data_set_owner =>  'amit@example.com' );
```

```
END;
/


BEGIN
   DBMS_CLOUD_LINK.REGISTER(
    schema_name    => 'CLOUDLINK',
    schema_object  => 'SALES_ALL',
    namespace      => 'TRUSTED_COMPARTMENT',
    name           => 'SALES',
    description    => 'Trusted Compartment, only accessible within my
compartment. Early sales data.',
    scope          => 'MY$COMPARTMENT',
    auth_required  =>  FALSE,
    data_set_owner =>  'amit@example.com' );
END;
/
```

The parameters are:

- `schema_name`: is the schema name (the object owner).
- `schema_object`: is the object's name. Valid objects are:
  - Tables (including Heap, External or Hybrid)
  - Views
  - Materialized Views
  - Cloud Tables (see Use Cloud Tables to Store Logging and Diagnostic Information for more information.)

> **Note:**
>
> Other objects such as analytic views or synonyms are not supported.

- `namespace`: is the namespace you provide as a name for access with Cloud Links (one part of the Cloud Link FQN).

  A `NULL` value specifies a system-generated `namespace` value, unique to the Autonomous Database instance.
- `name`: is the name you provide for access with Cloud Links (one part of the Cloud Link FQN).
- `description`: Specifies text to describe the data.
- `scope`: Specifies the scope. You can use one of the variables: `MY$REGION`, `MY$TENANCY`, or `MY$COMPARTMENT` or specify one or more OCIDs.

  See Data Set Scope, Access Control, and Authorization for more information.
- `auth_required`: A Boolean value that specifies if database level authorization is required for the data set, in addition to the scope access control. When this is set to `TRUE`, the data set enforces an additional authorization step. See Register a Data Set with Authorization Required for more information.

- `data_set_owner`: Text value specifies information about the individual responsible for the data set or a contact for questions about the data set. For example, you could supply an email address for the data set owner.

See REGISTER Procedure for more information.

For this example, after the registration completes the scope for the two registered objects are different, based on the scope parameter you supplied with `DBMS_CLOUD_LINK.REGISTER`:

- `MY$TENANCY`: Specifies tenancy level scope for `REGIONAL_SALES.SALES_AGG`.
- `MY$COMPARTMENT`: Specifies the more restrictive compartment level scope within the tenancy for `TRUSTED_COMPARTMENT.SALES`.

You can update some of the values for a data set's attributes after you register a data set. See Update Data Set Registration Attributes for more information.

If you want to revoke remote access to a registered data set, unregister the data set.

For example:

```
BEGIN
   DBMS_CLOUD_LINK.UNREGISTER(
    namespace      => 'TRUSTED_COMPARTMENT',
    name           => 'SALES');
END;
/
```

See UNREGISTER Procedure for more information.

- Notes for Registering or Unregistering a Data Set
  Provides notes for registering a data set with `DBMS_CLOUD_LINK.REGISTER` and unregistering a data set with `DBMS_CLOUD_LINK.UNREGISTER`.

Notes for Registering or Unregistering a Data Set

Provides notes for registering a data set with `DBMS_CLOUD_LINK.REGISTER` and unregistering a data set with `DBMS_CLOUD_LINK.UNREGISTER`.

- After you register an object, users may need to wait up to ten (10) minutes to access the object with Cloud Links.
- When you register a data set and you want consumers in a remote regions to access the data set, you must perform additional steps to make the data set available in a remote region. See Register or Unregister a Data Set in a Different Region for more information.
- Use the procedure `DBMS_CLOUD_LINK.UPDATE_REGISTRATION` to change the attributes for an existing data set.

  The wait time for the update to complete can be up to 10 minutes for a registration change to be propagated and accessible through Cloud Links. This delay can impact the accuracy of the data in the both the `DBA_CLOUD_LINK_REGISTRATIONS` and `DBA_CLOUD_LINK_ACCESS` views.
- You can register a table or view that resides in another user's schema when the you have `READ WITH GRANT OPTION` privileges for the table or view.
- Autonomous Database does not perform hierarchical validity checks at registration time and registrations that are outside the scope will never be visible or accessible.

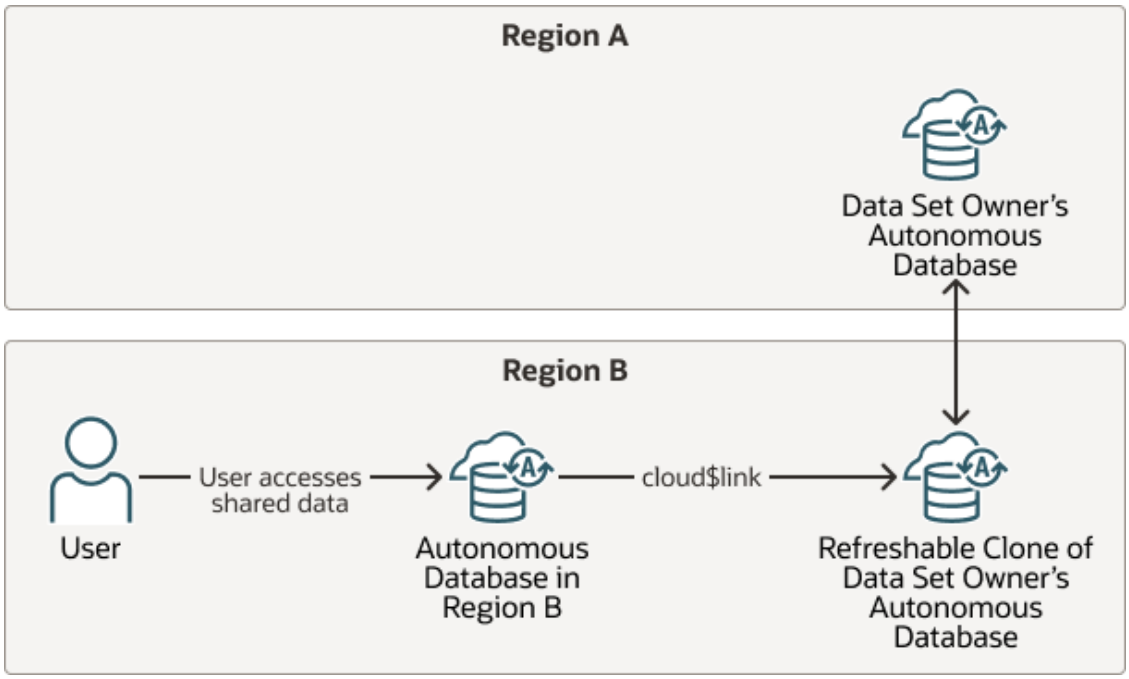
  For example, consider the following sequence:

1. A user with scope `MY$COMPARTMENT` registers an object with a scope that specifies an individual database OCID.
2. When a user requests access to the registered data set, Autonomous Database performs the check to see that the database OCID of the database where the request originates is in the OCID list specified with the `scope` when the data set was registered.
3. After this, the `namespace.name` object will be discoverable, visible, and usable in the database where the request originated.

- `DBMS_CLOUD_LINK.UNREGISTER` may take up to ten (10) minutes to fully propagate, after which the data can longer be accessed remotely.

Register or Unregister a Data Set in a Different Region

You can use Cloud Links in multiple regions, where the source region contains the data set's source database and one or more remote regions contain refreshable clones of the source database.



To use Cloud Links with a data set in a different region:

1. Create a cross-region Refreshable Clone of the source database in a remote region.

   See Create a Cross Tenancy or Cross-Region Refreshable Clone for information on adding a cross-region refreshable clone.
2. On the source database, register the data set.

   See Register or Unregister a Data Set for more information.
3. Refresh the refreshable clone.

   See Refresh a Refreshable Clone on Autonomous Database for more information.
4. On the remote Refreshable Clone register the data set using the same arguments you used to register the data set in the source region.

For example, assuming there is a schema `CLOUDLINK` on your Autonomous Database instance, after you the register `SALES_ALL` on the source database, register `SALES_ALL` on the refreshable clone:

```
BEGIN
   DBMS_CLOUD_LINK.REGISTER(
    schema_name    => 'CLOUDLINK',
    schema_object  => 'SALES_ALL',
    namespace      => 'TRUSTED_COMPARTMENT',
    name           => 'SALES',
    description    => 'Trusted Compartment, only accessible within my
compartment. Early sales data.',
    scope          => 'MY$COMPARTMENT',
    auth_required  =>  FALSE,
    data_set_owner =>  'amit@example.com' );
END;
/
```

The parameters are:

- `schema_name`: is the schema name (the object owner).
- `schema_object`: is the object's name. Valid objects are:
  - Tables (including Heap, External or Hybrid)
  - Views
  - Materialized Views

  > **Note:**
  >
  > Other objects such as analytic views or synonyms are not supported.

- `namespace`: is the namespace you provide as a name for access with Cloud Links (one part of the Cloud Link FQN).

  A `NULL` value specifies a system-generated `namespace` value, unique to the Autonomous Database instance.
- `name`: is the name you provide for access with Cloud Links (one part of the Cloud Link FQN).
- `description`: Specifies text to describe the data.
- `scope`: Specifies the scope. You can use one of the variables: `MY$REGION`, `MY$TENANCY`, or `MY$COMPARTMENT` or specify one or more OCIDs.

  See Data Set Scope, Access Control, and Authorization for more information.
- `auth_required`: A Boolean value that specifies if database level authorization is required for the data set, in addition to the scope access control. When this is set to `TRUE`, the data set enforces an additional authorization step. See Register a Data Set with Authorization Required for more information.
- `data_set_owner`: Text value specifies information about the individual responsible for the data set or a contact for questions about the data set. For example, you could supply an email address for the data set owner.

See REGISTER Procedure for more information.

After the registration completes on the refreshable clone, the scope for the registered object is `MY$COMPARTMENT`: Specifies the more restrictive compartment level scope for my compartment within my tenancy for `TRUSTED_COMPARTMENT.SALES`.

You can unregister a remote data set only in the remote regions, or in both the remote regions and in the source region:

To unregister a data set in a remote region and disable remote access to the data set:

1. On the refreshable clone, unregister the data set.

   For example:

```
BEGIN
   DBMS_CLOUD_LINK.UNREGISTER(
    namespace       => 'TRUSTED_COMPARTMENT',
    name            => 'SALES');
END;
/
```

2. Refresh the refreshable clone.

   See Refresh a Refreshable Clone on Autonomous Database for more information.

To unregister a data set on the source database and unregister the data set on remote region refreshable clones:

1. On the remote refreshable clone if there is only one, or on multiple refreshable clones in remote regions if there are more than one, unregister the data set.

   For example:

```
BEGIN
   DBMS_CLOUD_LINK.UNREGISTER(
    namespace       => 'TRUSTED_COMPARTMENT',
    name            => 'SALES');
END;
/
```

2. On the source database, unregister the data set.

   See Register or Unregister a Data Set for more information.

3. Refresh the refreshable clones.

   See Refresh a Refreshable Clone on Autonomous Database for more information.

- Notes for Registering or Unregistering a Data Set in a Remote Region
  Provides notes for registering a data set in a remote region.

Notes for Registering or Unregistering a Data Set in a Remote Region

Provides notes for registering a data set in a remote region.

- When you register a data set on a refreshable clone in a remote region, the invocation of `DBMS_CLOUD_LINK.REGISTER` on the remote region clone must use the same parameters with the same values as on the source database, with the exception of the `offload_targets` parameter.

  For example, when you run `DBMS_CLOUD_LINK.REGISTER` with scope set to `MY$COMPARTMENT` on the source Autonomous Database instance, run the procedure again on the cross-region refreshable clone with the same scope parameter value (`MY$COMPARTMENT`).

- If you specify the `offload_targets` parameter with `DBMS_CLOUD_LINK.REGISTER` on source, you should omit this parameter when you register the data set on the refreshable clone.
- After you register an object, users may need to wait up to ten (10) minutes to access the object with Cloud Links.
- The following actions require that you refresh the refreshable clone:
  - If you add a VPD policy to the data set in the source, you must refresh the refreshable clone.
  - If you perform a grant or a revoke for the data set on the source database, you must refresh the refreshable clone.

  See Refresh a Refreshable Clone on Autonomous Database for more information.

Register a Data Set with Authorization Required

Optionally, when you register a data set, in addition to the scope you can specify that database level authorization is required to access a data set.

Compared to the previous example where you set `auth_required` to `FALSE`, in this example you set `auth_required` to `TRUE`. When `auth_required` is `TRUE`, additional steps are required to specify one or more databases from which access to the data set is authorized.

> **Note:**
>
> You can only grant authorization, as shown in these steps, if you are authorized. The ADMIN grants authorization privileges with `DBMS_CLOUD_LINK_ADMIN.GRANT_AUTHORIZE`.

1. Use `DBMS_CLOUD_LINK.REGISTER` to register a data with authorization required.

   Assuming there is a schema `CLOUDLINK` on your Autonomous Database instance and you register the object `SALES_VIEW_AGG` and set `auth_required` to `TRUE`, then in addition to defining the scope, you must preform additional steps to determine how the object is accessed.

```
BEGIN
   DBMS_CLOUD_LINK.REGISTER(
    schema_name    => 'CLOUDLINK',
    schema_object  => 'SALES_VIEW_AGG',
    namespace      => 'REGIONAL_SALES',
    name           => 'SALES_AGG',
    description    => 'Aggregated regional sales information.',
    scope          => 'MY$TENANCY',
    auth_required  =>  TRUE,
    data_set_owner =>  'amit@example.com' );
END;
/
```

   The parameters are:

   - `schema_name`: is the schema name (the object owner).
   - `schema_object`: is the object's name. Valid objects are:
     - Tables (including Heap, External or Hybrid)

  - Views
  - Materialized Views

> **Note:**
>
> Other objects such as analytic views or synonyms are not supported.

- `namespace`: is the namespace you provide as a name for access with Cloud Links (one part of the Cloud Link FQN).

  A `NULL` value specifies a system-generated `namespace` value, unique to the Autonomous Database instance.
- `name`: is the name you provide for access with Cloud Links (one part of the Cloud Link FQN).
- `scope`: Specifies the scope. You can use one of the variables: `MY$REGION`, `MY$TENANCY`, or `MY$COMPARTMENT` or specify one or more OCIDs.

  See Data Set Scope, Access Control, and Authorization for more information.
- `auth_required`: A Boolean value that specifies if database level authorization is required for the data set, in addition to the scope access control. When this is set to `TRUE`, the data set enforces an additional authorization step.
- `data_set_owner`: Text value specifies information about the individual responsible for the data set or a contact for questions about the data set. For example, you could supply an email address for the data set owner.

2. Obtain the database ID for the database you want to grant authorization to (to allow the database to access to the data set).

   On the system you want to grant access the data set:

   ```
   SELECT DBMS_CLOUD_LINK.GET_DATABASE_ID FROM DUAL:
   ```

3. Use the database ID you obtained to grant authorization to a specified data set.

   You can only grant authorization, and run `DBMS_CLOUD_LINK.GRANT_AUTHORIZATION`, as shown in this step, if you are authorized. The ADMIN grants authorization with `DBMS_CLOUD_LINK_ADMIN.GRANT_AUTHORIZE`.

   ```
   BEGIN
      DBMS_CLOUD_LINK.GRANT_AUTHORIZATION(
       database_id    => '120xxxxxxx8506029999',
       namespace      => 'TRUSTED_COMPARTMENT',
       name           => 'SALES');
   END;
   /
   ```

   Perform these steps multiple times if you want to authorize additional databases.

You can update the value of the `auth_required` parameter after you register a data set. See Update Data Set Registration Attributes for more information.

If you want to revoke authorization for a database:

```
BEGIN
   DBMS_CLOUD_LINK.REVOKE_AUTHORIZATION(
    database_id    => '120xxxxxxx8506029999',
    namespace      => 'TRUSTED_COMPARTMENT',
    name           => 'SALES');
END;
/
```

See the following for more information:

- REGISTER Procedure
- UPDATE_REGISTRATION Procedure
- GRANT_AUTHORIZE Procedure
- REVOKE_AUTHORIZATION Procedure

Register a Data Set with Offload Targets for Data Set Access

Optionally, when you register a data set you can offload access to the data set to one or more Autonomous Database instances that are refreshable clones.

Use the optional `offload_targets` parameter with `DBMS_CLOUD_LINK.REGISTER` to specify that access is offloaded to refreshable clones. The source database for each refreshable clone is the Autonomous Database instance where you register the data set (data publisher).

The `offload_targets` value is a JSON document that defines one or more `CLOUD_LINK_DATABASE_ID` and `OFFLOAD_TARGET` key value pairs:

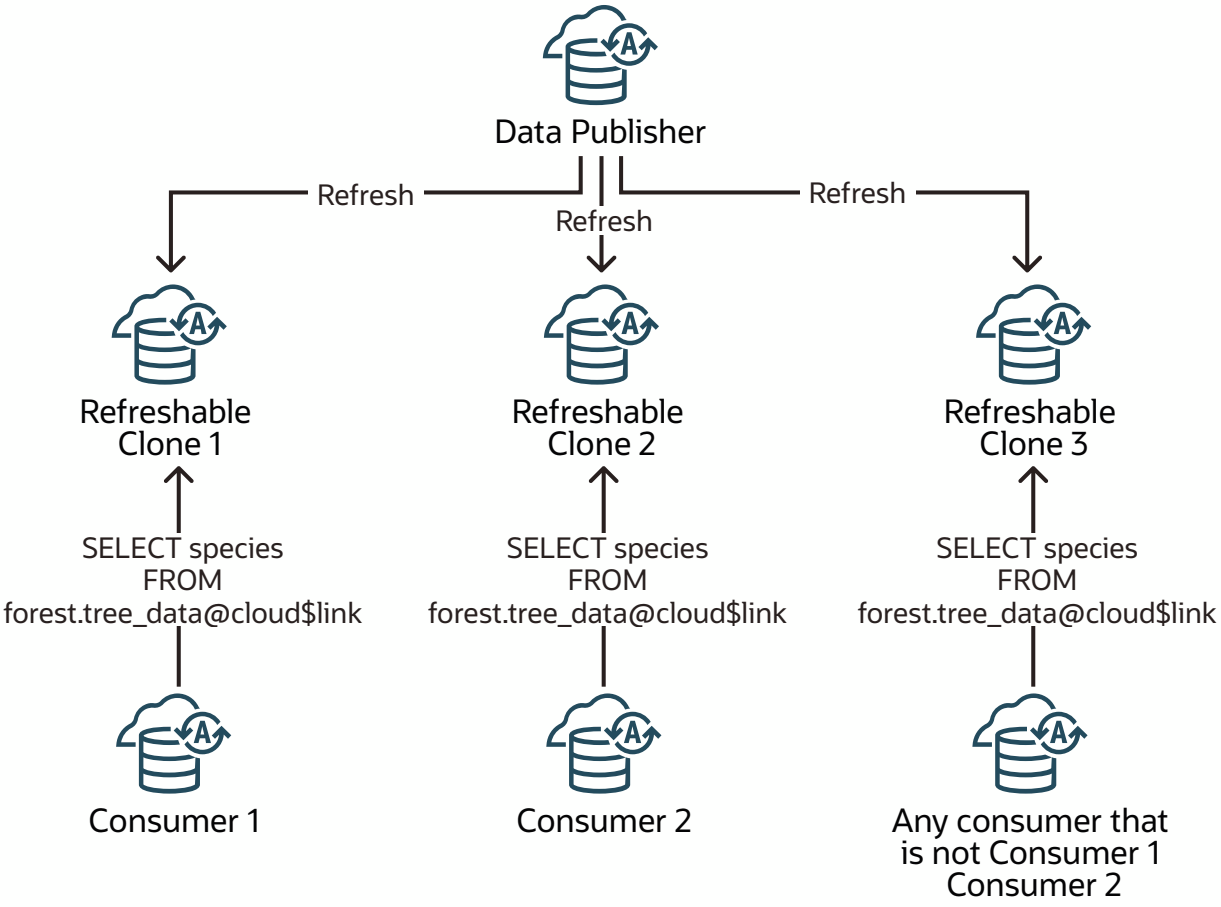
- `CLOUD_LINK_DATABASE_ID` is one of:
  - A database ID: This specifies a database ID for the data set consumer whose request is offloaded to the corresponding refreshable clone specified with the `OFFLOAD_TARGET` value.

    Obtain the database ID by running `DBMS_CLOUD_LINK.GET_DATABASE_ID`. See GET_DATABASE_ID Function for more information.
  - `ANY`: This specifies that any data set consumer's request is offloaded to the corresponding offload target. A consumer's data set request will be routed to the corresponding offload target.

    If you specify `ANY` without specifying database IDs, all data set requests from consumers are offloaded to the refreshable clone specified with the `OFFLOAD_TARGET` value.

    If you specify both database IDs and `ANY`, data set requests from consumers that do not match a database ID are offloaded to the refreshable clone specified with the `OFFLOAD_TARGET` value.
- `OFFLOAD_TARGET` is the OCID for an Autonomous Database instance that is a refreshable clone.

The following figure illustrates using offload targets.



OFFLOAD_TARGETS JSON	
CLOUD LINK DATABASE ID	**OFFLOAD TARGET**
Consumer 1's Database ID	Refreshable Clone 1
Consumer 2's Database ID	Refreshable Clone 2
ANY	Refreshable Clone 3

When a data set consumer requests access to a data set that you register with `offload_targets` and the Autonomous Database instance's database ID matches the value specified in `CLOUD_LINK_DATABASE_ID`, access is offloaded to the refreshable clone identified with `OFFLOAD_TARGET` in the supplied JSON.

For example, the following shows a JSON sample with three `OFFLOAD_TARGET`/ `CLOUD_LINK_DATABASE_ID` value pairs:

```
{
  "OFFLOAD_TARGETS": [
    {
      "CLOUD_LINK_DATABASE_ID": "34xxxxx69708978",
      "OFFLOAD_TARGET":
"ocid1.autonomousdatabase.oc1..xxxxx3pv6wkcr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqs
```

```
tifsfabc"
    },
    {
      "CLOUD_LINK_DATABASE_ID": "34xxxxx89898978",
      "OFFLOAD_TARGET":
"ocid1.autonomousdatabase.oc1..xxxxx3pv6wkcr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqs
tifsfdef"
    },
    {
      "CLOUD_LINK_DATABASE_ID": "34xxxxx4755680",
      "OFFLOAD_TARGET":
"ocid1.autonomousdatabase.oc1..xxxxx3pv6wkcr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqs
tifsfghi"
    }
  ]
}
```

When a data set consumer requests access to a data set that you register with `offload_targets` that includes the `ANY` keyword, access is offloaded to the refreshable clone identified with `OFFLOAD_TARGET` in the supplied JSON (except for requests from consumers that have a matching database ID entry in the supplied JSON).

For example, the following shows a JSON sample with one explicit `OFFLOAD_TARGET`/ `CLOUD_LINK_DATABASE_ID` value pair, and one `ANY` value with a corresponding `OFFLOAD_TARGET`:

```
{
  "OFFLOAD_TARGETS": [
    {
      "CLOUD_LINK_DATABASE_ID": "ANY",
      "OFFLOAD_TARGET":
"ocid1.autonomousdatabase.oc1..xxxxx3pv6wkcr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqs
tifsfdef"
    },
    {
      "CLOUD_LINK_DATABASE_ID": "34xxxxx4755680",
      "OFFLOAD_TARGET":
"ocid1.autonomousdatabase.oc1..xxxxx3pv6wkcr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqs
tifsfghi"
    }
  ]
}
```

To register a data set and specify offload targets, do the following:

1. Obtain the OCID for one or more refreshable clones where you want to offload data set access. Refreshable clone OCIDs are available on the Oracle Cloud Infrastructure Console on a refreshable clone.

> **Note:**
>
> It can take up to 10 minutes after you create a refreshable clone for the refreshable clone to be visible as an offload target. This means you may need to wait up to 10 minutes after you create a refreshable clone for the refreshable clone to be available for Cloud Links offload registration.

2. Obtain the database ID for one or more Autonomous Database instances that you want to access the data set using data provided by the refreshable clone.

   On the system you want to access the data set from a refreshable clone, run the command:

   ```
   SELECT DBMS_CLOUD_LINK.GET_DATABASE_ID FROM DUAL:
   ```

3. Register a data set and specify the `offload_targets` parameter.

   For example, assuming there is a schema `CLOUDLINK` on your Autonomous Database instance, you can register `SALES_VIEW_AGG` and specify the refreshable clones that provide access to the data set:

   ```
   BEGIN
      DBMS_CLOUD_LINK.REGISTER(
       schema_name     => 'CLOUDLINK',
       schema_object   => 'SALES_VIEW_AGG',
       namespace       => 'REGIONAL_SALES',
       name            => 'SALES_AGG',
       description     => 'Aggregated regional sales information.',
       scope           => 'MY$TENANCY',
       auth_required   =>  FALSE,
       data_set_owner  =>  'amit@example.com',
       offload_targets => '{
     "OFFLOAD_TARGETS": [
       {
         "CLOUD_LINK_DATABASE_ID": "34xxxx754755680",
         "OFFLOAD_TARGET":
   "ocid1.autonomousdatabase.oc1..xxxxxaaba3pv6wkcr4jqae5f44n2b2m2yt2j6rx32uzr
   4h25vqstifsfghi"
       }
     ]
   }');
   END;
   /
   ```

   The parameters are:

   - `schema_name`: is the schema name (the object owner).
   - `schema_object`: is the object's name. Valid objects are:
     - Tables (including Heap, External or Hybrid)
     - Views
     - Materialized Views

> **Note:**
>
> Other objects such as analytic views or synonyms are not supported.

- `namespace`: is the namespace you provide as a name for access with Cloud Links (one part of the Cloud Link FQN).

  A `NULL` value specifies a system-generated `namespace` value, unique to the Autonomous Database instance.
- `name`: is the name you provide for access with Cloud Links (one part of the Cloud Link FQN).
- `scope`: Specifies the scope. You can use one of the variables: `MY$REGION`, `MY$TENANCY`, or `MY$COMPARTMENT` or specify one or more OCIDs.

  See Data Set Scope, Access Control, and Authorization for more information.
- `auth_required`: A Boolean value that specifies if database level authorization is required for the data set, in addition to the scope access control. When this is set to `TRUE`, the data set enforces an additional authorization step. See Register a Data Set with Authorization Required for more information.
- `data_set_owner`: Text value specifies information about the individual responsible for the data set or a contact for questions about the data set. For example, you could supply an email address for the data set owner.
- `offload_targets`: Specifies one or more Autonomous Database OCIDs of refreshable clones where access to data sets is offloaded, from the Autonomous Database where the data set is registered.

  For each data set consumer, one of the following can match to offload the request to a refreshable clone:

  – When there is matching the value of the specified `cloud_link_database_id`, that is the values matches the consumer's database ID, access is offloaded to the refreshable clone identified by the OCID specified in `offload_target`.

  – When the `ANY` keyword is specified, if there is not a matching the value of a specified `cloud_link_database_id`, access is offloaded to the refreshable clone identified with the ANY entry by the OCID specified in the corresponding `offload_target`.

See the following for more information:

- REGISTER Procedure
- UPDATE_REGISTRATION Procedure
- Use Refreshable Clones with Autonomous Database

Update Data Set Registration Attributes

After you register a data set you can update some data set attributes. You cannot update the schema name, schema object, namespace, or name attributes.

To update data set attributes:

1. The user who registered a data set can update its attributes with `DBMS_CLOUD_LINK.UPDATE_REGISTRATION`.

If you do not have privileges to update data set attributes, you need to obtain grant register privileges from the ADMIN user.

See Grant Cloud Links Access for Database Users for more information.

2. Update one or more attributes for a data set.

   For example, use `DBMS_CLOUD_LINK.UPDATE_REGISTRATION` to update the `description`, `scope`, and `auth_required` attributes for the data set in the namespace `REGIONAL_SALES` with the name `SALES_AGG`:

```
BEGIN
   DBMS_CLOUD_LINK.UPDATE_REGISTRATION(
    namespace      => 'REGIONAL_SALES',
    name           => 'SALES_AGG',
    description    => 'Updated description for aggregated regional sales
information.',
    scope          => 'MY$COMPARTMENT',
    auth_required  =>  TRUE );
END;
/
```

   The required parameters are:

   - `namespace`: is the data set's namespace you provided when you registered the data set.
   - `name`: is the data set's name you provided when you registered the data set.

   Following is a list of the optional parameters. If `NULL` is passed in for an optional parameter value, the attribute is not modified.

   - `description`: Specifies the updated text to describe the data.
   - `scope`: Specifies the scope. You can use one of the variables: `MY$REGION`, `MY$TENANCY`, or `MY$COMPARTMENT` or specify one or more OCIDs.

     See Data Set Scope, Access Control, and Authorization for more information.
   - `auth_required`: A Boolean value that specifies if database level authorization is required for the data set, in addition to the scope access control. When this is set to `TRUE`, the data set enforces an additional authorization step. See Register a Data Set with Authorization Required for more information.
   - `data_set_owner`: Text value specifies information about the individual responsible for the data set or a contact for questions about the data set. For example, you could supply an email address for the data set owner.
   - `offload_targets`: Specifies one or more Autonomous Database OCIDs of refreshable clones where access to data sets is offloaded, from the Autonomous Database where the data set is registered. See Register a Data Set with Offload Targets for Data Set Access for more information.

   You cannot update the `schema_name` or `schema_object` attributes.

   See UPDATE_REGISTRATION Procedure for more information.

When a data set is registered in one or more cross-region Refreshable Clones, any changes to the registration in the source database should be propagated to the remote regions.

Note the following to propagate changes to cross-region Refreshable Clones:

- If a Producer has N cross-region Refreshable Clones in a region, for example in region A, run `DBMS_CLOUD_LINK.UPDATE_REGISTRATION` on exactly one Refreshable Clone in region A.
- If the same Producer has M cross-region Refreshable Clones in a different remote region, for example in region B, run `DBMS_CLOUD_LINK.UPDATE_REGISTRATION` on exactly one Refreshable Clone in region B.

To update attributes when a data set is registered in one or more cross-region Refreshable Clones:

1. On the source database, update the data set registration.
2. On a remote Refreshable Clone in the remote region if the there is only one remote region, or on a remote Refreshable Clone in each remote region if there are replicated Refreshable Clones in multiple regions, update the data set registration with the same values you used to update the source database, with the exception of the `offload_targets` parameter.

   In any given remote region, you only need to run `DBMS_CLOUD_LINK.UPDATE_REGISTRATION` on exactly one Refreshable Clone in that region (if the region has more than one Refreshable Clone associated with the same data set, you only need to run the procedure once to propagate the changes to all of the Refreshable Clones in an individual remote region).
3. Refresh the refreshable clones.

   See Refresh a Refreshable Clone on Autonomous Database for more information.

Find Data Sets and Use Cloud Links

A user who is granted access to read Cloud Links can search for data sets available to an Autonomous Database instance and can access and use registered data sets with their queries.

After the ADMIN user runs `GRANT_READ`, a user can search for and use Cloud Links.

1. Find the available data sets on your Autonomous Database instance.

   For example, search for all data sets that contain the string, "TREE":

```
DECLARE
   result CLOB DEFAULT NULL;
BEGIN
   DBMS_CLOUD_LINK.FIND('TREE', result);
    DBMS_OUTPUT.PUT_LINE(result);
END;
/

[{"name":"TREE_DATA","namespace":"FOREST","description":"Urban tree data
including county, species and height"}]
```

   The parameters are:

   - `search_string`: The string to search for. The search string is not case sensitive.
   - `search_result`: A JSON document that includes the namespace, name, and description values for the data set.

   See FIND Procedure for more information.

The DBMS_CLOUD_LINK.FIND procedure provides the FQN you can use with Cloud Links. In this case, FOREST.TREE_DATA.

2. Use the view DBA_CLOUD_LINK_ACCESS to list available data sets:

```
SELECT * FROM DBA_CLOUD_LINK_ACCESS;

NAMESPACE            NAME
---------            --------------
FOREST               TREE_DATA
REGIONAL_SALES       SALES_AGG
TRUSTED_COMPARTMENT  SALES
```

3. Use the procedure DBMS_CLOUD_LINK.DESCRIBE to find out more details about a registered data set.

```
SELECT DBMS_CLOUD_LINK.DESCRIBE('FOREST','TREE_DATA') FROM DUAL;

DBMS_CLOUD_LINK.DESCRIBE('FOREST','TREE_DATA')
-------------------------------------------------
Urban tree data including county, species and height
```

4. Use the registered data set in a query.

```
SELECT * FROM FOREST.TREE_DATA@cloud$link;
```

> **Note:**
>
> It can take up to 10 minutes after you register a data set with DBMS_CLOUD_LINK.REGISTER for the data set to be visible and accessible.

Cloud Links support private and public synonyms. For example, you can define and use a synonym for FOREST.TREE_DATA@cloud$link:

```
CREATE SYNONYM S1 for FOREST.TREE_DATA@cloud$link;
CREATE PUBLIC SYNONYM S2 for FOREST.TREE_DATA@cloud$link;

SELECT * FROM S1;

SELECT * FROM S2;
```

See CREATE SYNONYM for more information.

Use Cloud Links Consumer Options

You can set the service name mapping to use to access data from a consumer database and you can enable caching on a data set consumer for the results of a query or for a query fragment that accesses Cloud Link data.

- Set Database Service Name Mapping for Cloud Links Consumers
  You can set the service name mapping to use when Cloud Links consumers access data from a data set owner.

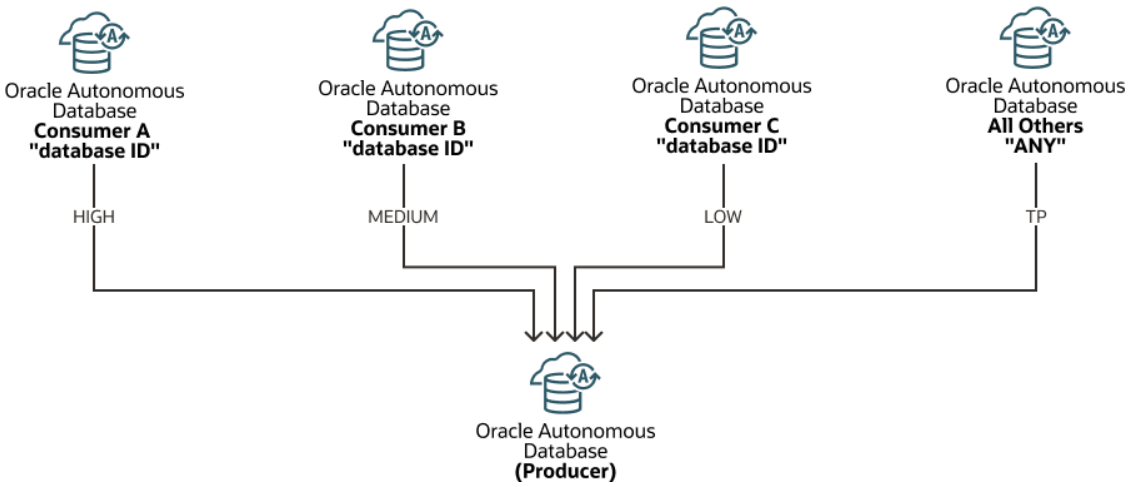
- Enable Caching for a Cloud Link Consumer
  You can enable caching on a data set consumer for the results of a query or for a query fragment that accesses Cloud Link data.

Set Database Service Name Mapping for Cloud Links Consumers

You can set the service name mapping to use when Cloud Links consumers access data from a data set owner.

Cloud Links rely on database resources in the Autonomous Database instance that is the data set producer, or the resources of a refreshable clone, to access shared data. By default remote connectivity for consumers to access Cloud Links data uses the `MEDIUM` database service.

Use `DBMS_CLOUD_LINK_ADMIN.ADD_SERVICE_MAPPING` to set the database service mapping for a consumer. In this procedure you supply either a database ID or the keyword `ANY` to specify the consumer service mapping. For example, the following figure shows a mapping for Consumer A to the HIGH service, Consumer B to the MEDIUM service, Consumer C to the LOW service, and a mapping for ANY to the TP service, meaning all other consumers access Cloud Links using the TP service.



See Database Service Names for Autonomous Database for more information on the characteristics of database services.

Perform the following steps to set the database service to use for a Cloud Links consumer:

1. Obtain the database ID for the database that you want to set a service mapping.

   On the consumer database run `GET_DATABASE_ID` to obtain the consumer's database ID. For example:

   ```
   SELECT DBMS_CLOUD_LINK.GET_DATABASE_ID FROM DUAL:
   ```

   See GET_DATABASE_ID Function for more information.

2. On the data set owner's Autonomous Database instance, specify a service mapping.

   Perform this step on the data set owner's Autonomous Database instance (that is, on the producer database).

   ```
   BEGIN
      DBMS_CLOUD_LINK_ADMIN.ADD_SERVICE_MAPPING(
   ```

```
    database_id      => 'database_id',
    service_name     => 'HIGH');
END;
/
```

Where the *database_id* parameter value is the database ID you obtained in step 1.

Specify the value `ANY` for the `database_id` to use the specified `service_name` with all consumer databases that do not have a `service_name` associated with their `database_id`. That is, any `database_id` whose `service_name` has not been set with `DBMS_CLOUD_LINK_ADMIN.ADD_SERVICE_MAPPING`.

See ADD_SERVICE_MAPPING Procedure for more information.

Only the ADMIN user and schemas with `PDB_DBA` role can run `DBMS_CLOUD_LINK_ADMIN.ADD_SERVICE_MAPPING`.

3. Verify the database IDs and service mapping by listing the Cloud Links service mappings.

   For example:

   ```
   SELECT * FROM  DBA_CLOUD_LINK_SERVICE_MAPPINGS;
   ```

   See DBA_CLOUD_LINK_SERVICE_MAPPINGS View for more information.

Notes for setting and changing service mappings:

- Service mappings take effect at the time connections are established. If a particular consumer's service mappings are changed, the new mappings take effect only for new sessions from the consumer.
- Any service mapping configured in a data set owner for a specific Consumer is honored even if access from the Consumer is offloaded to a Refreshable Clone. The Refreshable Clone must be refreshed to a point in time after the time when service mappings were configured in the data set owner. Note that offload to a Refreshable Clone is configured with the argument `offload_targets` during data set registration.

  See Register a Data Set with Offload Targets for Data Set Access for more information.
- Use the procedure `DBMS_CLOUD_LINK_ADMIN.REMOVE_SERVICE_MAPPING` to remove a service mapping for a specified `database_id`.

  After running `DBMS_CLOUD_LINK_ADMIN.REMOVE_SERVICE_MAPPING` a consumer either uses the default `MEDIUM` database service or the `service_name` you specify if you run `DBMS_CLOUD_LINK_ADMIN.ADD_SERVICE_MAPPING` with the `database_id` value `ANY`. See REMOVE_SERVICE_MAPPING Procedure for more information.
- Set Database Service Name Mapping for Cloud Links Consumers in Remote Region
  A data set that is registered in a source region can be accessed with Cloud Links from a remote region when you create cross-region Refreshable Clone in the remote region.

Set Database Service Name Mapping for Cloud Links Consumers in Remote Region

A data set that is registered in a source region can be accessed with Cloud Links from a remote region when you create cross-region Refreshable Clone in the remote region.

In this case, the service mappings for consumers in the remote region must be added twice, in the source database, and in the refreshable clone in the remote region.

Perform the following steps to set the service mappings for Cloud Links consumers in a remote region.

1. Create a cross-region Refreshable Clone of the source database (the Refreshable Clone is a clone of the Cloud Links data set owner in the remote region).

   See Create a Cross Tenancy or Cross-Region Refreshable Clone for more information.

2. Register the data set.

   See Register a Data Set for more information.

3. Add the service mappings on the data set owner.

   See Set Database Service Name Mapping for Cloud Links Consumers for more information.

4. Refresh the Refreshable Clone.

   See Refresh a Refreshable Clone on Autonomous Database for more information.

5. On the remote Refreshable Clone, register the data set using the same arguments you used to register the data set in the source region (that is, use the same arguments you used in Step 2).

6. On the remote Refreshable Clone, add the service mappings using the same arguments that you used in the source region (that is, use the same arguments you used in Step 3).

When a consumer in the remote region access Cloud Links data, the access will use the same service mappings as you added in source region's data set owner database.

Enable Caching for a Cloud Link Consumer

You can enable caching on a data set consumer for the results of a query or for a query fragment that accesses Cloud Link data.

To enable caching on a data set consumer use the `RESULT_CACHE` hint with the `SHELFLIFE` option. With the `SHELFLIFE` option you supply a value that indicates the duration, in seconds, that a query result is cached. After the `SHELFLIFE` interval has passed, the cached result is marked as invalid. As long as the cached result is valid, a query retrieves the cached data from the cache on the consumer database, which avoids a round trip to the data set owner's database.

Use the `RESULT_CACHE` hint with the `SHELFLIFE` option when the data set is static or when the consumer can tolerate stale results. The value of `SHELFLIFE` allows the Cloud Link data set consumer to control the time, in seconds, when the data in the cache is valid (the acceptable degree of staleness).

If the query result is large and does not fit in memory, you can use the `RESULT_CACHE` hint with the `SHELFLIFE` option and the `TEMP` option to indicate that the result should be written to disk in the temporary tablespace.

To cache Cloud Link data with the `RESULT_CACHE` hint:

1. In a query, specify the `RESULT_CACHE` hint with the `SHELFLIFE` option.

   For example:

   ```
   SELECT /*+ RESULT_CACHE (SHELFLIFE=120) */ * FROM
   FOREST.TREE_DATA@cloud$link;
   ```

   This `RESULT_CACHE` specifies a `SHELFLIFE` value of 120. This indicates that the result should be cached in memory for 120 seconds. After 120 seconds, the cached result is marked as invalid.

The SHELFLIFE value must be a positive integer. The maximum SHELFLIFE value is 4294967295.

2. If the query result is large and does not fit in memory, use both the SHELFLIFE and the TEMP options to indicate that the result should be written to disk in the temporary tablespace.

```
SELECT /*+ RESULT_CACHE (TEMP=true SHELFLIFE=360) */ * FROM
FOREST.TREE_DATA@cloud$link;
```

See RESULT_CACHE_MODE for details on using the result cache with Autonomous Database.

See RESULT_CACHE Hint for more information on RESULT_CACHE with SHELFLIFE.

See DBMS_RESULT_CACHE for information on procedures to manage the result cache and to invalidate objects in the result cache.

Monitor and View Cloud Links Information

Autonomous Database provides views that allow you to monitor and audit Cloud Links.

View	Description
V$CLOUD_LINK_ACCESS_STATS and GV$CLOUD_LINK_ACCESS_STATS Views	Use to track access to each registered data set on the Autonomous Database instance. These views track elapsed time, CPU time, the number of rows retrieved, and additional information about registered data sets. You can use the information in these views to audit Cloud Links data set access and usage.
DBA_CLOUD_LINK_REGISTRATIONS and ALL_CLOUD_LINK_REGISTRATIONS Views	Use to list details of the registered data sets on an Autonomous Database instance.
DBA_CLOUD_LINK_ACCESS and ALL_CLOUD_LINK_ACCESS Views	Use to retrieve details of registered data sets a database is allowed to access.
DBA_CLOUD_LINK_AUTHORIZATIONS View	Provides information about which databases are authorized to access which data sets. This applies to data sets where the auth_required parameter is TRUE.
DBA_CLOUD_LINK_PRIVS and USER_CLOUD_LINK_PRIVS Views	Provides information on the Cloud Link specific privileges, REGISTER, READ, or AUTHORIZE, granted to all users or to the current user.
DBA_CLOUD_LINK_SERVICE_MAPPINGS View	Displays details of all service mappings for Cloud Links consumer databases. Each service mapping consists of a Cloud Link database ID and a database service.

Define a Virtual Private Database Policy to Secure a Registered Data Set

By defining Virtual Private Database (VPD) policies for a registered data set, you can provide fine-grained Cloud Link access control so that only a subset of data (rows) is visible for specific remote sites.

Oracle Virtual Private Database (VPD) is a security feature that lets you control data access dynamically at row level for users and applications by applying filters on the same data set.

Any user who is granted access to read Cloud Links can access and use registered data sets if they are within the scope specified when the data set is registered, and if the extra authorization required parameter is set for the data set, the access is from an authorized

database. Each remote access is done in the context of the remote Autonomous Database instance accessing the registered data set (on the database where the data set was registered).

You use the function `DBMS_CLOUD_LINK.GET_DATABASE_ID` on the remote system to get the database's unique ID. By defining a VPD policy on the database that registered a data set, you now can use the remote database's identifier as `SYS_CONTEXT` rule to provide more fine-grained control. You can define rules for remote databases accessing your registered data set and limit access beyond what is possible by specifying a Cloud Link scope.

Consider an example where `REGIONAL_SALES.SALES_AGG` is made available at the tenancy level. If you want to restrict access to all databases except one specific database, only allowing full access to the specified database, you can add a VPD policy on the registered data set.

For example:

1. Obtain the unique Cloud Link database ID for the Autonomous Database instance where you want to provide full access.

```
DECLARE
     DB_ID NUMBER;
 BEGIN
     DB_ID := DBMS_CLOUD_LINK.GET_DATABASE_ID;
     DBMS_OUTPUT.PUT_LINE('Database ID:' || DB_ID);
 END;
 /
```

2. Create VPD policy on the database that registered the data set by only allowing full access to the one specific database whose identifier you got on the database in question in Step 1.

```
CREATE OR REPLACE FUNCTION vpd_policy_sales(
     owner IN VARCHAR2,
     object IN VARCHAR2)
     RETURN VARCHAR2 IS
BEGIN
  IF SYS_CONTEXT('USERENV', 'CLOUD_LINK_DATABASE_ID')  <>
'12121212163948244901' THEN
    RETURN 'time_id >= trunc(sysdate,''year'')';
  ELSE
    RETURN '';
  END IF;
END;
/
```

See DBMS_RLS for more information.

3. Register the VPD policy for your registered data set to limit full access to only the database identified in step 1.

```
BEGIN
  DBMS_RLS.ADD_POLICY(object_schema => 'CLOUDLINK',
            object_name => 'SALES_VIEW_AGG',
            policy_name => 'THIS_YEAR',
            function_schema => 'ADMIN',
            policy_function => 'VPD_POLICY_SALES',
```

```
            statement_types => 'SELECT',
            policy_type => DBMS_RLS.SHARED_CONTEXT_SENSITIVE);
END;
/
```

See DBMS_RLS for more information.

See Using Oracle Virtual Private Database to Control Data Access for more information.

Notes for Cloud Links

Provides notes and restrictions for using Cloud Links.

- There is a limit of 4096 on the number of data sets you can register.

  Each Autonomous Database instance can register no more than 4096 data sets. This limit applies to every Autonomous Database instance, regardless of the number of ECPUs (OCPUs if your database uses OCPUs) or the storage size of the instance. The limit is a fixed value and setting the **ECPU count** to a larger value does not allow you to register more data sets.
- You can register an object in another schema if you have `READ WITH GRANT OPTION` privilege on the object.
- To register data sets or see and access remote data sets, you have to have granted the appropriate privilege for either registering or reading data sets. This is true for ADMIN as well; however ADMIN can grant this privilege to themself.
- To use either `DBMS_CLOUD_LINK.REGISTER` or `DBMS_CLOUD_LINK.UPDATE_REGISTRATION`, you must have execute privilege on the `DBMS_CLOUD_LINK` package, in addition to the register privilege assigned with `DBMS_CLOUD_LINK_ADMIN.GRANT_REGISTER`. Only the ADMIN user and schemas with `PDB_DBA` have this privilege by default.
- If you drop and re-create a table that was registered, you need to re-register the table for remote access.
- Only the ADMIN user and users with the role `PDB_DBA` are privileged to access the following views:
  - `DBA_CLOUD_LINK_ACCESS`
  - `DBA_CLOUD_LINK_REGISTRATIONS`
  - `DBA_CLOUD_LINK_AUTHORIZATIONS`
  - `DBA_CLOUD_LINK_PRIVS`

  See DBMS_CLOUD_LINK Views for more information.
- Accessing registered, remote data requires the remote database to be open. If the remote database is closed or in restricted mode, you will not be able to access the data and an Oracle error will be returned.
- There is a limit of a maximum of four open database links per session. Exceeding this limit can lead to `ORA-02020` or `ORA-12545`.
- If result cache is enabled, like the default behavior in Autonomous Database with **Data Warehouse** workload, you need to ensure that result cache is not used when realtime data is required.
- When you update your license type from Free to Paid, you must re-register Cloud Links data sets. See Update Always Free Instance to Paid with Autonomous Database for more information.

- Cloud Links remote connectivity by default uses the `MEDIUM` database service. You can change the default with `DBMS_CLOUD_LINK_ADMIN.ADD_SERVICE_MAPPING` by using `ANY` as the value for `DATABASE_ID`. You can change the database service for a consumer with `DBMS_CLOUD_LINK_ADMIN.ADD_SERVICE_MAPPING` by specifying the consumer's `DATABASE_ID`. See Set Database Service Name Mapping for Cloud Links Consumers for more information.

  You can see the remote connections as user `C##DATA$SHARE` in `V$SESSION` and the Cloud Links views V$CLOUD_LINK_ACCESS_STATS and GV$CLOUD_LINK_ACCESS_STATS Views provide more details about remote connections.
- All interfaces are case-sensitive unless otherwise noted, as follows:
  - Anything you enter that exists in the database, for example user names and table names, case is significant and must be entered with uppercase letters.
  - Predefined variables, for example the predefined scope values must be entered in uppercase letters.
  - Anything that you specify for the setup of Cloud Links, for example namespaces or names of tables in a namespace, must be specified as entered. For example, if you define a namespace as `trees`, you must enclose the namespace in double quotes, as `"trees"`, when accessing it with SQL.
- You can share Cloud Links when a data set resides on an Autonomous Database instance in Read-Only mode. See Use Cloud Links from a Read Only Autonomous Database Instance for more information.
- It can take up to 10 minutes after you create a refreshable clone for the refreshable clone to be visible as an offload target. This means you may need to wait up to 10 minutes after you create a refreshable clone for the refreshable clone to be available for Cloud Links offload registration.

  See Register a Data Set with Offload Targets for Data Set Access and Create a Refreshable Clone for an Autonomous Database Instance for more information.
- Use Cloud Links from a Read Only Autonomous Database Instance
  You can share Cloud Links when a data set resides on a Read-Only Autonomous Database instance.

Use Cloud Links from a Read Only Autonomous Database Instance

You can share Cloud Links when a data set resides on a Read-Only Autonomous Database instance.

To share Cloud Links from an Autonomous Database instance that is in Read-Only mode:

1. Change the database mode to Read/Write mode.

   See Change Autonomous Database Operation Mode to Read/Write Read-Only or Restricted for more information.
2. With the database in Read/Write mode, perform the steps to register a data set.
   a. Grant Cloud Links Access for Database Users
   b. Register or Unregister a Data Set
3. After you register one or more data sets, change the database to read-only mode.

   See Change Autonomous Database Operation Mode to Read/Write Read-Only or Restricted for more information.

Use Table Hyperlinks for Read Only Data Access on Autonomous Database

You can generate and manage Table Hyperlinks to access data on Autonomous Database.

Using a Table Hyperlink allows you to easily retrieve data from the database, without requiring you to provide additional information other than the Table Hyperlink to access the data. Any user can access the data by supplying the Table Hyperlink in a browser or using a REST client, subject to security controls enforced by the database. Users can view the retrieved data in tabular format from a browser by appending a query parameter to the Table Hyperlink.

You can read and access data in a Microsoft Excel sheet using a Table Hyperlink through the Oracle Autonomous Database Add-in for Excel, which allows direct retrieval and viewing of tabular data from the Autonomous Database within Excel. See Read and Access Data Using Table Hyperlinks in Excel Sheet for more information.

- About Table Hyperlinks on Autonomous Database
  Depending on how you generate a Table Hyperlink, a Table Hyperlink provides access to data in tables or views or by running a SQL query.
- Generate, List, Invalidate, and Monitor Table Hyperlinks
  Shows you the steps to generate. list, invalidate, or monitor Table Hyperlinks.
- Use Table Hyperlink to Access Data
  Table Hyperlink data is retrieved and returned in JSON format and is paginated.
- Use Table Hyperlinks to Access Data in Table Format
  Table Hyperlink data can be viewed in tabular format when accessed from a browser. The table format supports scrolling, sorting, filtering, cell colors, and Group By where you select a column name to see column data based on the values in the column.
- Use Table Hyperlinks to Create an External Table
  You can create an external table over an Autonomous Database Table Hyperlink.

About Table Hyperlinks on Autonomous Database

Depending on how you generate a Table Hyperlink, a Table Hyperlink provides access to data in tables or views or by running a SQL query.

When you generate a Table Hyperlink you specify an expiration, either as an expiration time, for example set the Table Hyperlink to expire after 120 minutes, or as an expiration count, for example the Table Hyperlink expires after the Table Hyperlink is used 10 times.

Table Hyperlinks provide the following:

- **Public Access**: Using a Table Hyperlink a data recipient on the public internet can access data when the data resides on an Autonomous Database instance in a private subnet.
- **Access with ACL Checking**: You can generate a Table Hyperlink with ACL checking enabled. When you generate a Table Hyperlink and you enable ACL checking, Table Hyperlink data is only accessible to a consumer database after the incoming request's IP address is validated based on the ACLs defined on the producer database.
- **Expiration**: A data provider specifies expiration for a Table Hyperlink, meaning the Table Hyperlink has a limited time before it expires (up to a maximum of 90 Days).
- **Expiration Use Count Limits**: A data provider can specify a limit on how many times a recipient can use a Table Hyperlink to access data.
- **Endpoint Transparency**: A data provider is able to hide the Autonomous Database name so that it is not visible in a Table Hyperlink.

- **Column Level Options for Tabular Format Features and Grouping**: When you generate a Table Hyperlink you can specify the following options by column: sorting, filtering, allow cell colors, or the group by option where you can select to view data grouped by the values in a specified column.

> **Note:**
>
> There is a limit of 128 active Table Hyperlinks on an Autonomous Database instance.

Table Hyperlink Use Cases

Generating and providing a Table Hyperlink supports the following use cases:

Use Case	Description
Within Organization Collaboration	You can use Table Hyperlinks for emergency data access. In situations where a rapid response is needed, such as during a critical incident investigation, you provide a Table Hyperlink to allow immediate and temporary access to specific data without the need to create new database accounts or modify existing permissions.
BB (Business to Business) Applications	A business partner can easily access data. Using a Table Hyperlink, a business can provide a business partner with a simple way to access data or reports. This can eliminate the need for manual report generation and email distribution.
Third-party Audits and Reviews	When an external auditor or reviewer requires access to specific data for a limited time, a Table Hyperlink can give them the access they need without compromising the overall security of the database.
Data as a Product (Digital Commerce)	Vendors can grant limited or single-use access to purchased content or data using a Table Hyperlink. Once accessed, the URL expires, protecting the product's exclusivity and ensuring efficient, secure delivery.

Security Best Practices for Table Hyperlinks

Following are some best practices for generating and using Table Hyperlinks:

- **Set a Short Expiration Time**: A Table Hyperlink should only be valid for the minimum time required. The shorter the validity period, the lower the risk if the Table Hyperlink is compromised.
- **PAR Invalidation**: Invalidate a Table Hyperlink immediately when it is no longer required.
- **Use Appropriate Permissions**: A Table Hyperlink runs using the privileges granted to the database user who generates the Table Hyperlink. The user that generates a Table Hyperlink should have the minimum privileges required for providing access to the data.
- **Content Security**: To mitigate risk of sharing unintended dynamic data:
  - Create a view on top of the data that you want to share in a Table Hyperlink, and monitor that the view definition is up to date.
  - As needed, create a VPD policy when you generate a Table Hyperlink. You can use VPD policies to restrict the rows visible to the Table Hyperlink users.
  - As needed, enable ACL checking when you generate a Table Hyperlink.
- **Load Monitoring**: Monitor the Table Hyperlink query load using PerfHub and SQL monitoring.

Enable **Compute auto scaling** and make sure that the CPU count is appropriately sized for the data set size and Table Hyperlink query load.

Generate, List, Invalidate, and Monitor Table Hyperlinks

Shows you the steps to generate. list, invalidate, or monitor Table Hyperlinks.

- Generate a Table Hyperlink for a Table or a View
  Shows you the steps to generate a Table Hyperlink that you can use to share access for a schema object (table or view).
- Generate a Table Hyperlink with a Select Statement
  Shows you the steps to generate a Table Hyperlink that provides access to data using a SQL query statement.
- Generate a Table Hyperlink with UI Features Specified on Columns
  When you generate a Table Hyperlink you can use the `column_lists` parameter to specify UI features for specified columns.
- List Table Hyperlinks
  You can list the active Table Hyperlinks that you generated on an Autonomous Database instance and the ADMIN user can list all active Table Hyperlinks.
- Invalidate Table Hyperlinks
  At any time a user with appropriate privileges can invalidate a Table Hyperlink.
- Define a Virtual Private Database Policy to Secure Table Hyperlink Data
- Monitor and View Table Hyperlink Usage
  Autonomous Database provides views that allow you to monitor Table Hyperlink usage.

Generate a Table Hyperlink for a Table or a View

Shows you the steps to generate a Table Hyperlink that you can use to share access for a schema object (table or view).

When a Table Hyperlink runs it uses the privileges granted to the database user who generates the Table Hyperlink. The user that generates a Table Hyperlink should have the minimum privileges required for providing access to the data. To maintain security, Oracle recommends that you do not run `DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL` as the ADMIN user.

To use a Table Hyperlink to provide access to data as a schema object (table or view):

1. Identify the table or view that you want to share.

   If there are restrictions on the data you want to make available, use the `application_user_id` parameter when you generate the Table Hyperlink and create a VPD policy to restrict the data that you expose. See Define a Virtual Private Database Policy to Secure Table Hyperlink Data for more information.

2. Run `DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL` to generate the Table Hyperlink.

```
DECLARE
   status CLOB;
   BEGIN
      DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL(
          schema_name => 'ADMIN',
          schema_object_name    => 'TREE_DATA',
          expiration_minutes    => 360,
          service_name          => 'HIGH',
          inherit_acl            => FALSE,
```

```
            result                => status);
        dbms_output.put_line(status);
    END;
/
```

Use the optional `service_name` parameter to specify that the generated Table Hyperlink is serviced with a specific service-level guarantee and resources. For example, use the `service_name` parameter to specify that access to the object is mapped to the HIGH service.

Use the optional `inherit_acl` parameter if you want to limit access to Table Hyperlink data. Set this parameter to `TRUE` to inherit the ACL. When inherit ACL is true, an incoming Table Hyperlink consumer's IP address is validated with the ACLs on the producer database before allowing access to data. If the producer database does not have ACLs configured, the `inherit_acl` value is ignored and data access is allowed without any ACL checks.

See Configuring Network Access with Access Control Rules (ACLs) and GET_PREAUTHENTICATED_URL Procedure for more information.

3. Check the result.

   In this example `status` contains the result that includes information about the Table Hyperlink.

```
{
"status":"SUCCESS","id":"wPY0uxyx-gioxOndiKVlqVF585xqJs14CIp9M1qH",
"preauth_url":"https://dataaccess.adb.us-
ashburn-1.oraclecloudapps.com/adb/p/-pPLdrbUTcLUrZLExyzxyzabc_1w/data",
"expiration_ts":"2024-10-22T22:37:18.805Z",
"expiration_count":null
}
```

Notes for generating a Table Hyperlink with `DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL`:

- When you generate a Table Hyperlink on an Autonomous Database instance with a private endpoint, the result includes a name `private_preauth_url` with the value of the form: `"https://`*`private-endpoint`*`/adb/p/`*`parurl-token`*`/data"`.

  When you generate a Table Hyperlink on an Autonomous Database instance with a private endpoint and the private endpoint is configured with **Allow public access** enabled, the result includes both the `preauth_url` for the public endpoint and `private_preauth_url`.

  See Configure Private Endpoints and Use a Private Endpoint with Public Access Allowed for more information.

- Use `DBMS_DATA_ACCESS.LIST_ACTIVE_URLS` to show Table Hyperlinks. See List Table Hyperlinks for details.
- Use `DBMS_DATA_ACCESS.EXTEND_URL` to extend the life of a Table Hyperlink. See EXTEND_URL Procedure for more information.

Generate a Table Hyperlink with a Select Statement

Shows you the steps to generate a Table Hyperlink that provides access to data using a SQL query statement.

When a Table Hyperlink runs it uses the privileges granted to the database user who generates the Table Hyperlink. The user that generates a Table Hyperlink should have the minimum

privileges required for providing access to the data. To maintain security, Oracle recommends that you do not run `DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL` as the ADMIN user.

To use a Table Hyperlink to provide to access to data as an arbitrary SQL query statement:

1. Identify the table or view that contains the information you want to share, as well as the `SELECT` statement on the table or view that you want to use.

   If there are restrictions on the data you want to make available, use the `application_user_id` parameter when you generate the Table Hyperlink and create a VPD policy to restrict the data that you expose. See Define a Virtual Private Database Policy to Secure Table Hyperlink Data for more information.

2. Run `DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL` to generate the Table Hyperlink.

   You have two options, depending on whether you want to include default bind variable values (if you include bind variables in the select statement).

   a. Provide a select statement and do not use the optional `default_bind_values` parameter.

      For example:

```
DECLARE
   status CLOB;
   BEGIN
      DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL(
          sql_statement        => 'SELECT species, height FROM
TREE_DATA',
          expiration_count     => 10,
          service_name         => 'HIGH',
          inherit_acl            => FALSE,
          result               => status);
       dbms_output.put_line(status);
    END;
/
```

      The `sql_statement` value must be a `SELECT` statement. The `SELECT` statement supports bind variables. If bind variables are included in the select statement, the bind variable values must be appended to the generated Table Hyperlink as a query parameter when accessing the data.

   b. Provide a select statement and include the `default_bind_values` parameter.

      For example:

```
DECLARE
   status CLOB;
   BEGIN
      DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL(
          sql_statement => 'select * FROM TREE_DATA WHERE COUNTY
= :countyNAME',
          default_bind_values  => '{"countyNAME" : "First"}',
          expiration_count     => 10,
          service_name         => 'HIGH',
          inherit_acl            => FALSE,
          result               => status);
       dbms_output.put_line(status);
```

```
    END;
/
```

The `sql_statement` parameter must be a `SELECT` statement. The `SELECT` statement supports bind variables.

The `default_bind_values` parameter specifies values for one or more bind variables in the `sql_statement`. When `default_bind_values` is included with `DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL`:

- For bind variables specified in `default_bind_values`, you can omit the bind variable values when you access the data. The default value is used for a specified bind variable if an override is not provided in the Table Hyperlink URL as a query parameter when you access the data.
- You can override a default bind variable value specified in `default_bind_values` if you append the bind variable value in the Table Hyperlink URL as a query parameter when you access the data.
- If a bind variable is included in the select statement and a default value for the bind variable is not specified in the `default_bind_values` parameter, you must append a bind variable value to the generated Table Hyperlink URL as a query parameter when you access the data.

Bind variable support is available for `NUMBER` and `VARCHAR2` column types.

In these examples the `expiration_count` parameter specifies that the Table Hyperlink expires and is invalidated after 10 uses. When an `expiration_time` is not specified, the expiration time is set to the default value, 90 days.

Use the optional `service_name` parameter to specify that the generated Table Hyperlink is serviced with a specific service-level guarantee and resources. For example, use the `service_name` parameter to specify access to the SQL statement is mapped to the HIGH service.

Use the optional `inherit_acl` parameter if you want to limit access to Table Hyperlink data. Set this parameter to `TRUE` to inherit ACLs. When inherit ACL is true, an incoming Table Hyperlink consumer's IP address is validated with the ACLs on the producer database before allowing access to data. If the producer database does not have ACLs configured, the `inherit_acl` value is ignored and data access is allowed without any ACL checks.

See Configuring Network Access with Access Control Rules (ACLs) and GET_PREAUTHENTICATED_URL Procedure for more information.

3. Check the result.

   In this example `status` contains the result that includes the Table Hyperlink.

```
{
"status":"SUCCESS","id":"LCvtpALqZgcHGL4Lxyzabcxyza-
QVEFngwh1UGhg8jjuFAHOQJLGFi",
"preauth_url":"https://dataaccess.adb.us-
ashburn-1.oraclecloudapps.com/adb/p/RIabc_xyz7Kw8ZLyZ2JzCXZQXpc/data",
"expiration_ts":"2025-01-20T16:42:49.527Z",
"expiration_count":10}
}
```

Notes for generating a Table Hyperlink with `DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL`:

- When you generate a Table Hyperlink on an Autonomous Database instance with a private endpoint, the result includes a name `private_preauth_url` with the value of the form: "`https://`*`private-endpoint`*`/adb/p/`*`parurl-token`*`/data`".

  When you generate a Table Hyperlink on an Autonomous Database instance with a private endpoint and the private endpoint is configured with **Allow public access** enabled, the result includes both the `preauth_url` for the public endpoint and `private_preauth_url`.

  See Configure Private Endpoints and Use a Private Endpoint with Public Access Allowed for more information.
- Use `DBMS_DATA_ACCESS.LIST_ACTIVE_URLS` to show Table Hyperlinks. See List Table Hyperlinks for details.
- Use `DBMS_DATA_ACCESS.EXTEND_URL` to extend the life of a Table Hyperlink. See EXTEND_URL Procedure for more information.

Generate a Table Hyperlink with UI Features Specified on Columns

When you generate a Table Hyperlink you can use the `column_lists` parameter to specify UI features for specified columns.

The `column_lists` parameter is a JSON value that specifies options by column. The supported values in `column_lists` are one or more of the following:

column_lists Value	Description
`order_by_columns`	Specifies the columns that support sorting. The columns are specified in a JSON array.
`filter_columns`	Specifies the columns that support filtering. The columns are specified in a JSON array.
`default_color_columns`	Specifies to only use the default coloring for the specified columns. The columns are specified in a JSON array.
`group_by_columns`	Specifies that group by is allowed for the specified columns (viewing the data by grouping the specified column is allowed). The columns are specified in a JSON array.

To specify table view column level UI features for a Table Hyperlink:

1. Identify the table or view or select statement that you want to share.

   This example generates a Table Hyperlink using the `column_lists` parameter with a table. You can also use this parameter when you generate a Table Hyperlink with a `SELECT` statement.

   See Generate a Table Hyperlink for a Table or a View and Generate a Table Hyperlink with a Select Statement for more information.

2. Run `DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL` to generate the Table Hyperlink and specify group by columns option for viewing with a browser:

   For example, to specify group by columns:

```
DECLARE
   status CLOB;
   BEGIN
      DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL(
          schema_name => 'ADMIN',
          schema_object_name    => 'TREE_DATA',
```

```
          expiration_minutes    => 360,
          service_name          => 'HIGH',
          column_lists          => '{ "group_by_columns": ["COUNTY",
"SPECIES"] }',
          result                => status);

       dbms_output.put_line(status);
    END;
/
```

The `column_lists` parameter is JSON that contains a list of JSON arrays of columns defining Table Hyperlink functionality. Use this parameter to specify the columns for one or more of the options: `order_by_columns`, `filter_columns`, `default_color_columns`, or `group_by_columns`.

For example:

```
column_lists => '{
        "group_by_columns":["COUNTY", "SPECIES"],
        "order_by_columns":["COUNTY"] }'
```

See GET_PREAUTHENTICATED_URL Procedure for more information.

3. Check the result.

   In this example `status` contains the result that includes information about the Table Hyperlink.

```
{
"status":"SUCCESS",
"id":"LLUZjJ5Yy8d0txydMiuxCVL_j4abc_xyzV1989GmgRInaBYTw",
"preauth_url":"https://dataaccess.adb.us-
ashburn-1.oraclecloudapps.com/adb/p/Rt_yRfTxbcabc_xyzWlANP5RmM9Y/data",
"expiration_ts":"2024-10-25T02:30:39.810Z",
"expiration_count":null
}
```

   See Use Table Hyperlinks to Access Data with the Group By Option in Table Format for an example showing group by columns usage.

4. Run `DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL` to generate the Table Hyperlink and specify a UI format option.

   For example, to set the columns that can be sorted use the `column_lists order_by_columns` option:

```
DECLARE
   status CLOB;
   column_lists CLOB;
   BEGIN
      DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL(
          schema_name => 'ADMIN',
          schema_object_name    => 'TREE_DATA',
          expiration_minutes    => 360,
          service_name          => 'HIGH',
          column_lists          => '{ "order_by_columns": ["COUNTY",
```

```
"SPECIES"] }',
          result                => status);
       dbms_output.put_line(status);
    END;
/
```

The `column_lists` parameter is JSON that contains a list of JSON arrays of columns defining Table Hyperlink functionality. Use this parameter to specify the columns for one or more of the options: `order_by_columns`, `filter_columns`, `default_color_columns`, or `group_by_columns`.

See GET_PREAUTHENTICATED_URL Procedure for more information.

5. Check the result.

   In this example `status` contains the result that includes information about the Table Hyperlink.

```
{
"status":"SUCCESS",
"id":"LLUZjJ5Yy8d0txydMiuxCVL_j4abc_xyzV1989GmgRInaBYTw",
"preauth_url":"https://dataaccess.adb.us-
ashburn-1.oraclecloudapps.com/adb/p/Rt_yRfTxbcabc_xyzk3EmM9Y/data",
"expiration_ts":"2024-10-25T02:30:39.810Z",
"expiration_count":null
}
```

   See Use Table Hyperlinks in Table Format with Column Sorting for an example showing order by columns usage.

You can also:

- Use `DBMS_DATA_ACCESS.LIST_ACTIVE_URLS` to show Table Hyperlinks. See List Table Hyperlinks for details.
- Use `DBMS_DATA_ACCESS.EXTEND_URL` to extend the life of a Table Hyperlink. See EXTEND_URL Procedure for more information.
- Generate Table Hyperlinks that are serviced with different service-level guarantees and resources. For example, access to an object or SQL statement can be mapped to services HIGH or MEDIUM, whereas access to another object or SQL statement can be mapped to the LOW service. See GET_PREAUTHENTICATED_URL Procedure for more information.

List Table Hyperlinks

You can list the active Table Hyperlinks that you generated on an Autonomous Database instance and the ADMIN user can list all active Table Hyperlinks.

Run `DBMS_DATA_ACCESS.LIST_ACTIVE_URLS` to list the active Table Hyperlinks. For example:

```
DECLARE
    result CLOB;
    BEGIN
       result := DBMS_DATA_ACCESS.LIST_ACTIVE_URLS;
       dbms_output.put_line(result);
    END;
/
```

This example shows the following output:

```
[{"id":"pT36lYHFGA4s3UXSNBCRO13v3D4_example1",
"created_by":"SCOTT",
"service_name":"HIGH",
"expiration_time":"2025-07-28T16:38:02.723Z",
"expiration_count":10,
"access_count":0,
"created":"2025-04-29T16:38:02.977Z",
"inherit_acl":true,
"is_group_url":false,
"group_ids":[null],
"sql_statement":"select * FROM TREE_DATA WHERE COUNTY = :county"}]
```

See LIST_ACTIVE_URLS Function for more information.

> **Note:**
>
> The behavior of `DBMS_DATA_ACCESS.LIST_ACTIVE_URLS` is dependent on the invoker. If the invoker is ADMIN or any user with `PDB_DBA` role, the function lists all active Table Hyperlinks, regardless of the user who generated the Table Hyperlink. If the invoker is not the ADMIN user and not a user with `PDB_DBA` role, the list includes only the active Table Hyperlinks generated by the invoker.

Invalidate Table Hyperlinks

At any time a user with appropriate privileges can invalidate a Table Hyperlink.

To invalidate a Table Hyperlink, you need the Table Hyperlink `id`. Use `DBMS_DATA_ACCESS.LIST_ACTIVE_URLS` to list each Table Hyperlinks and its associated `id`.

Use `DBMS_DATA_ACCESS.INVALIDATE_URL` to invalidate a Table Hyperlink. For example:

```
DECLARE
    status CLOB;
    BEGIN
       DBMS_DATA_ACCESS.INVALIDATE_URL(
        id => 'Vd1Px7QWASdqDbnndiuwTAyyEstv82PCHqS_example',
        result => status);
       dbms_output.put_line(status);
    END;
/
```

See INVALIDATE_URL Procedure for more information.

Define a Virtual Private Database Policy to Secure Table Hyperlink Data

By defining Oracle Virtual Private Database (VPD) policies for data that you share with a Table Hyperlink, you can provide fine-grained access control so that only a subset of data, rows, is visible for a specific Table Hyperlink.

Oracle Virtual Private Database (VPD) is a security feature that lets you control data access dynamically at row level for users and applications by applying filters on the same data set.

When a Table Hyperlink is accessed, the value of `application_user_id` specified during Table Hyperlink generation is available through `sys_context('DATA_ACCESS_CONTEXT$', 'USER_IDENTITY')`. You can define VPD Policies that make use of the value of this Application Context to restrict the data, rows, visible to the application user.

Any user who is granted access to read data with a Table Hyperlink can access and use the data (either a table, a view, or the data provided with a select statement). By defining a VPD policy on the database that generated a Table Hyperlink, you can use the `application_user_id` value in a SYS_CONTEXT rule to provide more fine-grained control. Consider an example where data is made available with a Table Hyperlink. If you want to restrict access to some of the data you can add a VPD policy.

For example:

1. Obtain the `application_user_id` value that you specified when you generated the Table Hyperlink.
2. Create VPD policy on the database where you generated the Table Hyperlink.

```
CREATE OR REPLACE FUNCTION limit_sal (v_schema IN VARCHAR2, v_objname IN
VARCHAR2)
  RETURN VARCHAR2 authid current_user AS
BEGIN
     RETURN 'employee_id = SYS_CONTEXT(''DATA_ACCESS_CONTEXT$'',
''USER_IDENTITY'')';
END;
```

   See DBMS_RLS for more information.

3. Register the VPD policy.

```
BEGIN
   DBMS_RLS.ADD_POLICY(
        object_schema        => 'HR',
        object_name          => 'EMPLOYEE',
        policy_name          => 'POL',
        policy_function      => 'LIMIT_SAL');
END;
/
```

   See DBMS_RLS for more information.

See Using Oracle Virtual Private Database to Control Data Access for more information.

Monitor and View Table Hyperlink Usage

Autonomous Database provides views that allow you to monitor Table Hyperlink usage.

Views	Description
V$DATA_ACCESS_URL_STATS and GV$DATA_ACCESS_URL_STATS Views	These views track Table Hyperlink usage, including elapsed time, CPU time, and additional information.

Use Table Hyperlink to Access Data

Table Hyperlink data is retrieved and returned in JSON format and is paginated.

> **Note:**
>
> Table Hyperlink data can be viewed in tabular format when accessed from a browser. See Use Table Hyperlinks to Access Data in Table Format for details.

You can access Table Hyperlink data using a browser or using any REST client. The data returned is paginated to allow you to access a maximum of 100 records at a time, with the total data size in the response limited to 1 MB. You can provide the `limit` query parameter to limit the number of records fetched. Table Hyperlink data retrieval is blocked if Table Hyperlink authentication fails or if the requested Table Hyperlink has expired.

For example, use a Table Hyperlink:

```
curl https://dataaccess.adb.us-chicago-1.oraclecloudapps.com/adb/p/K6XExample/
data
```

The Table Hyperlink response includes links for any previous or next pages, when the data includes more than one page. This allows you to navigate in either direction while fetching data. The JSON also includes a `self` link that points to the current page, as well as a `hasMore` attribute that indicates if there is more data available to fetch.

The following is the response format:

```
{
    "items": [],                <-- Array of records from database
    "hasMore": true OR false,   <-- Indicates if there are more records to
fetch or not
    "limit": Number,            <-- Indicates number of records in the page.
Maximum allowed number is 100.
    "offset": Number,           <-- Offset indicating the start of the
current page
    "count": Number,            <-- Count of records in the current page
    "links": [
        {
            "rel": "self",
            "href": "{Link to table hyperlink url for the current page}"
        },
        {
            "rel": "previous",
            "href": "{Link to table hyperlink url for the previous page}"
        },
        {
            "rel": "next",
            "href": "{Link to table hyperlink url for the next page}"
        }
    ]
}
```

For example, the following is a sample response from a Table Hyperlink (with newlines added for clarity):


```
{"items":[
{"COUNTY":"Main","SPECIES":"Alder","HEIGHT":45},
{"COUNTY":"First","SPECIES":"Chestnut","HEIGHT":51},
{"COUNTY":"Main","SPECIES":"Hemlock","HEIGHT":17},
{"COUNTY":"Main","SPECIES":"Douglas-fir","HEIGHT":34},
{"COUNTY":"First","SPECIES":"Larch","HEIGHT":12},
{"COUNTY":"Main","SPECIES":"Cedar","HEIGHT":21},
{"COUNTY":"First","SPECIES":"Douglas-fir","HEIGHT":10},
{"COUNTY":"Main","SPECIES":"Yew","HEIGHT":11},
{"COUNTY":"First","SPECIES":"Willow","HEIGHT":17},
{"COUNTY":"Main","SPECIES":"Pine","HEIGHT":29},
{"COUNTY":"First","SPECIES":"Pine","HEIGHT":16},
{"COUNTY":"First","SPECIES":"Spruce","HEIGHT":6},
{"COUNTY":"Main","SPECIES":"Spruce","HEIGHT":8},
{"COUNTY":"First","SPECIES":"Hawthorn","HEIGHT":19},
{"COUNTY":"First","SPECIES":"Maple","HEIGHT":16},
{"COUNTY":"Main","SPECIES":"Aspen","HEIGHT":35},
{"COUNTY":"First","SPECIES":"Larch","HEIGHT":27},
{"COUNTY":"First","SPECIES":"Cherry","HEIGHT":20},
{"COUNTY":"Main","SPECIES":"Pine","HEIGHT":37},
{"COUNTY":"Main","SPECIES":"Redwood","HEIGHT":78},
{"COUNTY":"Main","SPECIES":"Alder","HEIGHT":45},
{"COUNTY":"First","SPECIES":"Chestnut","HEIGHT":51},
{"COUNTY":"Main","SPECIES":"Hemlock","HEIGHT":17},
{"COUNTY":"Main","SPECIES":"Douglas-fir","HEIGHT":34},
{"COUNTY":"First","SPECIES":"Larch","HEIGHT":12},
{"COUNTY":"Main","SPECIES":"Cedar","HEIGHT":21},
{"COUNTY":"First","SPECIES":"Douglas-fir","HEIGHT":10},
{"COUNTY":"Main","SPECIES":"Redwood","HEIGHT":78}],

"hasMore":false,
"limit":100,
"offset":0,
"count":30,
"links":
[
{"rel":"self",
"href":"https://dataaccess.adb.us-ashburn-1.oraclecloudapps.com/adb/p/
F5Sn..._example/data"}
]}
```


Use Table Hyperlink to Access Data with Bind Variables

If there is a bind variable in the `SELECT` statement when you generate a Table Hyperlink, you must either use the `default_bind_values` parameter when you generate the Table Hyperlink to specify a default bind variable value or you must pass a bind variable value as a query parameter to access the Table Hyperlink data.

Bind variable support is available for `NUMBER` and `VARCHAR2` column types.

For example, when a Table Hyperlink is generated with the following SQL statement without the `default_bind_values` parameter:

```
sql_statement = 'SELECT * FROM TREE_DATA WHERE COUNTY = :countyName'
```

You must append a bind variable and specify a value as a query parameter to access the data. For example:

```
curl https://dataaccess.adb.us-chicago-1.oraclecloudapps.com/adb/p/
K6X...example/data?countyNAME=MAIN
```

The bind value does not need to be specified in the URL when the `DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL` includes the `default_bind_values` parameter:

```
sql_statement = 'SELECT * FROM TREE_DATA WHERE COUNTY = :countyNAME',
default_bind_values => '{"countyNAME" : "First"}',
```

In the case where `default_bind_values` defines a value for a variable you do not need to provide the value as a query parameter (the default value is supplied):

```
curl https://dataaccess.adb.us-chicago-1.oraclecloudapps.com/adb/p/
K6X...example/data

{"items":[
{"COUNTY":"First","SPECIES":"Pine","HEIGHT":16},
{"COUNTY":"First","SPECIES":"Spruce","HEIGHT":6},
{"COUNTY":"First","SPECIES":"Hawthorn","HEIGHT":19},
{"COUNTY":"First","SPECIES":"Cherry","HEIGHT":20},
{"COUNTY":"First","SPECIES":"Chestnut","HEIGHT":51}],
"hasMore":false,
"limit":100,
"offset":0,
"count":6,
"links":
[
{"rel":"self",
"href":"https://dataaccess.adb.us-ashburn-1.oraclecloudapps.com/adb/p/
gTlbq...example/data"}
]}
```

When `default_bind_values` defines a bind variable value you can override the bind variable value by explicitly specifying a value as a query parameter. For example:

```
curl https://dataaccess.adb.us-chicago-1.oraclecloudapps.com/adb/p/
K6X...example/data?countyNAME=MAIN
```

See GET_PREAUTHENTICATED_URL Procedure for more information.

Use Table Hyperlinks to Access Data in Table Format

Table Hyperlink data can be viewed in tabular format when accessed from a browser. The table format supports scrolling, sorting, filtering, cell colors, and Group By where you select a column name to see column data based on the values in the column.

- Use Table Hyperlinks and View Data in Default Table Format
  You can access data with a Table Hyperline from a browser and view the data in table format.
- Use Table Hyperlinks in Table Format with Column Sorting
  You can access data with a Table Hyperlink from a browser and view the data in table format and sort table data.
- Use Table Hyperlinks in Table Format with Column Filtering
  You can access data with a Table Hyperlink from a browser and view the data in table format and filter table data.
- Use Table Hyperlinks to Access Data in Table Format with Colors in Cells
  You can access data with a Table Hyperlink from a browser and view the data in table format and display cell values with color.
- Use Table Hyperlinks to Access Data with the Group By Option in Table Format
  You can access data with a Table Hyperlink from a browser and group the data by the values of cells in a column.

Use Table Hyperlinks and View Data in Default Table Format

You can access data with a Table Hyperline from a browser and view the data in table format.

To return data in table format, append the `?view=table` query parameter to any Table Hyperlink. For example, from a browser use a Table Hyperlink with `?view=table` appended:

```
https://dataaccess.adb.us-ashburn-1.oraclecloudapps.com/adb/p/F5Sn..._example/
data?view=table
```

Here is a sample response to a Table Hyperlink with `?view=table` appended:



COUNTY	SPECIES	HEIGHT
Main	Spruce	8
First	Douglas-fir	10
Main	Yew	11
First	Larch	12
First	Pine	16
First	Maple	16
Main	Hemlock	17
First	Willow	17
First	Hawthorn	19
First	Cherry	20

Column options are configurable when a Table Hyperlink is generated. See Generate a Table Hyperlink with UI Features Specified on Columns for more information.

Use Table Hyperlinks in Table Format with Column Sorting

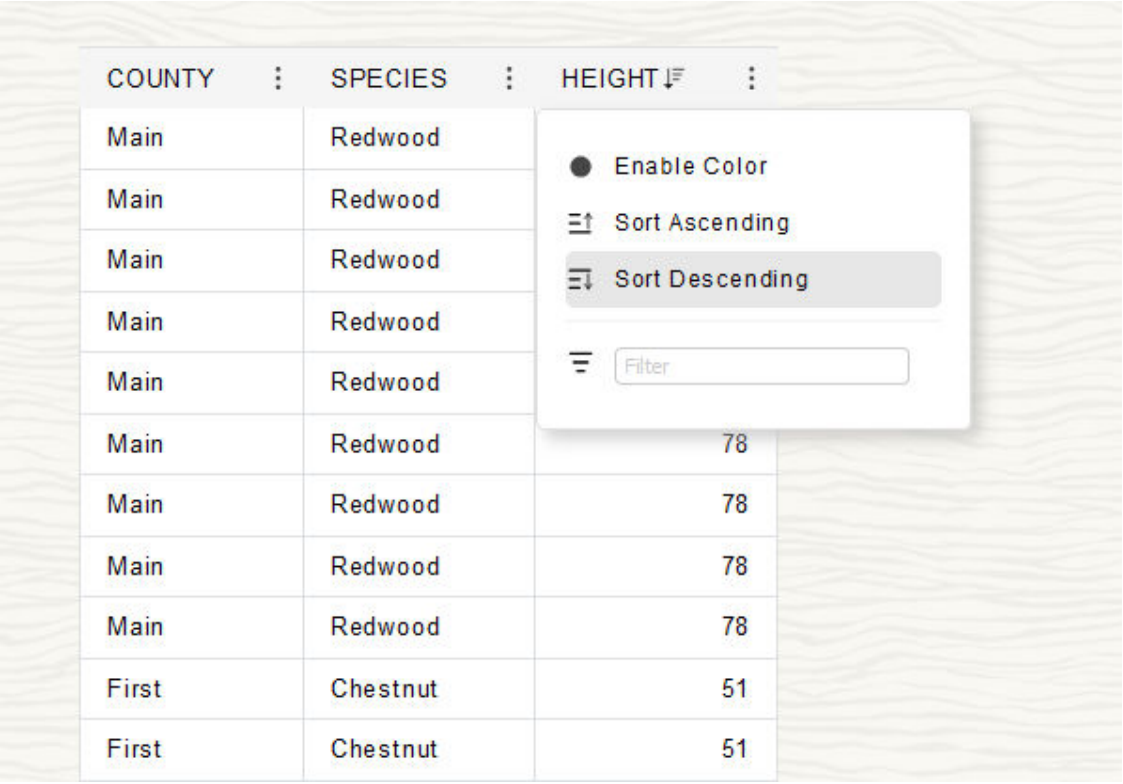
You can access data with a Table Hyperlink from a browser and view the data in table format and sort table data.

To return data in table format, append the `?view=table` query parameter to any Table Hyperlink. For example, from a browser use a Table Hyperlink with `?view=table` appended:

```
https://dataaccess.adb.us-ashburn-1.oraclecloudapps.com/adb/p/F5Sn..._example/
data?view=table
```

In table view you can sort Table Hyperlink columns in ascending or descending order by clicking  next to a column heading, and then select **Sort Ascending** or **Sort Descending**.

For example:



Column sorting is configurable by column when a Table Hyperlink is generated. See Generate a Table Hyperlink with UI Features Specified on Columns for more information.

Use Table Hyperlinks in Table Format with Column Filtering

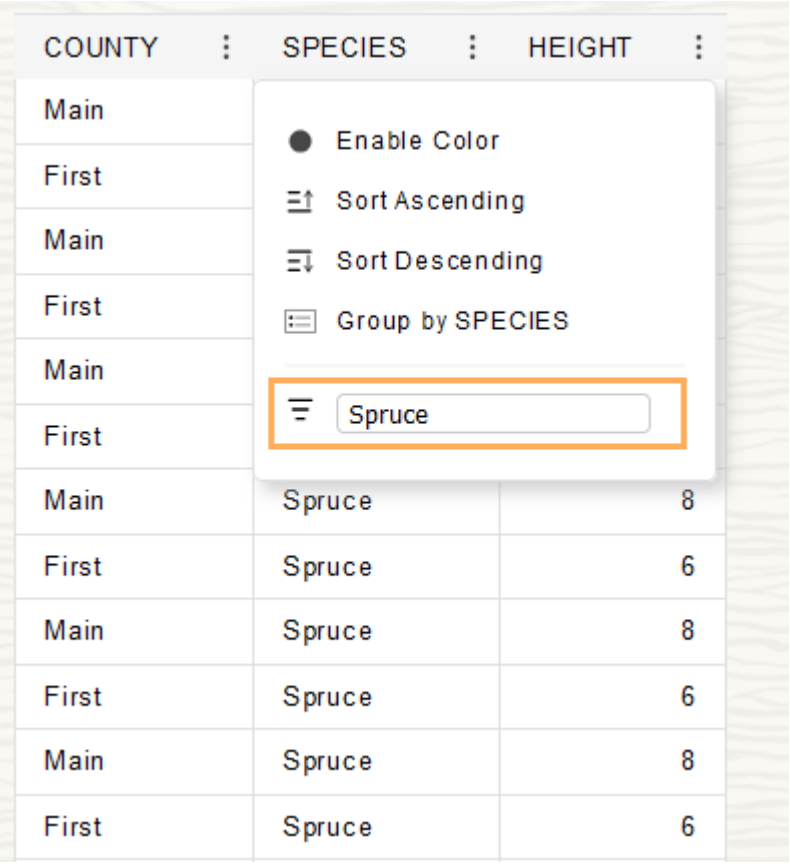
You can access data with a Table Hyperlink from a browser and view the data in table format and filter table data.

To return data in table format, append the `?view=table` query parameter to any Table Hyperlink. For example, from a browser use a Table Hyperlink with `?view=table` appended:

```
https://dataaccess.adb.us-ashburn-1.oraclecloudapps.com/adb/p/F5Sn..._example/
data?view=table
```

You can filter Table Hyperlink columns by clicking  next to a column heading, and then enter the filter text in the input text box, and type **Enter**.

The following is a screenshot of a sample table with the `SPECIES` column filtered to only display data for Spruce trees:



Column filtering is configurable by column when a Table Hyperlink is generated. See Generate a Table Hyperlink with UI Features Specified on Columns for more information.

Use Table Hyperlinks to Access Data in Table Format with Colors in Cells

You can access data with a Table Hyperlink from a browser and view the data in table format and display cell values with color.

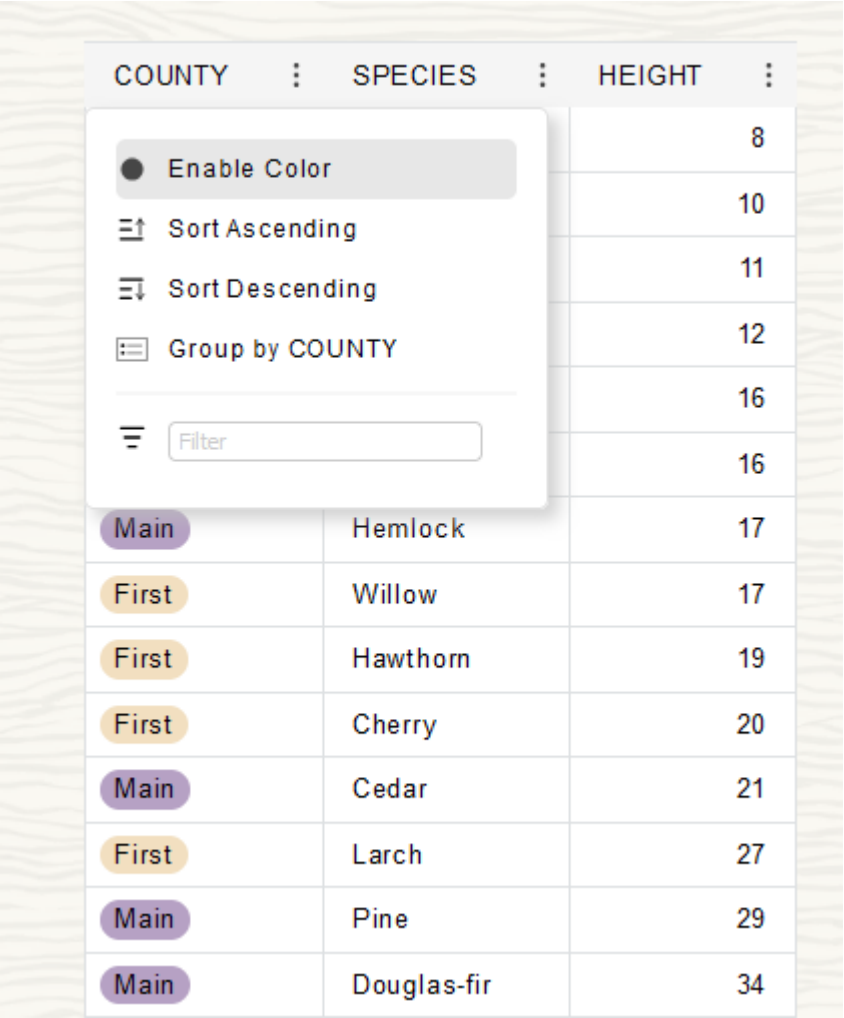
You can access Table Hyperlink data with a browser and view the returned data in database table like format. The table format supports scrolling, sorting, filtering and cell colors.

To return data in table format, append the `?view=table` query parameter to any Table Hyperlink. For example, from a browser use a Table Hyperlink with `?view=table` appended:

```
https://dataaccess.adb.us-ashburn-1.oraclecloudapps.com/adb/p/F5Sn..._example/
data?view=table
```

By default columns do not include coloring for cell values. You can add color to cell values in Table Hyperlink table view by clicking  next to a column heading, and then select **Enable Color**.

The following is a screenshot of a sample table with the `SPECIES` column showing color with values in cells:

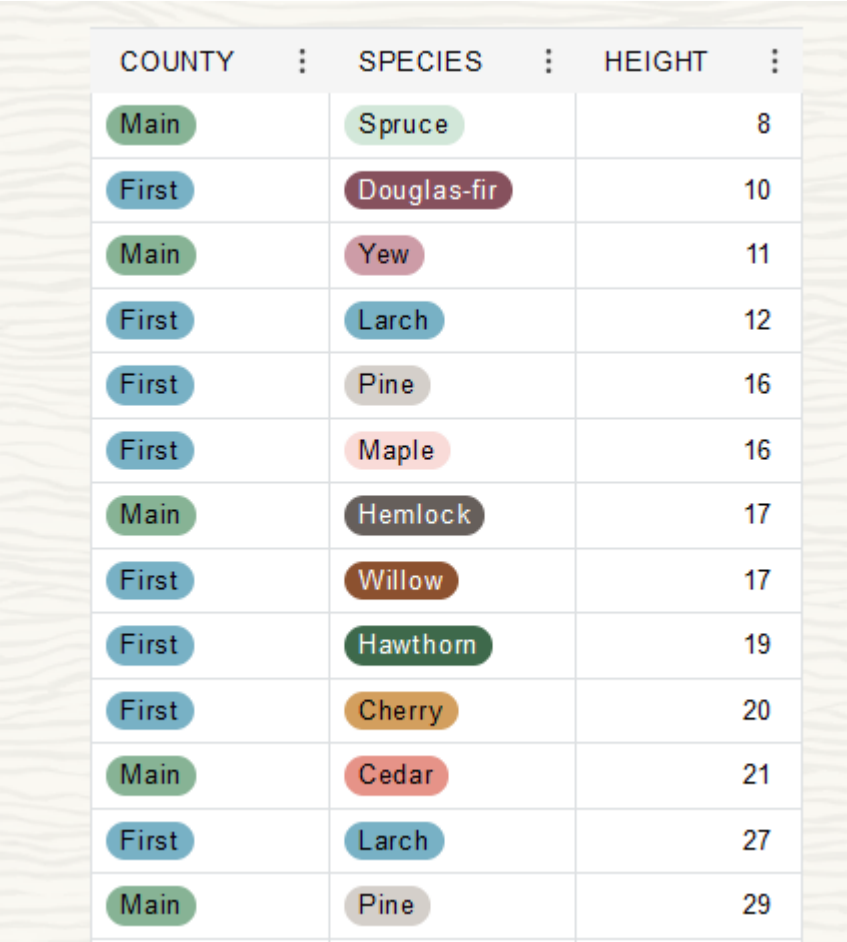


Color Column Data with Query Parameters

You can also use query parameters select the columns you want to add color to. To specify the columns with color, include the column names with the `colored_column_names` query parameter. For example, to color the `COUNTY` and `SPECIES` columns:

```
https://dataaccess.adb.us-ashburn-1.oraclecloudapps.com/adb/p/F5Sn..._example/
data?view=table&colored_column_names=COUNTY,SPECIES
```

The following is a screenshot of a sample table with COUNTY and `SPECIES` columns showing color with cell values:



COUNTY	SPECIES	HEIGHT
Main	Spruce	8
First	Douglas-fir	10
Main	Yew	11
First	Larch	12
First	Pine	16
First	Maple	16
Main	Hemlock	17
First	Willow	17
First	Hawthorn	19
First	Cherry	20
Main	Cedar	21
First	Larch	27
Main	Pine	29

Color Column Data Types with Query Parameter

You can also use query parameter to select a specific column data type to add color in a table by providing the data type with the `colored_column_types` query parameter.

This parameter supports `colored_column_types=VARCHAR` to color all the string (VARCHAR) columns, and `colored_column_types=NONE` to not color any data types.

For example:

```
https://dataaccess.adb.us-ashburn-1.oraclecloudapps.com/adb/p/F5Sn..._example/
data?view=table&colored_column_types=VARCHAR
```

Use Table Hyperlinks to Access Data with the Group By Option in Table Format

You can access data with a Table Hyperlink from a browser and group the data by the values of cells in a column.

The Table Hyperlink generation procedures supports column list options that provides support specifying columns and viewing column values with the group by option.

For example, when group by is specified when a Table Hyperlink is generated, do the following to view data by group:

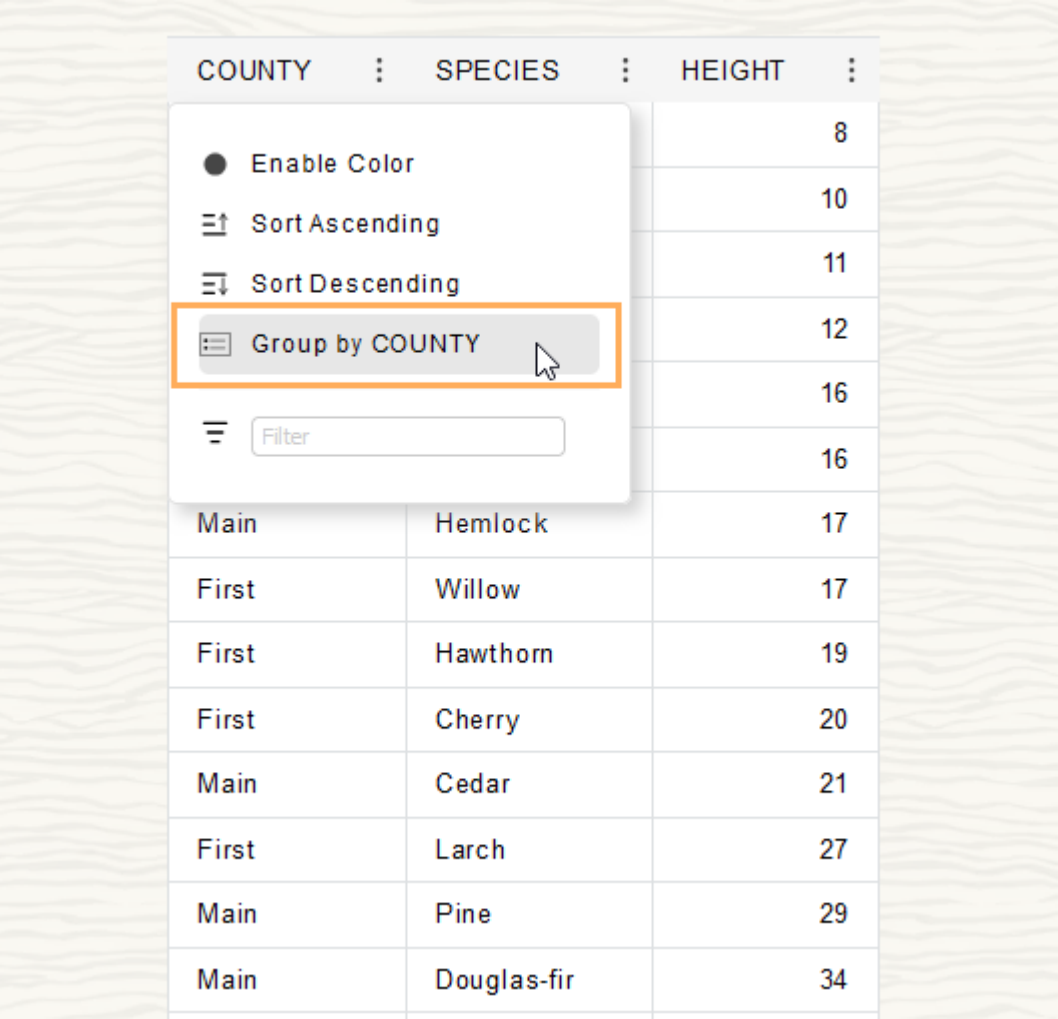
1. From a browser use a Table Hyperlink with `?view=table` appended.

For example:

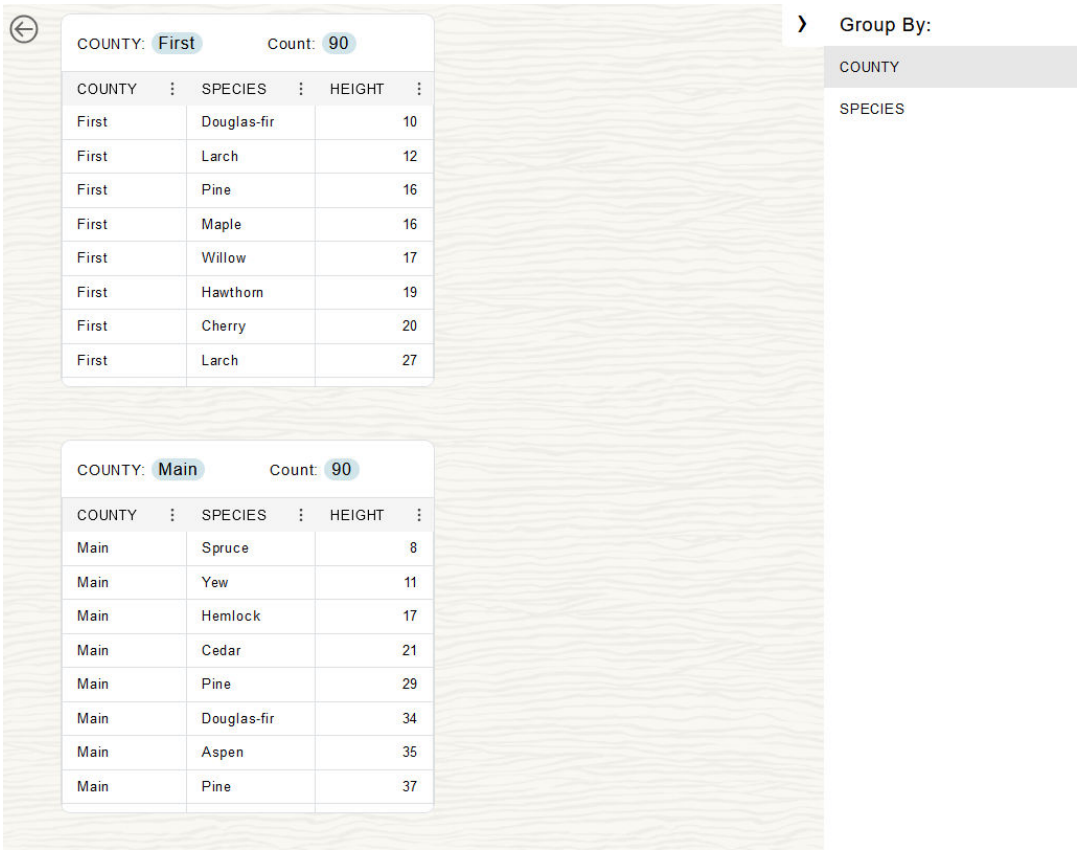
```
https://dataaccess.adb.us-ashburn-1.oraclecloudapps.com/adb/p/
F5Sn..._example/data?view=table
```

2. From a column that shows the group by option, select **Group by**.

   For example:



This shows the Group by display, and allows you to select another column if another column is available for group by display (depending on how the Table Hyperlink was generated).



By default the Table Hyperlink columns do not support group by options. This option is configurable by column when a Table Hyperlink is generated. See Generate a Table Hyperlink with UI Features Specified on Columns for more information.

Use Table Hyperlinks to Create an External Table

You can create an external table over an Autonomous Database Table Hyperlink.

You can use a single Autonomous Database Table Hyperlink or a comma-delimited list of Autonomous Database Table Hyperlinks to create an external table. See Query Data that Uses Table Hyperlink Source for more information.

Share Data with Data Studio

The Data Share tool enables you to share data and metadata with other users. The Data Share page displays information about the different types of shares available in the Oracle Autonomous Database.

See The Data Share Tool for more information.

- The Data Share Tool
  The Data Share tool allows you to share Oracle data and metadata with other databases and non-database tools. The Data Share page displays information about the different types of shares available in the Oracle Autonomous Database.

The Data Share Tool

The Data Share tool allows you to share Oracle data and metadata with other databases and non-database tools. The Data Share page displays information about the different types of shares available in the Oracle Autonomous Database.

- Overview of the Data Share Tool
  Oracle Autonomous Database enables you to create shares using the share tool.
- Access and Enable the Data Share Tool
  To share or consume data using the Data Share tool, navigate to the Data Studio, launch the Data Share tool, and enable sharing.
- Provide Share
  Use the Provide Share page to view the number of shares of the database, and the number of share recipients, search for a share entity, search for a recipient, create a new share, and create a new recipient.
- Consume Share
  Once the providers share the objects, there are a few steps the recipients need to follow to consume the share.
- Data Share Limitations on Autonomous Database
  This section summarizes the limitations of Data Share (both Versioned and Live Share) included in the Oracle Autonomous Database.
- Troubleshooting Errors with the Data Share Tool
  This chapter provides instructions on what issues cause slow performance and how to troubleshoot the issues.
- FAQs on the Data Share Tool
  The following are frequently asked questions on the Data Share tool.

Overview of the Data Share Tool

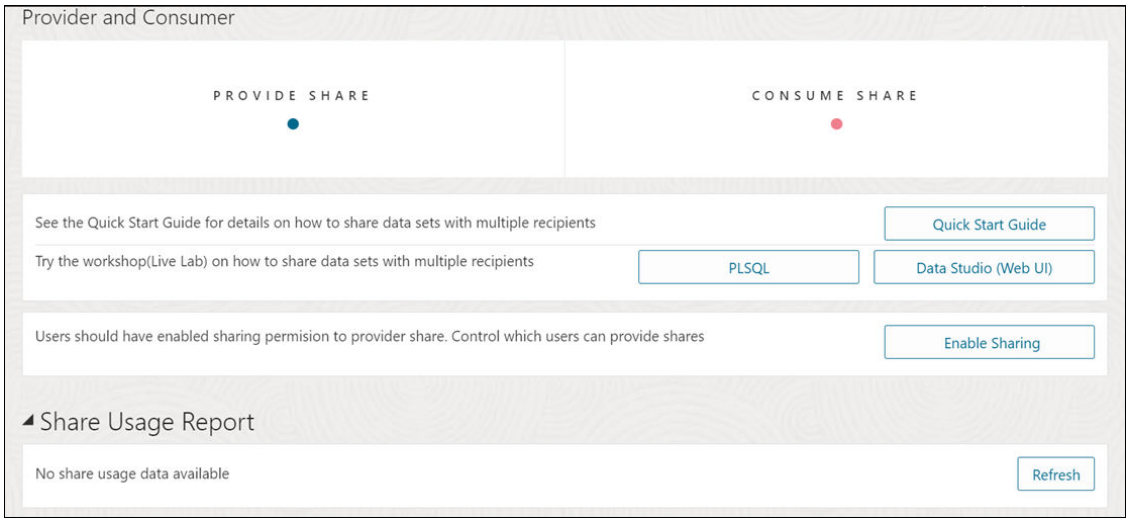
Oracle Autonomous Database enables you to create shares using the share tool.

Sharing objects requires two steps. The Provider provides data share for access and the Consumer role consumes (or receives) access to the published shares. The provider creates a share with the objects to share in the desired cloud object location. The provider also adds the recipient. The recipient accepts and receives the configured shared objects from the provider for consumption.

> **Note:**
>
> You must have the correct privileges to create or consume a data share. In case the Data Share card is disabled, click on the tool tip and follow the steps for the administrator to grant you the required privilege.

Select the **Data Share** menu from the Data Studio suite in the Database Actions homepage to access this tool. This opens the **Data Share** homepage. It consists of widgets that enable you to provide and consume share objects.

> **Note:**
>
> This is the homepage you view after you have enabled sharing and set Provider identification details.



> **Note:**
>
> If you do not see the Data Share tool card then your database user is missing the required DWROLE role.

Click Quick Start Guide to familiarize yourself with the Data Share tool.

Click PLSQL or Data Studio (Web UI) to try Data Sharing with PL/SQL or Data Studio without creating an account on Oracle Cloud tenancy.

Click **Enable Sharing** to grant sharing permission to you as a provider. See Access and Enable the Data Share Tool for more details.

The widgets are defined in the following sections:

- Provide Share
- Consume Share

Share Terminology

Provider: The Autonomous Database Serverless enables the provider to share existing objects. The share can contain a single table, a set of related tables, or a set of tables with some logical grouping. It could be a person, an institution, or a software system that shares the objects.

Example: An institution, such as NASA, that makes a data set available via data.gov.

Recipient: A Share recipient is an entity that associates an individual, an institution or a software system that receives a share from a provider. A recipient can have access to multiple shares. If you remove a recipient, that recipient loses access to all shares it could previously access.

Example: An external system, such as Microsoft Power BI, that supports the Delta Sharing REST API.

Share: A Share is a named entity in the provider's instance. It can be a group of datasets shared as a single entity.

Example: A SALES table that needs to be shared within an organization.

Overview of Providers and Recipients

A Data Share is the logical container that contains objects (such as tables) that share recipients will get access to a share and all tables within this share. A Data Share also implements security mechanisms on a high object level which simplifies the authorization for a set of individual objects. A provider creates and publishes share of a versioned type. The recipient is given access to a share. The provider can modify shares (both data and metadata) after the provider publishes the share to the recipients.

Use Case of Data Share

A marketing agency can share sales information with multiple interested parties. The Data Analysis tool analyses the data, generates insights and then the application shares the information with interested parties.

How the Data Share tool works?

Data is made accessible by the data sharing provider (that is an Oracle Autonomous Database) to the data sharing recipient at query time in parquet format for a versioned share. A live share uses cloud links and can only be consumed in an Oracle database The provider can only share data which they have access to when they log into an autonomous database instance.

As a data provider, you create a share and select other additional entities to share. The Oracle Data Sharing for general recipients is based on the open delta sharing standard protocol, providing a simple REST-based API to share data in parquet format. For near real-time access to shared data, customers can use Live Shares accessed using the consumer's ADB-S instance.

The Autonomous Database Serverless Versioned Sharing protocol works as follows:

- The provider creates and publishes a share that can be shared with one or multiple recipients. Every recipient will get a personal activation link to download their own JSON profile with the necessary information to access their share.
- The versioned share recipient registers with the share server by entering the URL for the end point along with a client ID, secret key and a bearer token.
- The versioned share recipient retrieves data from the share by calling the /shares/../tables/../query endpoint to obtain a list of URLs. The recipient then sends a GET request on these URLs to obtain the parquet files.

The Autonomous Database Serverless Live Sharing protocol works as follows:

- For a Live Share, the intended recipient will copy the sharing ID from the consumer page and publish the share that can be shared with recipients. This is the case when provider shares to only one database.

- A provider can share also share to **ALL_REGIONS**, **ALL_TENANCY**, or **ALL_COMPARTMENTS**.

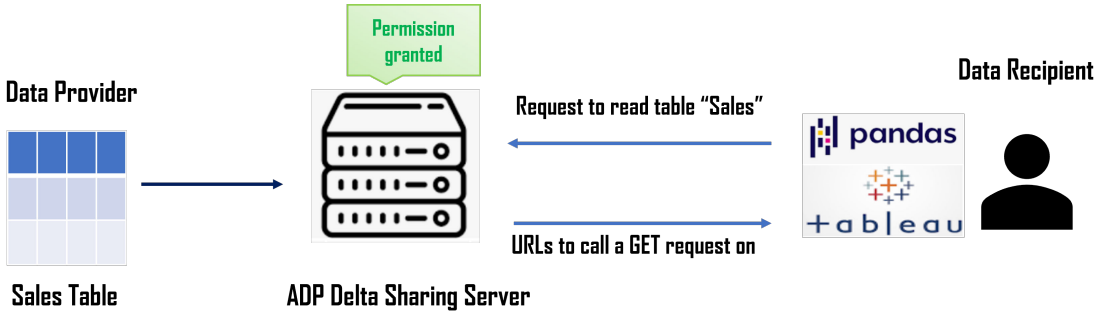
Features of Autonomous Database Serverless Share

With Autonomous Database Share you can:

- Share objects easily across Autonomous Databases and all tools or APIs that support the open delta sharing protocol.
- Share versioned data with many recipients without data replication for all recipients.
- Establish secure and centrally managed data sharing and collaboration within and across organizations.

Share Architecture

The following diagram is a generalized flow diagram of the architecture of the Data Share.



Prerequisites for Share Providers

Here are some prerequisites for a share provider to use the share tool:

- For a versioned share, you must have read and write access to a bucket to store or cache your shares.
- The schema you wish to use to create and publish shares must be enabled by an `ADMIN` user.

Prerequisites for Share Recipients

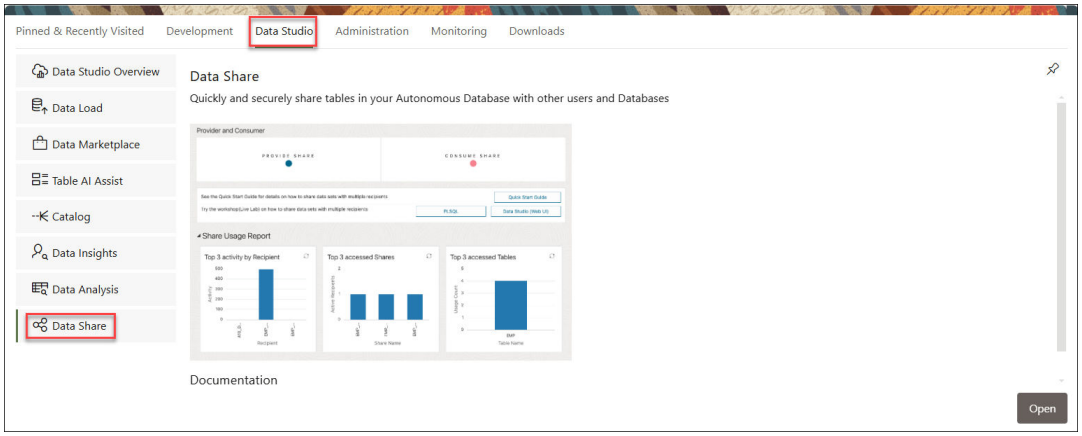
The share recipient must have a valid email address a provider can use to register the recipient to use the share tool. Oracle Data Share allows to share the information about a recipient's activation link by email.

Access and Enable the Data Share Tool

To share or consume data using the Data Share tool, navigate to the Data Studio, launch the Data Share tool, and enable sharing.

Follow these steps:

1. Navigate to the **Launchpad** on your Database Actions instance.
2. Click the **Data Studio** tab and select the **Data Share** pane.

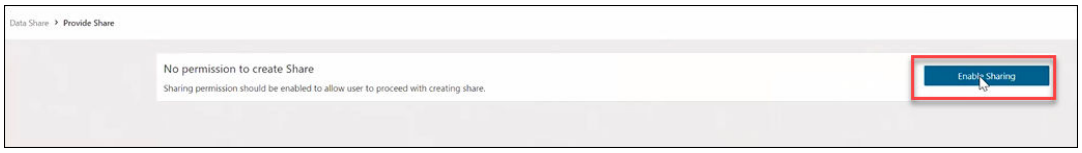


3. Clicking the Data Share pane opens a Provider and Consumer page where you can view **Provide Share** and **Consume Share** widgets.



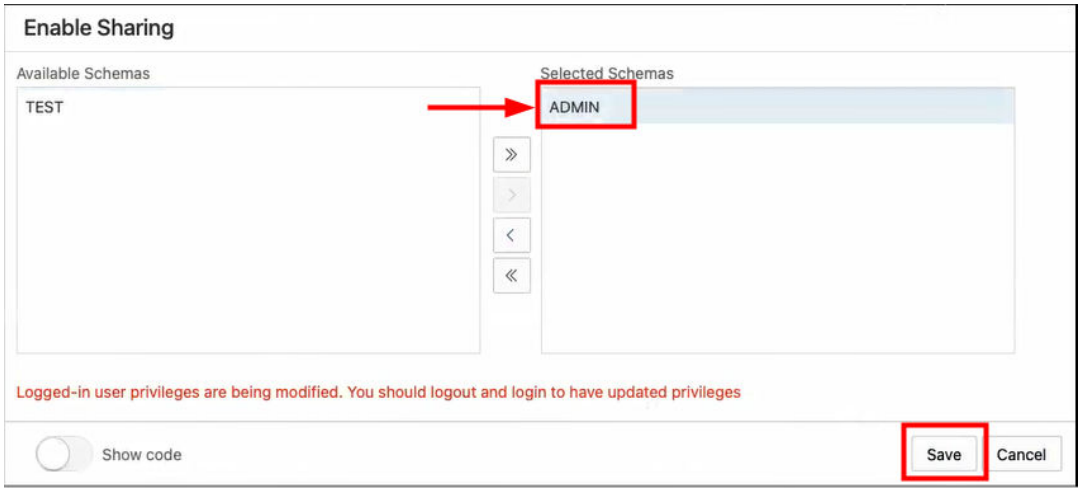
Click **Provide Share**.

4. On the **Provide Share** page, click **Enable Sharing**.



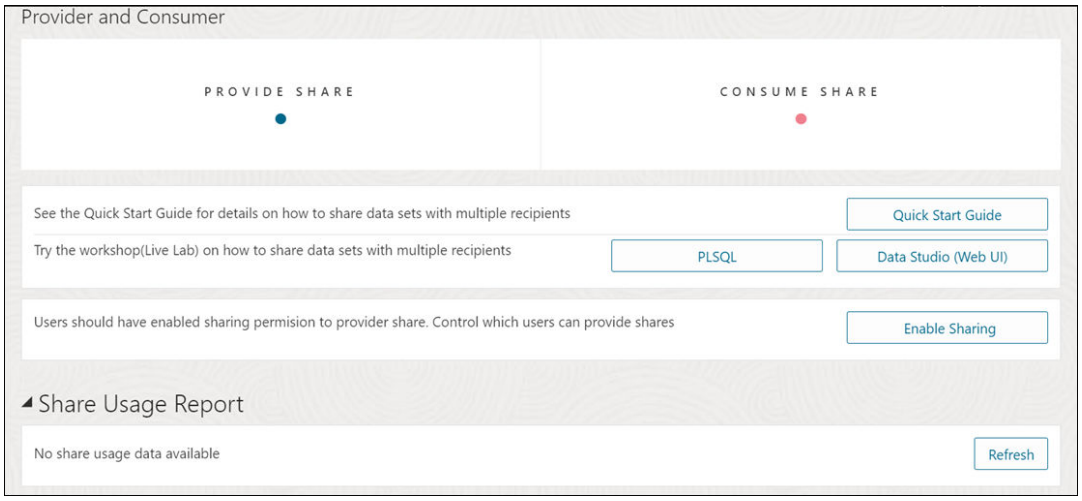
Clicking the **Enable Share** icon takes you to the Enable Sharing screen where you must select the user schema you want to enable and move it from the Available Schemas column to the Selected Schemas column.

5. Click **Save**.



6. Log out and log in again to update privileges.

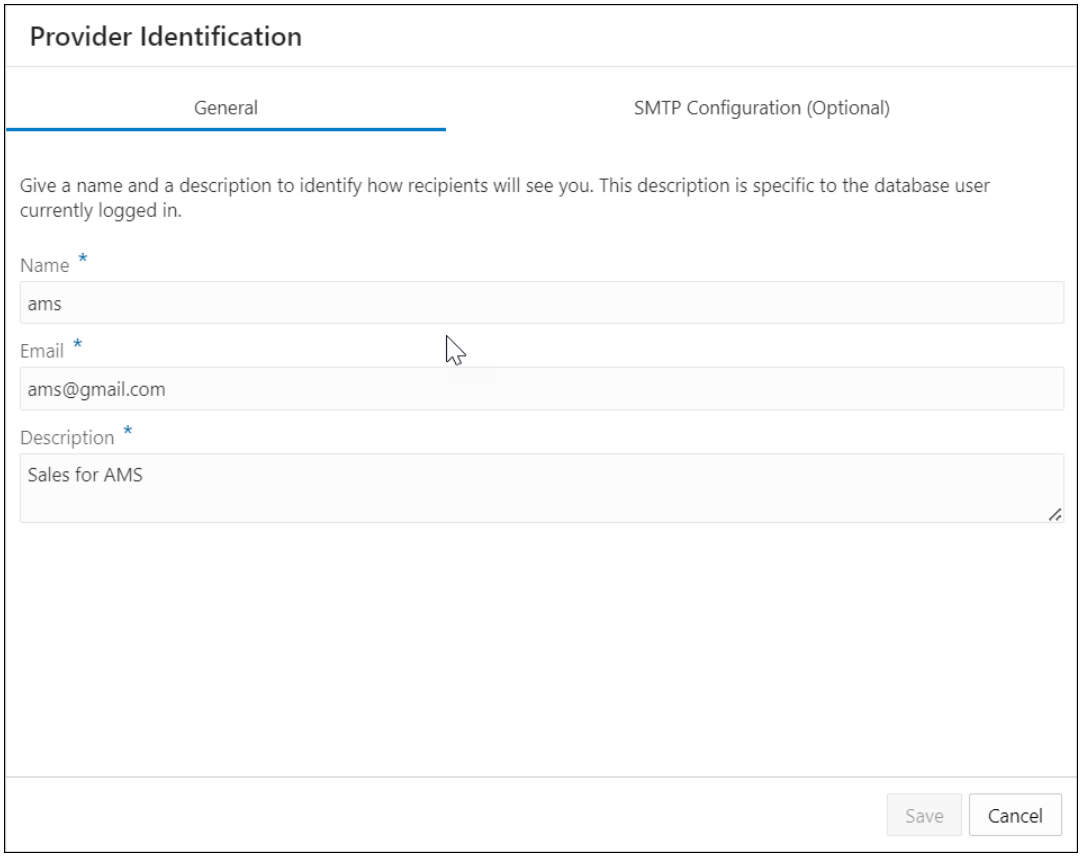
7. You now view the following **Provider** and **Consumer** page.



You have now enabled the share. Click **Provide Share**. On the Provide Share page, click **Provider Identification** to create a Provider ID. This provides information to the recipient on how to identify you.

This field provides the details of the provider before you share the data. The share provider identification will be available to recipients with whom you grant the share.

Selecting this button opens a Provider Identification wizard. Specify the following fields on the General tab of the Provider Identification wizard:



- **Name**: Enter the name of the provider. For example, *ams*.
- **Email**: Enter the email address of the provider. For example, *ams@gmail.com*.
- **Description**: Enter a description of the provider. For example, *AMS sharing stuff*.

Progress to the Simple Mail Transfer Protocol (SMTP) Configuration tab of the wizard. This tab configures email service like a local email client on your system (e.g., Thunderbird).

> **Note:**
>
> You must configure SMTP only once and then the system uses the saved configuration from that point forward.

Specify the following field values in the SMTP Configuration tab of the wizard:

- **Server host**: Enter the endpoint used to send the email. For example, *internal-mailrouter.oracle.com*.
- **Server port**: Enter the SMTP port used to accept an email. Email Delivery supports TLS on port 25. Sender: Enter the email address of the sender. For example, *oaapgens_us@oracle.com*.
- **Server Encryption**: This field indicates if TLS, the standard means of performing encryption in transit for emails, is being used. Providers must encrypt emails while it's in transit to the Oracle Cloud Infrastructure Email Delivery service. Encrypted emails are protected from being read during transit. If there is no encryption, enter, *None*.

Select credential to use for SMTP connection from the drop-down. If there is no available credential in the drop-down, you can create credential by clicking **Create Credential** icon. Refer to Create Oracle Cloud Infrastructure Native Credentials for more details.

Click **Test** to test the SMTP configuration. You will view a screen that asks you to run the ACL script. You can run the script if you have the ADMIN rights. This is just a first time setting. After you receive a successful message of SMTP Test, you can save the SMTP configuration.

You can now start sharing your data.

Provide Share

Use the Provide Share page to view the number of shares of the database, and the number of share recipients, search for a share entity, search for a recipient, create a new share, and create a new recipient.

The Provide Share page consists of widgets Shares and Recipients and icons to create Shares and Recipients. The page is described in the following sections:

- Provide Share Overview
  Use the Provide Share to create and publish shares.
- Provide Versioned Share
  In this type of share, data is shared and published as well-defined, known, as-of-snapshots. At the time of publication, the tool generates and stores the data share as parquet files in the specified bucket. The recipient can directly access the share in the object store.
- Provide Live Share
  In this type of share the recipient accesses data directly from the Oracle table or view. This share gives the recipient the latest data. Live Share works by querying data over a cloud

link, which is implemented on a database link. The database link allows you to query objects over the network on a foreign machine and schema. The advantage of this mode is that the data is up to date as of the time of query.

- View Share Entity Details
  Use the **Actions** icon at the right of the share entity entry to view details about the live share or delta share entity you create.
- Create Share Recipient
  Use the **+ Create Share Recipient** icon on the Provide Share page to create a consumer of the data share. When there is already a recipient created, use the **+ Create Recipient** icon on the Recipients widget of the Provide Share page to create a consumer of the data share.
- View Share Recipient Entity Details
  Use the **Actions** icon at the right of the share recipient entity entry to view details about the live share or delta share recipient entity you create.
- Provide Share Overview
  The Provide Share page enables you to view, create, edit, and access information on the shares and recipients you create from this page.
- Provide Versioned Share
  You can share data as a set of distinct versions, such as at the end of each day. The recipient will only see changes to the data when you publish a new version. As a versioned share Provider, you must create an OCI native credential and associate the bucket's URL with the credential.
- Provide Live Share
  You can share data as of the latest database commit with Autonomous Databases in the same region. The recipient will always see the latest data.
- View Share Entity Details
  To view details about an entity, click the **Actions** icon at the right of the Share entity entry, then click **View Details**.
- Create Share Recipient
  On the Provide Shares page, click the **Recipients** widget and select **+ Create Recipient** to create a new Share Recipient.
- View Share Recipient Entity Details
  Click the **Actions** icon at the right of the Share Recipient entity entry, then click **View Details**.

Provide Share Overview

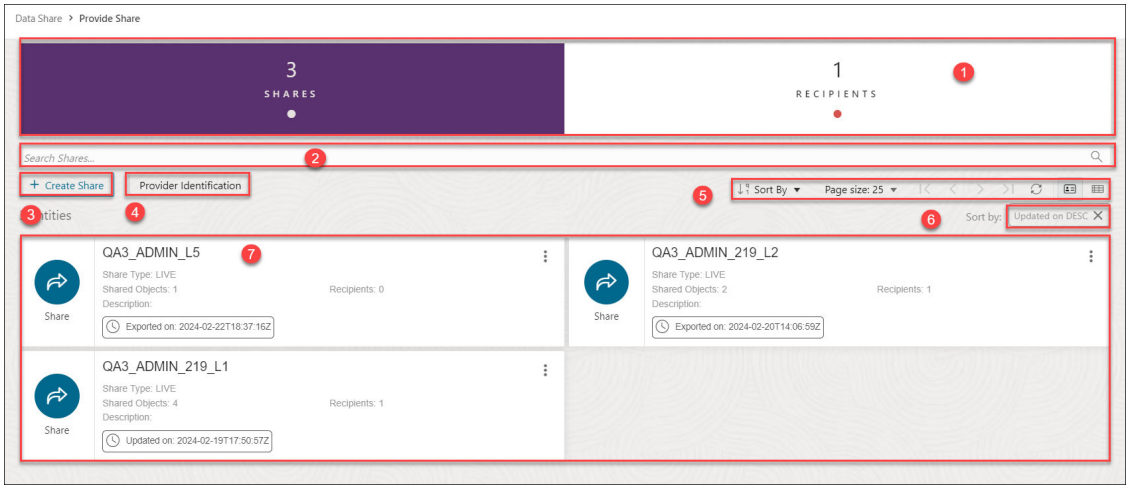
The Provide Share page enables you to view, create, edit, and access information on the shares and recipients you create from this page.

Click the **Provide Share** menu under Data Share in the Data Studio menu to reach the Provide Share page. You can switch the display of the Provide Share screen based on the widget you select from the Shares and Recipients section.

The search field and the display area of the **Provide Share** page varies when you switch between the **SHARES** and **RECIPIENTS** widgets.

Shares Page

The following image displays the Provide Share page when you select the **SHARES** widget.



The Provide Share page when you select the Shares widget consists of the following:

1. The top section of the Provide Share page displays Shares and Recipients section with two widgets. The widgets are defined in the following sections:
   - **SHARES**- Select this widget to view, search and perform actions on the Shares you create.
   - **RECIPIENTS**- Select this widget to view, search and perform actions on the Share Recipients you create.
2. **Search Shares** field
   Selecting Shares widget enables you to search for the share you create. You can click the field and type or paste the name of the share or the recipient you are looking for. Click the magnifier icon to return the search input. Click **X** to clear your searches in the Search Shares field.
3. **+ Create Share** button
   You can use the **Create Share** wizard to define data shares based on object storage data and associate them with recipients. A single data share can consist of multiple shared objects that are shared with one or multiple recipients.

   You can create a new share and a new recipient from this button. Refer to the Provide Live Share, Provide Versioned Share and Create Share Recipients section to explore more on what each button does.
4. **Provider Identification** button : See Access and Enable the Data Share tool to view more about this icon for a first time user. If you have already configured SMTP, you can edit the Provider ID or set SMTP using this button.
5. **Toolbar**
   The toolbar consists of the following buttons:
   - **Sort By**
     To select sorting values, click the **Sort By** button to open the list of options. Then click the **Ascending** or **Descending** icon next to one or more of the sorting values.

     For example, if you select the **Ascending** icon next to **Entity name** and the **Descending** icon next to **Entity type**, the entities will be sorted in alphabetical order by entity name and then in reverse alphabetical order by entity type.

     Click **Reset** in the list to clear the choices in the list.

     The sorting values you choose are listed next to the **Sort by** label beneath the toolbar. Click the **X** icon on a sorting value to remove it.
   - **Page size**

By default, up to 25 entities are displayed on the page. If you want more entities on a page, select a number from this list.

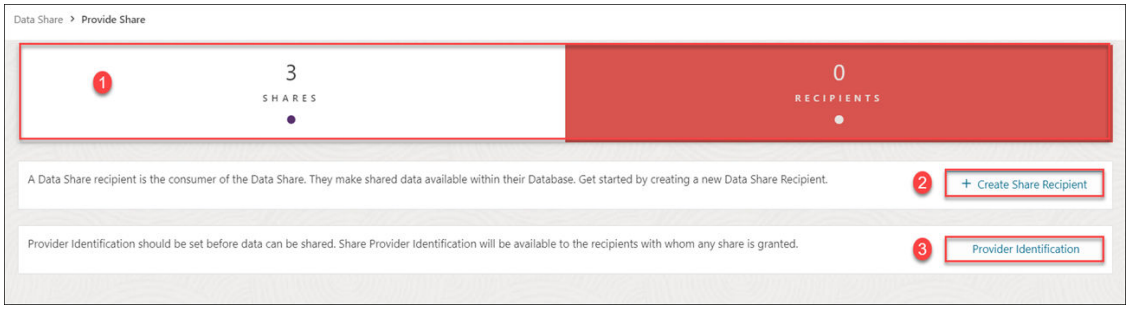
- **Previous** and **Next**
  If the search results are displayed on multiple pages, click these buttons to navigate through the pages.
- **Refresh**
  Click to refresh the entities shown on the page, based on the current search string.
- **Entity view options**
  Choose one of these three options to set how entities are displayed on the page.

  Click **Open Card view** to display entities as card arranged into one or two columns; click **Open Grid View** to display entities as rows in a table; or click **Open List View** to display entities in a single column of panels.

6. **Sort by settings**
   When you set sorting values by using the Sort By control in the toolbar, the settings are displayed in small boxes beneath the toolbar. You can delete a setting by clicking the X icon in the box. Or you can change the settings by returning to the Sort By control in the toolbar.
7. **Display area**
   The area below the Search field displays the entities returned by a search. If you select the Shares widget, you will view the Share entities in the display area. If you select the Recipient widget, you will view the Share Recipient entities in the display area. This area shows all providers you are subscribed to.

Recipients Page

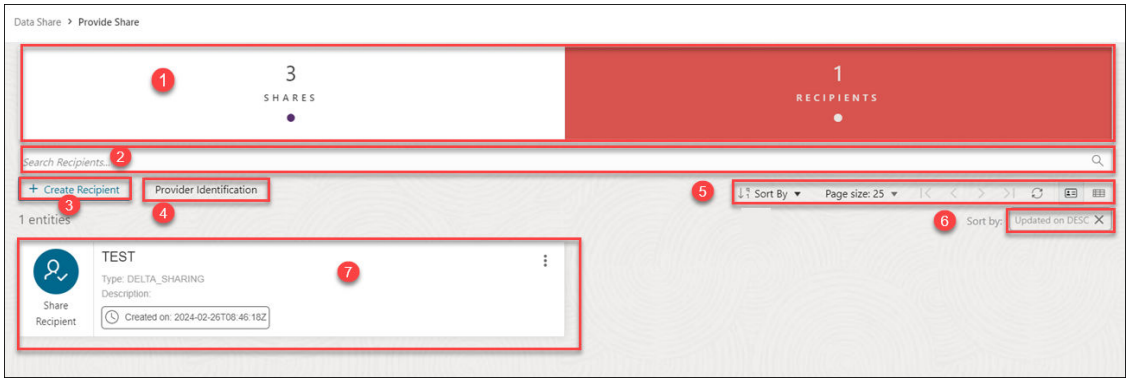
The following image displays the Provide Share page when you select the RECIPIENTS widget and ***there is no RECIPIENT created***



Click **+ Create Share Recipient** to create a share recipient. See Create Share Recipient for more details.

Click **Provider Identification** to specify the name and description to identify how the recipients will view you. You can also configure SMTP settings.

The following image displays the Provide Share page when you select the RECIPIENTS widget and ***there is minimum one RECIPIENT already created***



The Provide Share page consists of the following:

1. The top section of the Provide Share page displays Shares and Recipients section with two widgets. The widgets are defined in the following sections:
   - **SHARES**- Select this widget to view, search and perform actions on the Shares you create.
   - **RECIPIENTS**- Select this widget to view, search and perform actions on the Share Recipients you create.
2. **Search Recipients** field
   Searching for a share or recipient depends on the widget you select from the Shares and Recipients section. If you select the Recipients widget, you can search for the recipients you create. You can click the field and type or paste the name of the share or the recipient you are looking for. Click the magnifier icon to return the search input. Click **X** to clear your searches in the Search Shares or Search Recipients field.
3. **+ Create Recipient** button
   You can create a new share and a new recipient from this button. Refer to the Provide Live Share, Provide Versioned Share and Create Share Recipients section to explore more on what each button does.
4. **Provider Identification** button: See Access and Enable the Data Share tool to view more about this icon for a first time user. If you have already configured SMTP, you can edit the Provider ID or set SMTP using this button.
5. **Toolbar**
   The toolbar consists of the following buttons:
   - **Sort By**
     To select sorting values, click the **Sort By** button to open the list of options. Then click the **Ascending** or **Descending** icon next to one or more of the sorting values.

     For example, if you select the **Ascending** icon next to **Entity name** and the **Descending** icon next to **Entity type**, the entities will be sorted in alphabetical order by entity name and then in reverse alphabetical order by entity type.

     Click **Reset** in the list to clear the choices in the list.

     The sorting values you choose are listed next to the **Sort by** label beneath the toolbar. Click the **X** icon on a sorting value to remove it.
   - **Page size**
     By default, up to 25 entities are displayed on the page. If you want more entities on a page, select a number from this list.
   - **Previous** and **Next**
     If the search results are displayed on multiple pages, click these buttons to navigate through the pages.

- **Refresh**
  Click to refresh the entities shown on the page, based on the current search string.
- **Entity view options**
  Choose one of these three options to set how entities are displayed on the page.

  Click **Open Card view** to display entities as card arranged into one or two columns; click **Open Grid View** to display entities as rows in a table; or click **Open List View** to display entities in a single column of panels.

6. **Sort by settings**
   When you set sorting values by using the Sort By control in the toolbar, the settings are displayed in small boxes beneath the toolbar. You can delete a setting by clicking the **X** icon in the box. Or you can change the settings by returning to the Sort By control in the toolbar.
7. **Display area**
   The area below the Search field displays the entities returned by a search. If you select the Shares widget, you will view the Share entities in the display area. If you select the Recipient widget, you will view the Share Recipient entities in the display area.

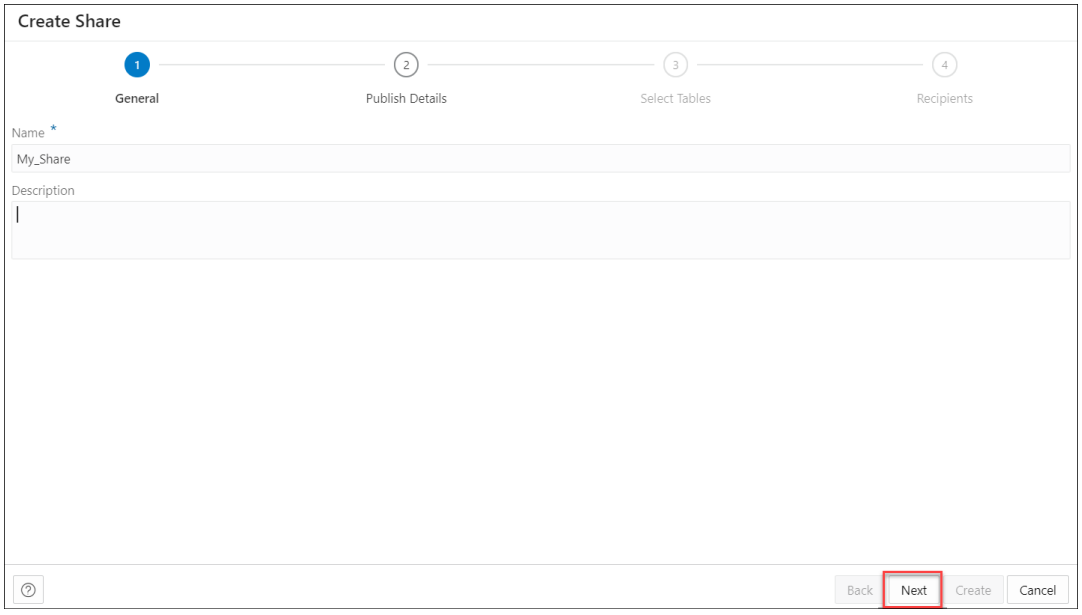
Provide Versioned Share

You can share data as a set of distinct versions, such as at the end of each day. The recipient will only see changes to the data when you publish a new version. As a versioned share Provider, you must create an OCI native credential and associate the bucket's URL with the credential.

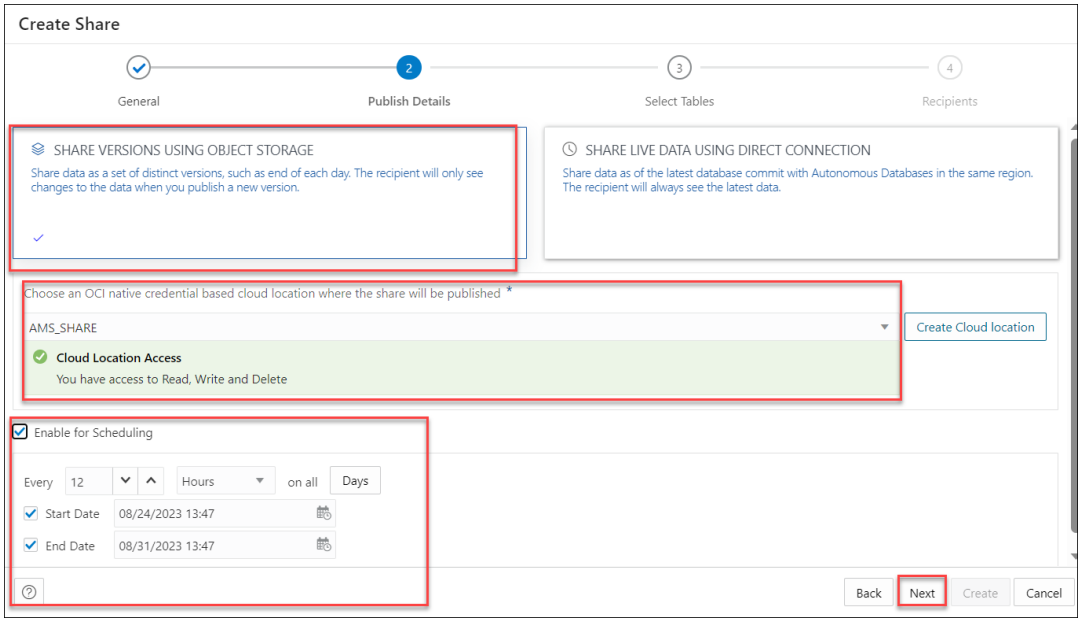
On the Provide Share page, select **+ Create Share** from the Create menu. This brings up the **Create Share** wizard.

1. On the Create Share wizard, in the **Name** field of the General tab, enter a name for the Share. For example: *My_Share*.
   In the **Description** field, enter a description for the link. For example: *Weekly Sales report*.

   Select **Next** to progress to the Publish Details tab of the Create Share wizard. You can alternatively click on the Publish Details tab to progress to the Publish Details screen of the wizard.



2. In the Publish Details Tables tab of the wizard, select **SHARE VERSIONS USING OBJECT STORAGE**.



Select a cloud location where you want the share to be published. If you already have an OCI cloud storage location available, use the **Select Cloud location** drop-down and select any available Cloud location from the drop-down. You will view a successful message if you have access to read, write, or delete in that cloud location.

> **Note:**
>
> Select **Create Cloud location** and follow the steps described in the Create Oracle Cloud Infrastructure Native Credentials if you do not have a cloud location available.

For sharing versioned data, you can enable the schedule for share. Select **Enable For Scheduling** option to set up a schedule for sharing. In the time interval fields, enter a number, and select a time type and the days on which to poll the bucket for new or changed files. For example, to poll every two hours on Monday, Wednesday, and Friday, enter **2**, select **Hours**, and then select **Monday**, **Wednesday**, **Friday** in the appropriate fields. You can select **All Days**, **Monday to Friday**, **Sunday to Thursday**, or **Custom** from the Week Days drop-down. The Custom field enables you to select Monday, Tuesday, Wednesday, Thursday and Friday in the appropriate fields.

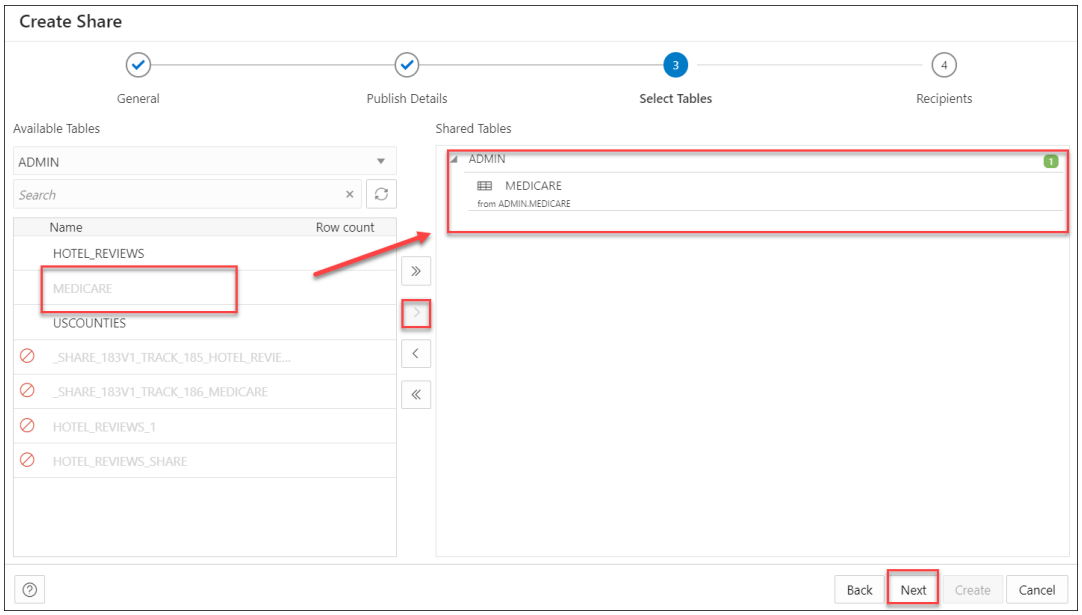
Select a start and end date. If you don't select a start date, the current time and date are used as the start date. The end date is optional. However, without an end date, the share will continue to poll. This is an optional step.

Click **Next** to progress to the Select Tables tab of the Create Share wizard.

3. On the **Select Tables** tab of the wizard, select the schema from the drop-down menu and click the table you wish to share from the Available Tables. Choose any of the available options:

   - **>**: This option enables you to move the table to Shared Tables.
   - **<**: To remove the selected table from Shared Tables select this option.
   - **>>**: This option allows you to move all the tables to the Shared Tables screen.

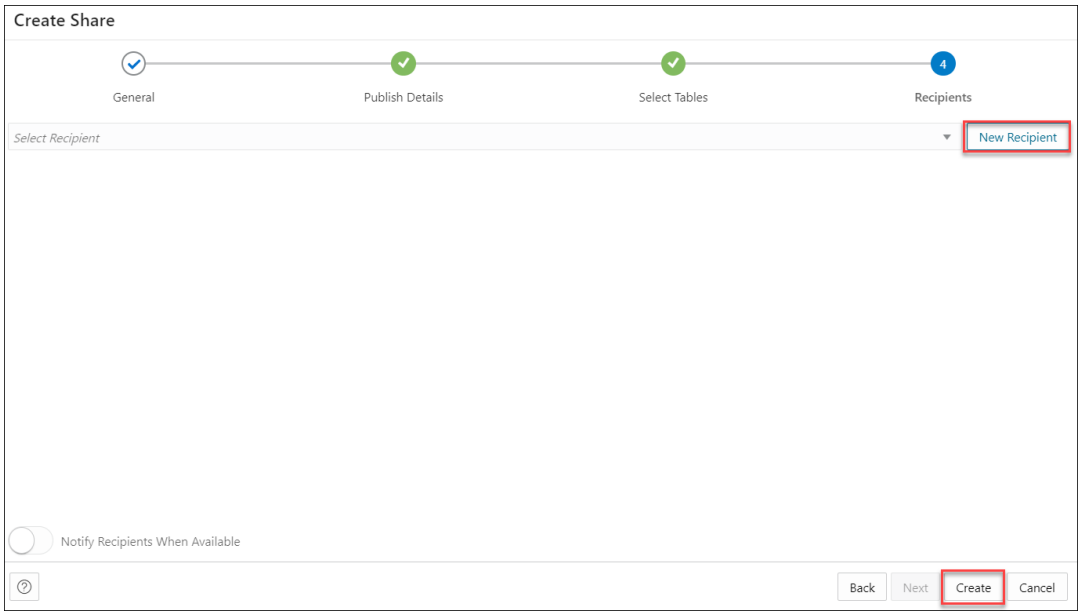
- **<<**: To remove all the selected tables from Shared Tables select this option.



Click **Next** to proceed to the Recipients tab of the Create Share wizard.

> **Note:**
>
> You may find it convenient to click on the specific tab to move back and forth. You can update changes on any screen this way.

4. On the Recipients tab of the Create Share wizard, select available Recipient from the drop-down.



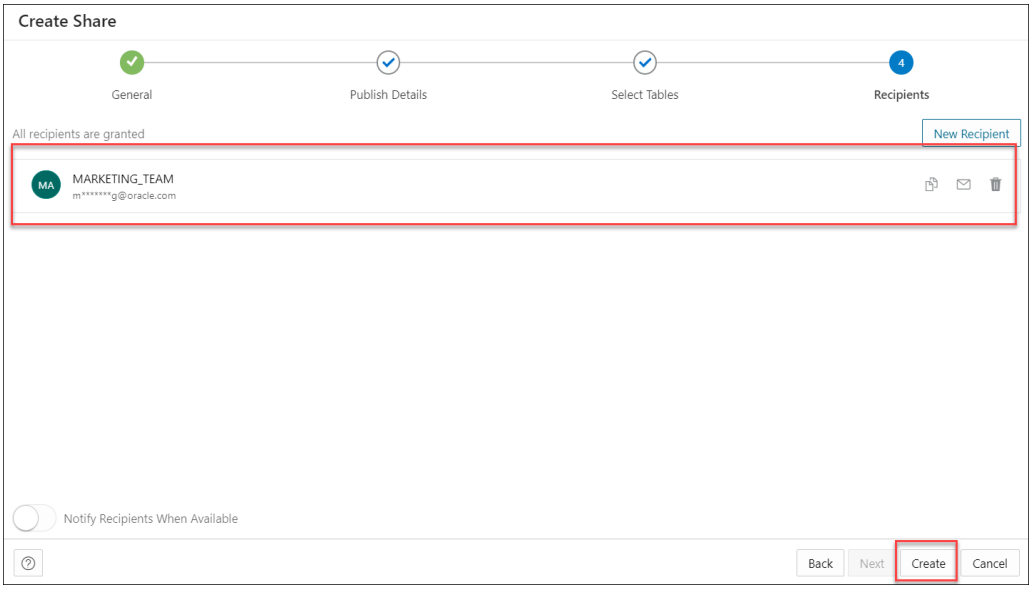
Select **Notify Recipients When Available** when you want to notify the recipients when they are available. This can be enabled when SMTP is configured to send an email notification to the recipient.

5. In case you want to create a new recipient, select **New Recipient**. Selecting New Recipient opens a **Create Share Recipient** wizard. See Create Share Recipient.

> **Note:**
>
> Ensure there are no spaces in the name of the recipient. This field does not support duplicate recipient names. If you try creating a recipient with a duplicate name, the creation of the share fails with the "Name is already used by an existing object" error.

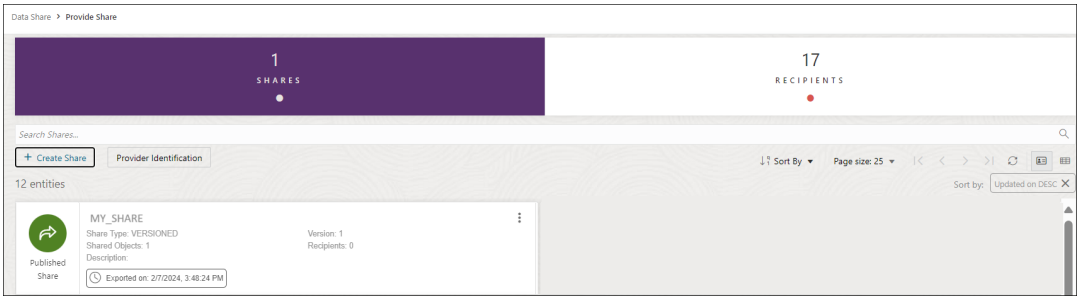
Click **Create** to finish creating the share recipient. The newly added recipient is displayed in the list of recipients. Click **Cancel** to cancel the ongoing process of recipient creation.

6. After you create a new share recipient, in the granted recipient's list you can:
   - Select **copy profile activation link to clipboard** icon to copy the activation link.
   - Select **Email recipient profile download link** icon to mail the profile link to the recipient you select in this step. Once you select it, the Share tool triggers an activation mail with the profile link to the recipient's email address. Click **Send** to share the profile link with the recipient. See *About the Consume Share Page* to know more about the activation mail and how to proceed as a recipient once you receive the activation mail.
   - Click **deselect Recipient** icon to remove the selected recipient.



Select **Create** to create the share. Note that the status of the share changes from unpublished to published depending on the size of the table in the share. You will receive a successful share notification message at the top right of the display that says the share is successful. The Provide Share page displays the share along with its details such as the entity type, owner, shared objects and the recipients.

> **Note:**
>
> The Data Share tool can create only two delta share profiles for the same recipient. If you try to create a third profile, the tool invalidates the first profile, and the newly created profile becomes active.



> **Note:**
>
> A versioned share will show different colors and states depending upon what state the share is in.

Provide Live Share

You can share data as of the latest database commit with Autonomous Databases in the same region. The recipient will always see the latest data.

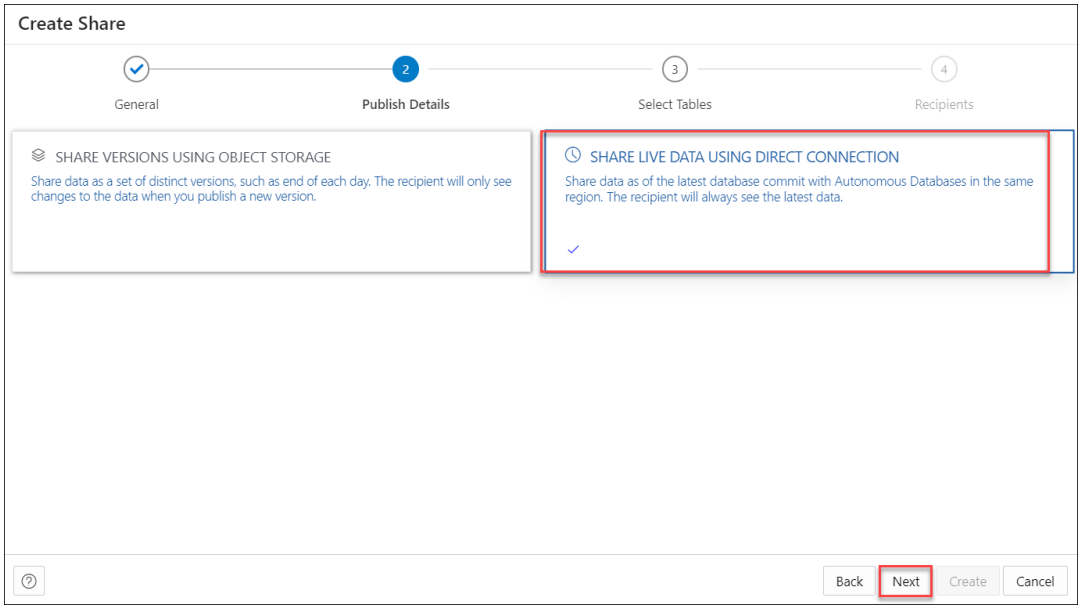
To provide access to live shared data, you must define data shares and grant recipients access to consume them. Use the **Create Share** wizard to create and publish shares for live data.

On the Provide Share page, select **+ Create Share** from the Create menu. This brings up the Create Share wizard.

1. On the Create Share wizard, in the **Name** field of the General tab, enter a name for the Share. For example: *My_Share*.
   In the **Description** field, enter a description for the data you are sharing. For example: *Weekly Sales report*.

   Select **Next** to progress to the Publish Details tab of the Create Share wizard. You can alternatively click on the Publish Details tab to progress to the Publish Details screen of the wizard.

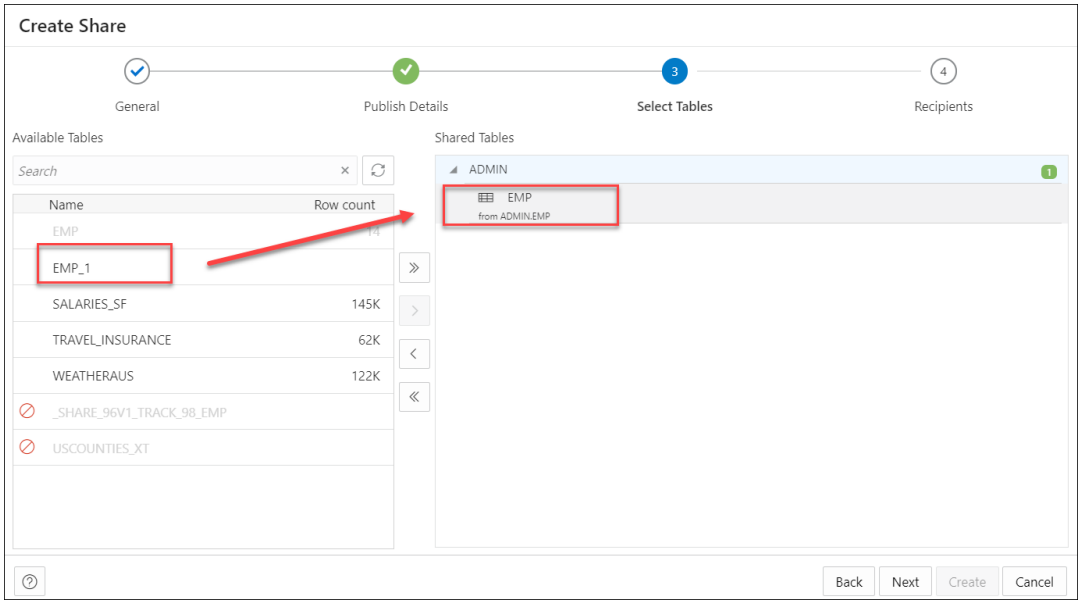
2. In the Publish Details Tables tab of the wizard, select **SHARE LIVE DATA USING DIRECT CONNECTION**.



Click **Next** to progress to the Select Tables tab of the Create Share wizard.

3. On the **Select Tables** tab of the wizard, select the schema from the drop-down menu and click the table you wish to share from the Available Tables. Choose any of the available options:

- **>**: This option enables you to move the table to Shared Tables.
- **<**: To remove the selected table from Shared Tables select this option.
- **>>**: This option allows you to move all the tables to the Shared Tables screen.
- **<<**: To remove all the selected tables from Shared Tables select this option.

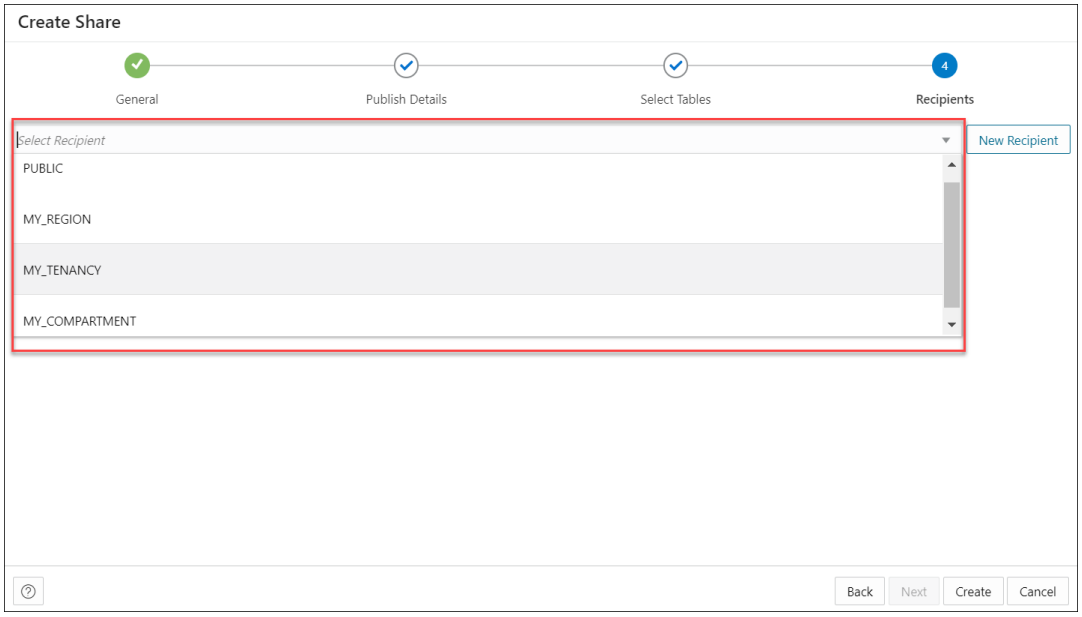


> **Note:**
>
> If a live share data provider is sharing multiple tables which would normally be joined together, it is recommended that the producer create a view that performs the joins and then only shares the view.

Click **Next** to proceed to the Recipients tab of the Create Share wizard.

> **Note:**
>
> You may find it convenient to click on the specific tab to move back and forth. You can update changes on any screen this way.

4. On the Recipients tab of the Create Share wizard, select available Recipients from the drop-down. The list of recipients available in the drop-down depends on the scope of the share. You can select any or all the following values depending on with whom you intend to share the data:
   - **MY_REGION**: Access to the data share is allowed for databases in the same region as the data share provider.
   - **MY_TENANCY**: Access to the data share is allowed for databases in the same tenancy as the data share provider.
   - **MY_COMPARTMENT**: Access to the data share is allowed for databases in the same compartment as the data share provider.

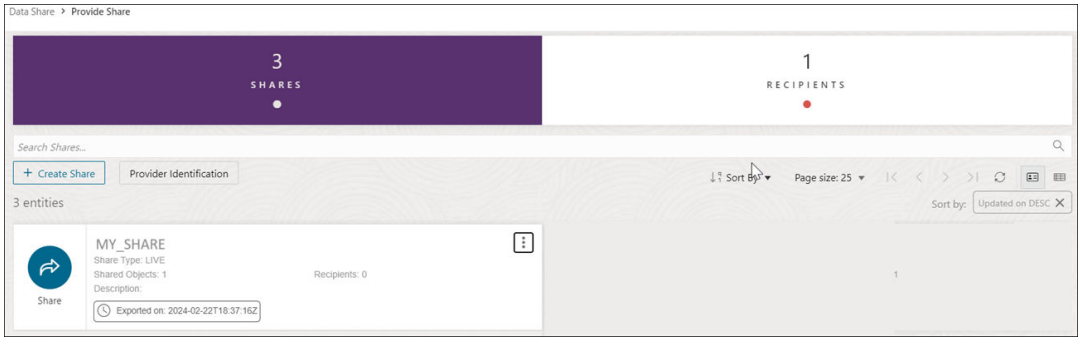


5. In case you want to create a new recipient, select **New Recipient**. Selecting New Recipient opens a **Create Share Recipient** wizard. Enter the recipient's Name, Description (optional), Sharing ID, Email in their respective fields.

> **Note:**
>
> - Ensure there are no spaces in the name of the recipient. This field does not support duplicate recipient names. If you try creating a recipient with a duplicate name, the creation of the share fails with the "Name is already used by an existing object" error.
> - A Sharing ID is a unique provider for your Autonomous Database. The Data Share tool uses it to share data with you. Copy the Sharing ID from the Consume Share page to the clipboard and paste it in the Sharing ID field of the **Create Share Recipient** wizard.

Click **Create** to finish creating the share recipient. The newly added recipient is displayed in the list of recipients. Click **Cancel** to cancel the ongoing process of recipient creation.

6. Select **Create** to create the share. You will receive a successful share notification message at the top right of the display that says the share is successful. The Provide Share page displays the share along with its details such as the entity type, owner, shared objects and the recipients.



Click the name of the entity or click **Action** to view details about the share entity. See View Share Entity Details.

After share is created, you can view the new Share entity in the Provide Share page.

View Share Entity Details

To view details about an entity, click the **Actions** icon at the right of the Share entity entry, then click **View Details**.

For all entities, the details include General, Selected Tables, Publishing details, Recipients, Versions, Log Details, Lineage and Impact sections.

General

The General tab displays the Name and description of the share.

Selected Tables

The Selected Tables tab displays the list of tables you wish to share.

Publishing

This tab displays the publishing details you configure during creation of the share. It displays whether you select live data share or share versioned data.

Recipients

You can review the list of recipients with whom you intend to share objects. It displays the name of the recipient with their mail address.

Jobs

See the Jobs option below to view information about it.

Log Details

You can view the detailed status of the share in this tab. It provides information such as, the time, Status, Log Level, and Details of each activity that takes place during the running of the job, i.e., creation of the share. It also displays if any errors are encountered during the creation of the share.

Lineage

Lineage displays all known information about the dependencies of the entity, and therefore how the entity was created and how it is linked to other entities.

For example, for a share you create in the database, the lineage is just the cloud location. For a share that you create by adding files from cloud storage, the lineage includes the cloud location of the file you share.

Pointing to the name of an item in the lineage displays the share name, the application that creates it, the type of entity, the path to it, and the schema it is in.

Arrows point from an entity to the entity that it derives from. For example, for a share you create, an arrow points from the cloud location to the share. If you point to an arrow, then a Links Information box appears that shows information about the relationship between the two entities.

To view more details about an item, click the **Actions** icon for the item, then click **Expand**. For a cloud location, the files in the location are displayed. Pointing to the name of the file displays the name, application, type, path, and schema of the file. To collapse the display, click the Actions icon, then click Collapse.

Impact

Impact shows all known information about the downstream use of an entity, and therefore how a change in the definition of an entity may affect other entities that depend on it.

For a specific share entity, you can perform the following actions using the Actions context menu:

- **View Details:** See View Share Entity Details.
- **Objects:** Opens Select Tables dialog of the specific share where you can remove tables from the list of Shared Tables list or add tables from the list of Available Tables list to the list of Shared Tables.
- **Recipients and Profiles:** Opens the Recipients and Profiles dialog which enables you to create a new share recipient. You can also revoke the recipient right from the share by clicking X next to the recipient name. The Recipients and Profiles dialog enables you to update the list of recipients, copy profile activation link, send an activation mail to the selected recipient.
- **Jobs:** Opens the Jobs table and displays information related to it. A job is created whenever something needs to happen to a share. This includes the following:
  - PUBLISH VERSIONED SHARE
  - PUBLISH LIVE SHARE
  - DROP SHARE VERSIONS

Only one job can run at a time on a single share. E.g. if the user calls `publish_share` while a `PUBLISH` job is already running, then the second job will have to wait. The `STATUS` column reflects the same.

The `JOB TYPE` column displays the method of parquet files generation (only for a versioned share). It's values can be `ODI` or `DBMS_CLOUD`.

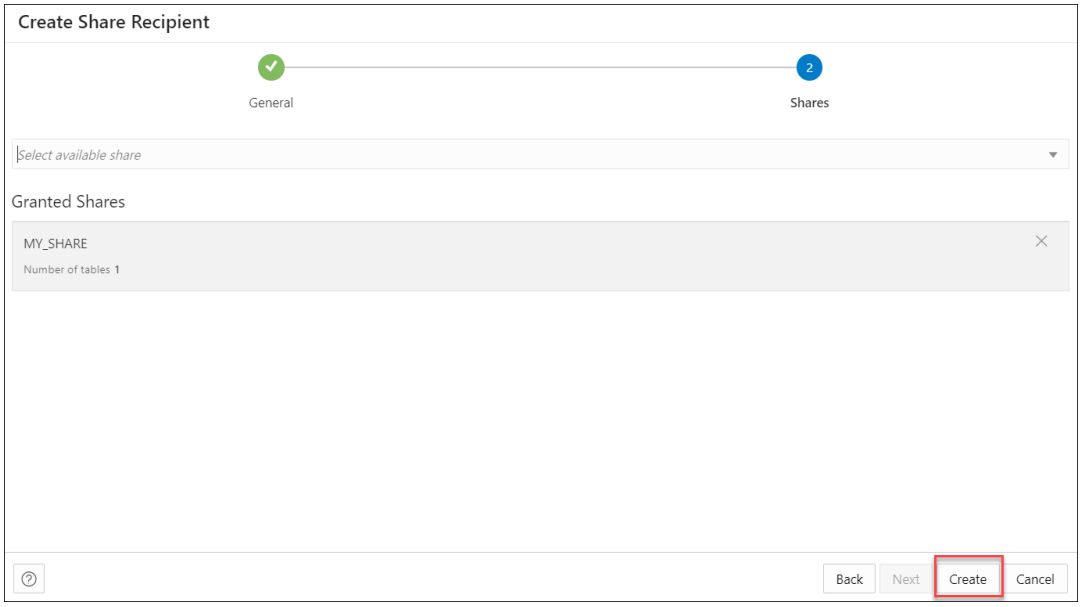
- **Publish:** Select **Publish** to publish the updated share.
- **Unpublish:** Select **Unpublish** to unpublish the share. This option prevents all recipients from accessing a single share, but lets them continue to access any other shares. You can restore access by publishing share again.
- **Delete:** Deletes the Share Provider Entity.

Create Share Recipient

On the Provide Shares page, click the **Recipients** widget and select **+ Create Recipient** to create a new Share Recipient.

Share Recipients are those with whom we intend to share objects with. You can share objects securely with users outside of Oracle workspace. Apart from creating recipients from the Create Share wizard, you can also create a share recipient from the Create Recipient icon on the Provide Share page and verify the details you configure while its creation.

1. On clicking the **+ Create Recipient** icon, a Create Share Recipient dialog box appears.



2. On the General tab, specify the following fields:
   - **Name**: Enter the name of the Share Recipient. For example, *Marketing_team.*

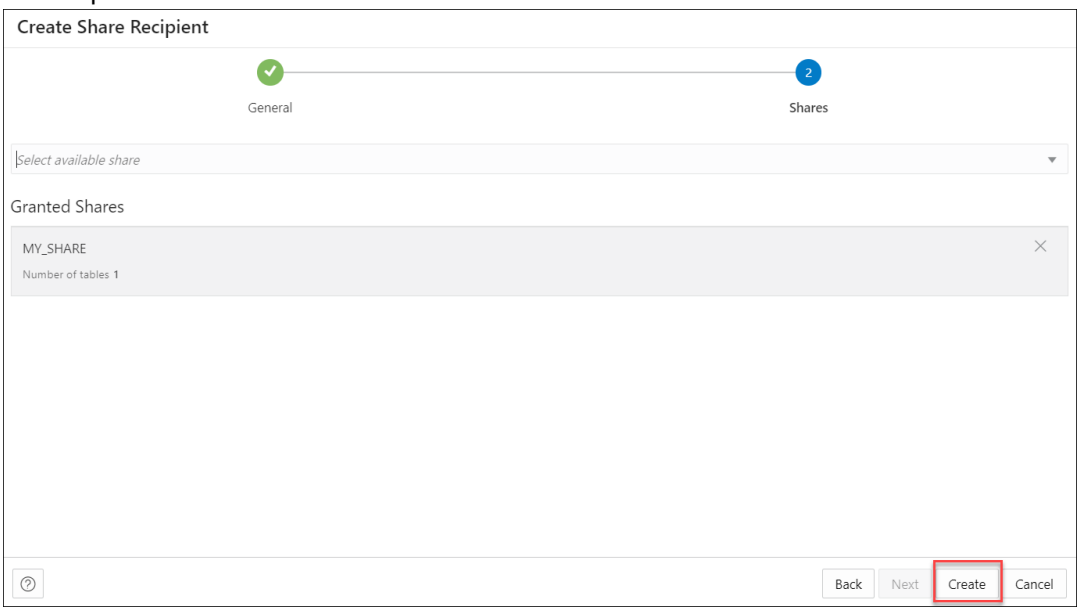
> **Note:**
>
> Ensure there are no spaces in the name of the recipient. This field does not support duplicate recipient names. If you try creating a recipient with a duplicate name, the creation of the share fails with the "Name is already used by an existing object" error.

   - **Email**: Enter the email ID of the recipient. For example, *marketing@oracle.com*.

- **Description**: Add a description. This is optional.
- **Sharing ID**: An intended Recipient should copy this ID from the consumer page and send it to the Producer to be used here. The Share Provider wont have access to the consumer page. **This field is required for Live Share Recipients only**.
  On the User token lifetime section, enter any value in the `TOKEN_LIFETIME` field for the share recipient. This property specifies for how long the generated token will be valid after which the recipient loses access to the data share and must request a new token. Enter 60 in the first text field. Select Minutes from the drop-down menu.

Click **Next** to progress to the Shares tab of the Create Share Recipient dialog box.

3. On the Shares tab, select an available share from the drop-down. As you select the share, it lists under the Granted Shares section. You have the choice to share multiple shares to the recipient.



Select **Create** to create the Share recipient.

Once you have created the Share recipient, you will be able to view the newly created Share Recipient entity in the display area of the Provide Share page. The new recipient card displays details such as the entity type and the owner.



Click the name of the entity or click **Action** to view details about the share recipient entity. See View Share Recipient Entity Details.

View Share Recipient Entity Details

Click the **Actions** icon at the right of the Share Recipient entity entry, then click **View Details**.

For all entities, the details include general details, Log details, Granted Shares, Lineage and Impact.

General

The General tab displays the Name and Email address of the Share Recipient.

Granted Shares

The Granted Shares tab displays the list of shares granted to the recipient.

Log Details

The Log details tab displays the Time, Status, Log Level, Share Name, Request, End Point, Remote address, User Agent and Details of the Recipient.

Lineage

Lineage displays all known information about the dependencies of the entity, and therefore how the entity was created and how it is linked to other entities.

For example, for a share you publish, the lineage includes the files in cloud storage along with the cloud location of the files you share.

To view more details about an item, click the **Actions** icon for the item, then click **Expand**. For a cloud location, the files in the location are displayed. Pointing to the name of the file displays the name, application, type, path, and schema of the file. To collapse the display, click the **Actions** icon, then click **Collapse**.

Impact

Impact shows all known information about the downstream use of an entity, and therefore how a change in the definition of an entity may affect other entities that depend on it.

Click **Actions** in the Share Recipient entity card to open the context menu. The actions available are:

- **View Details:** See View Share Recipient Entity Details.
- **Granted Share:** Opens Manage Granted Shares dialog where you can grant access to a share from list of available shares or you can revoke access from the current selection of share.
- **Get Bearer Token:** Opens an Auth Token dialog to authenticate the recipient. Use this option when your token expires. This enables you to generate a new token. A Bearer Token is part of the JSON file.
- **Send Activation Mail:** Opens the mail with the profile link.
- **Copy Profile Activation Link to Clipboard:** Copies the JSON profile activation link to clipboard.
- **Rename:** Select this option to rename the Recipient.
- **Delete:** Select this option to delete the recipient you created entirely. By doing so, access to all shares will end instantly and any existing delta credentials will become invalid.

Consume Share

Once the providers share the objects, there are a few steps the recipients need to follow to consume the share.

Use the Consume Share page to perform the following operations:

- Consume Share Overview:
  To consume data shares, you need to subscribe to them and create views of tables included in the live share.
- Consume Versioned Share:
  As a recipient, you will need to download your share profile, subscribe to the data share provider, register the shares and create external tables on top of your shares. The Data Share tool authorizes the access using the JSON profile sent to the recipient using the activation mail. After the access is granted, the Data Share tool links the shared objects with the Data link tool where the consumer can run the data link job and access the objects shared by the provider.
- Consume Live Share:
  This lets you as a recipient to consume live data from the database.
- View Share Provider Entity Details:
  Use the **Actions** icon at the right of the live share or delta share provider entity entry to view details about the live share or delta share provider entity you create.
- Consume Share Overview
  The Consume Share provides an overview on the list of share providers, search for share providers and add a share provider.
- Consume Versioned Share
  You must follow these steps to make shared versioned data available to you within Oracle Autonomous Database. Data shared with you through Delta Sharing is not automatically available and discoverable in your Autonomous Database.
- Consume Live Share
  Live data shared with you through data sharing is not automatically available for consumption.
- View Share Provider Entity Details
  To view details about the Share Provider entity, click the **Actions** icon at the right of the Share Provider entity entry, then click **View Details**.

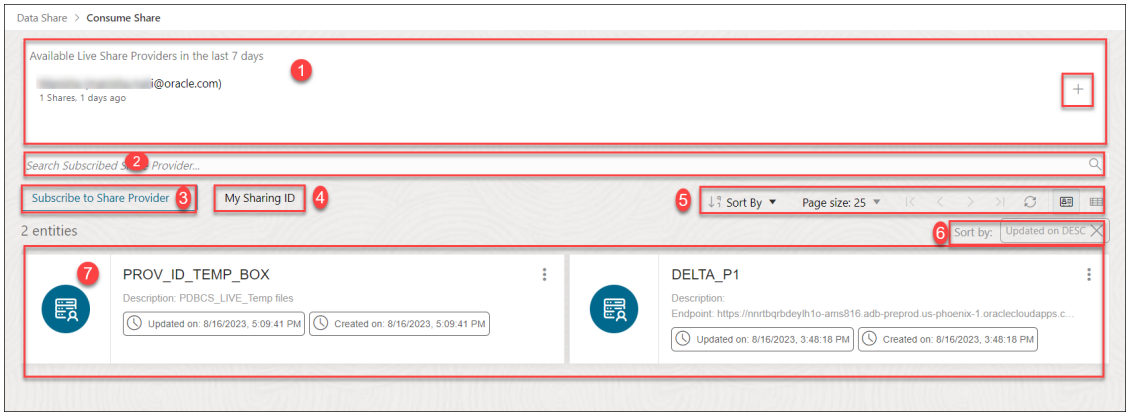
Consume Share Overview

The Consume Share provides an overview on the list of share providers, search for share providers and add a share provider.

To navigate to the Consume Share page, do either of the following:

- On the Data Studio menu, select **Consume Share** under the Data Share menu.
- On the Data Share page, click the **Consume Share** widget present in the Provider and Consumer section.

The Consume Share page contains:



1. **Available Live Share Providers in the last 7 Days**
   This area displays the list of available Live Share Providers you create in the last 7 days. You can update any of the fields as per your wish. You can also subscribe to Live Share using the + sign to the right of Shares listed in **Available shares in the last 7 days**.
2. **Search Subscribed Share Provider** field

   You can search for the Share Recipient you create by entering the name of the Subscribed Share Provider. Enter the name of the Subscribed Share Provider, for example, *REVIEW_PROVIDER* and click the **magnifier** icon to complete the search. The Share tool displays the search results in the display area.
3. Select **+ Subscribe to Share Provider** to subscribe to a new Share Provider. See Subscribe to Share Provider for exploring this icon.
4. **My Sharing ID**
   A Sharing ID is a unique provider for your Autonomous Database. Copy this ID to the clipboard and paste it in the Sharing ID field of the **Create Share Recipient** wizard. This enables a Live Share to be shared with a share provider.
5. **Toolbar**
   The toolbar consists of the sort by, page size, refresh and entity view options.
6. **Sort by settings**
   When you set sorting values by using the Sort By control in the toolbar, the settings are displayed in small boxes beneath the toolbar. You can delete a setting by clicking the **X** icon in the box. Or you can change the settings by returning to the Sort By control in the toolbar.
7. **Display area**

   The area beneath the Search Consumer Share Providers field displays the entities returned by a search and that match the filter criteria set in the Filters panel. You can sort the entities by clicking the **Sort By** button and then setting sort values.

Consume Versioned Share

You must follow these steps to make shared versioned data available to you within Oracle Autonomous Database. Data shared with you through Delta Sharing is not automatically available and discoverable in your Autonomous Database.

You need to perform several basic steps to subscribe and access the provided data shares. Here are the steps you need to perform:

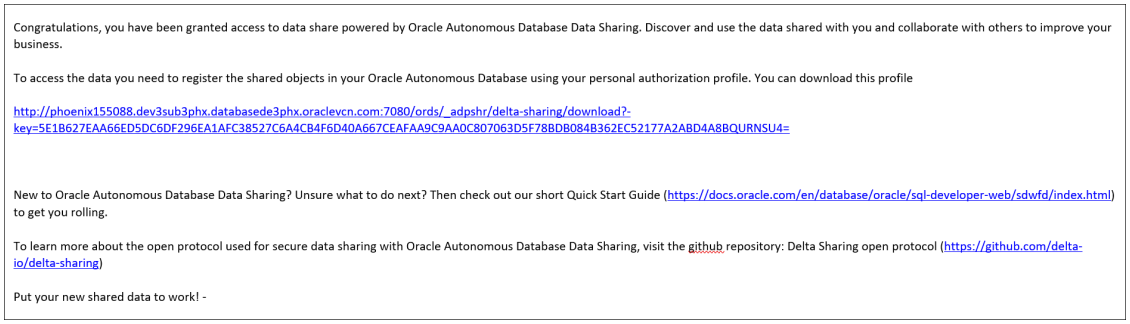
- Download the JSON profile.
- Subscribe to the data share provider.

- Register shares made available to you.
- Create external tables on top of your shares.

About JSON Profile

Profile files are JSON files containing a user's credentials to access a Delta Sharing Server. This enables you to authenticate yourself with the delta sharing server and to discover the data shares you are eligible to access. Download the profile using the URL included in the invitation email sent by the data share provider.

Once the Share tool creates a share object, the recipient receives an activation mail with subject titled *Oracle Autonomous Database Data Share*.



Congratulations, you have been granted access to data share powered by Oracle Autonomous Database Data Sharing. Discover and use the data shared with you and collaborate with others to improve your business.

To access the data you need to register the shared objects in your Oracle Autonomous Database using your personal authorization profile. You can download this profile

http://phoenix155088.dev3sub3phx.databasede3phx.oraclevcn.com:7080/ords/_adpshr/delta-sharing/download?-key=5E1B627EAA66ED5DC6DF296EA1AFC38527C6A4CB4F6D40A667CEAFAA9C9AA0C807063D5F78BDB084B362EC52177A2ABD4A8BQURNSU4=

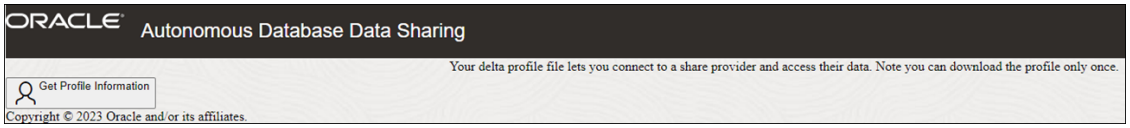
New to Oracle Autonomous Database Data Sharing? Unsure what to do next? Then check out our short Quick Start Guide (https://docs.oracle.com/en/database/oracle/sql-developer-web/sdwfd/index.html) to get you rolling.

To learn more about the open protocol used for secure data sharing with Oracle Autonomous Database Data Sharing, visit the github repository: Delta Sharing open protocol (https://github.com/delta-io/delta-sharing)

Put your new shared data to work! -

To access the share, you need to register the shared objects using personal authorization JSON profile.

You can click on the profile link to download the JSON profile. Clicking the profile link takes you to a new screen in the browser with a **Get Profile Information** button as shown below:



Select **Get Profile Information** to download the JSON profile to connect to the share provider.

> **Note:**
>
> You can click on the **Get Profile Information** button only once and also view the link to the Quick Start Guide. The share tool does not allow you to select **Get Profile Information** twice. Clicking on it twice brings up a screen which displays the list of causes of the failure to download the profile.

The below profile is an example of the JSON profile you download.

```
{  "
        shareCredentialsVersion ": 1,
        "endpoint": "https://myhost.us.example.com/ords/prov/
_delta_sharing/",
        "tokenEndpoint": "http://myhost.us.example.com:1234/ords/pdbdba/oauth/
token",
        "bearerToken": "-xxxxxxxxxxxxxxxxxxxx",
        "expirationTime": "2023-01-13T07:53:11.073Z",
```

```
"clientID": "xxxxxxxxxxxxxxxxxxxxx..",
"clientSecret": "xxxxxxxxxxxxxxxxxxxx.."}
```

The profile stores the credentials in an encrypted format. The parameters with their description are:

- **shareCredentialsVersion:** The version of the share you publish.
- **endpoint:** Specifies the share endpoint.
- **tokenEndpoint:** Specifies the token endpoint. The Share tool client use the token endpoint to refresh the timeout on your bearer token if you are consuming the share using Oracle.
- **bearerToken:** This is a cryptic string which the authentication server generates in response to a login request.
- **expirationTime:** This is the time taken for the authentication to expire.
- **ClientID:** Specifies the public identifier the authentication server generates when you register the instance for authentication.
- **clientSecret:** Specifies a secret identifier the authentication server generates for authorization.

Copy the JSON content of the profile in a notepad. You will need this JSON below to subscribe your share provider.

> **Note:**
>
> Ensure you copy the full JSON content profile, including the left brace and the right brace.

Security enhancements

As a share recipient, you must set up an access control list (ACL) to the share provider's machine by using the `APPEND_HOST_ACE` procedure as an ADMIN user, or another privileged user. This allows you to access the share via the Internet.

> **Note:**
>
> This must be done before using the Add Share Provider Wizard to add an access control entry (ACE) to the access control list (ACL) of the host (i.e. Share provider). You can find the host name from the JSON profile you downloaded in the previous step.

For example, if you wish to allow a database user, A_SHARE_USER to access the endpoints on a host (Share provider) named, here is a sample of PL/SQL procedure you will need to run in the SQL worksheet editor as an admin. As a prerequisite, extract the host name from the endpoint property in the delta sharing JSON profile, as provided in the example above. The hostname from the example is `myhost.us.example.com`.

```
BEGIN
        dbms_network_acl_admin.append_host_ace(
        host =>'myhost.us.example.com',
```

```
lower_port=>443,
upper_port=>443,
ace => xs$ace_type(
privilege_list => xs$name_list('http', 'http_proxy'),
principal_name =>'A_SHARE_USER',
principal_type => xs_acl.ptype_db));
COMMIT;
END;
/
```

Following are the parameters with their description:

- **host**- Specifies the name or the IP address of the host. The host or domain name is not case-sensitive.
- **lower port**- Specifies the lower port of an optional TCP port range.
- **upper port**- Specifies the upper port of an optional TCP port range.
- **ace** – The Access Control Entry.
- **privilege list**- Specifies the list of network privileges to be granted or denied.
- **principal_name**- It is the principal (database user or role) to whom the privilege is granted or denied. It is case-sensitive.
- **principal_type**- Specifies the type of principal you use.

Refer to the PL/SQL Packages and Types Reference document for more details on the `DBMS_NETWORK_ACL_ADMIN` package subprograms.

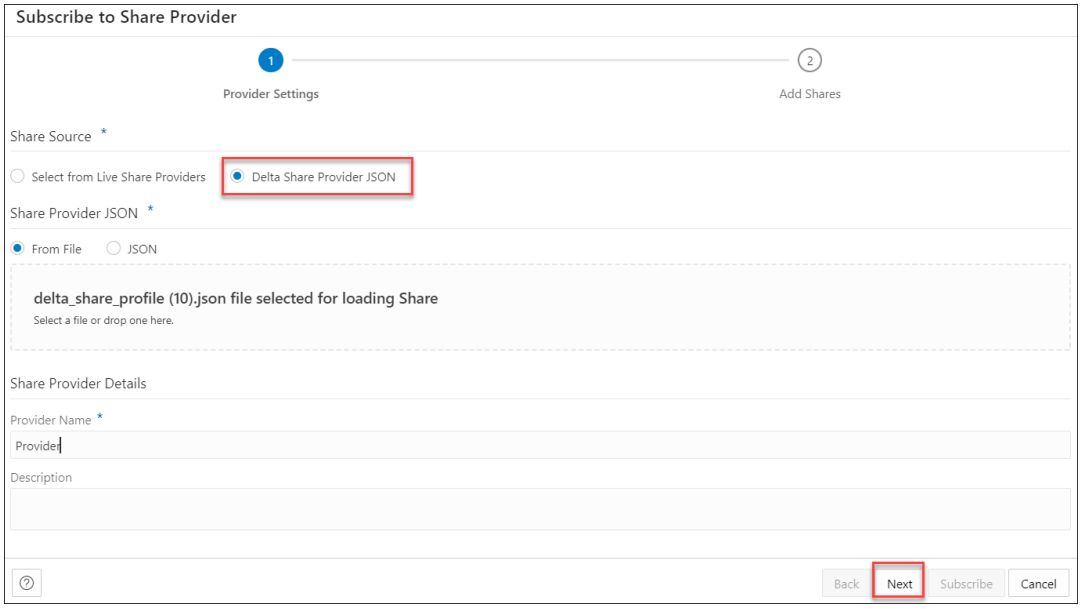
Grant an ACL to the user on the local ORDS endpoint. You will need this to generate bearer tokens on locally created shares.

```
PRIV_ORDS_ACL              CONSTANT PLS_INTEGER := 8;
```

In this process you will load the provider's JSON profile for configuration and credentials to enable access to the recipients.

1. Open the Consume Share page and click **+ Subscribe to Share Provider** to select **Subscribe to Delta Share Provider** from the drop-down. This opens the Subscribe to Share Provider dialog box.

2. On the Provider Settings pane of the Register Share Provider dialog box, specify the following details:



- **Provider Name**: *Provider*.
- **Description**: Add a description. This field is optional.

Under Share Source section, choose **Delta Share Provider JSON**.

Under Share Provider JSON field, you can share the JSON profile in the following ways:

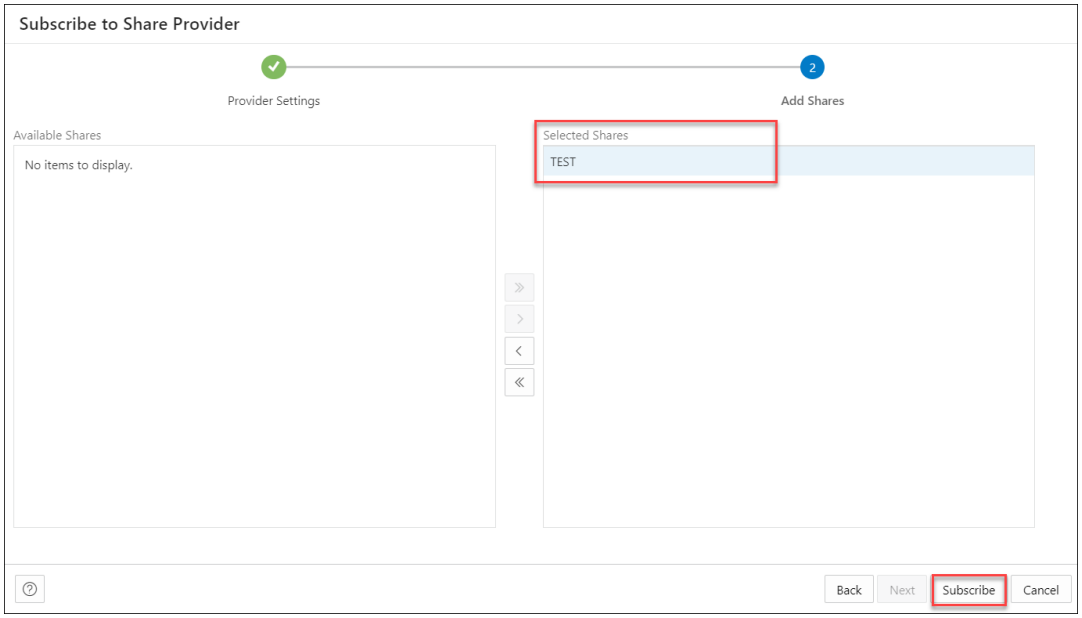
- **From File**: Select this option and click on the drop area titled "Delta Share Profile JSON". Clicking on the area opens your local repository where you can select the JSON profile you had downloaded.
- **JSON**: You can select this option and paste the JSON content of the profile you copy into the notepad.

Upload the JSON profile file and create a share provider subscription.

Click **Next** to progress to the Add Shares tab.

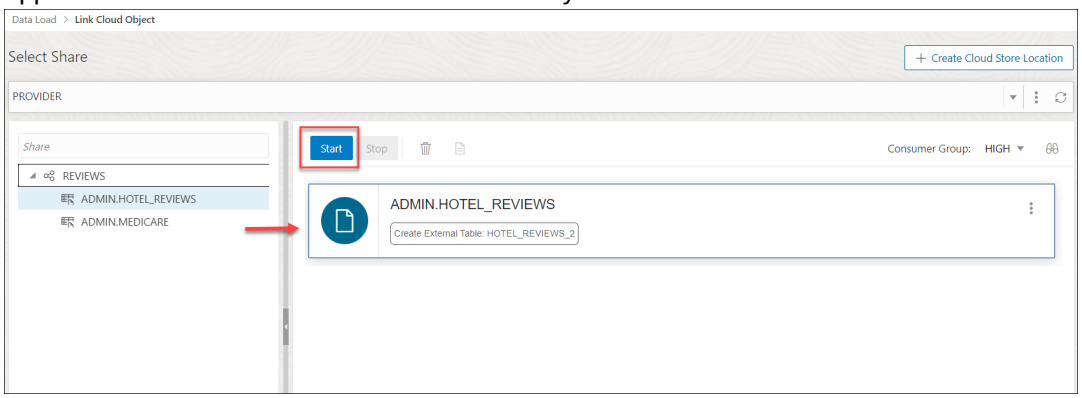
3. On the Add Shares tab of the dialog, you will view the list of available shares. Click the share you wish to consume from the Available Shares and select any of the available

options:



- **>**: This option enables you to move the Available Share to Selected Shares.
- **<**: Select this option to remove the selected share from Selected Shares.
- **>>**: This option allows you to move all the shares to the Selected Shares screen.
- **<<**: Select this option to remove all the selected shares from Selected Shares.

4. Click **Subscribe** to add the share. A confirmation prompt appears when the provider is created successfully. After successful creation of the provider, you will now view the Link Cloud Object screen of the Data Load page.
5. You can view the name of the Share provider in the cloud storage location field. The share appears in the source file location with the files you add to the share.

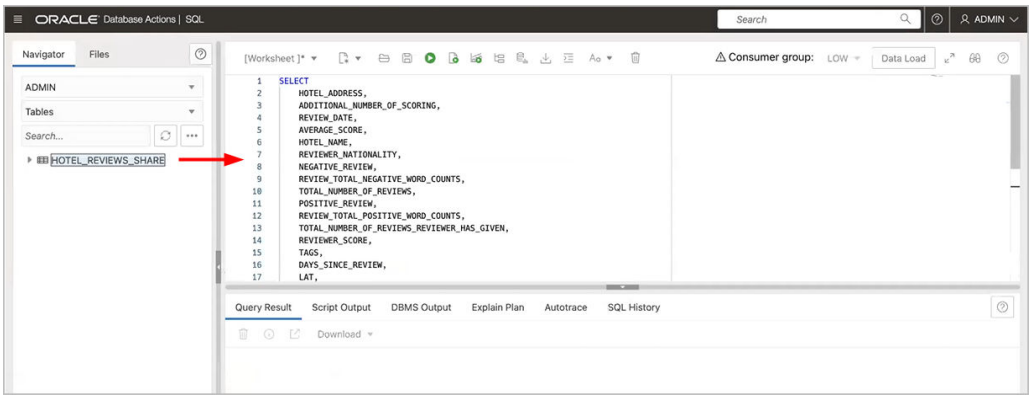


Expand the Share folder cart, drag and drop the file you share from the source to the Data Link cart.

Select **Start** in the Data link cart to run the data link job.

6. View the created tables from Database Actions.
   - Click on Database Actions, in the breadcrumb, to go back to the Database Actions launchpad.
   - Click on the SQL tile.
   - Select the external table, drag and drop it into the worksheet.

The SQL Select statement for the table appears. This SQL statement can be run to consume the shared data.

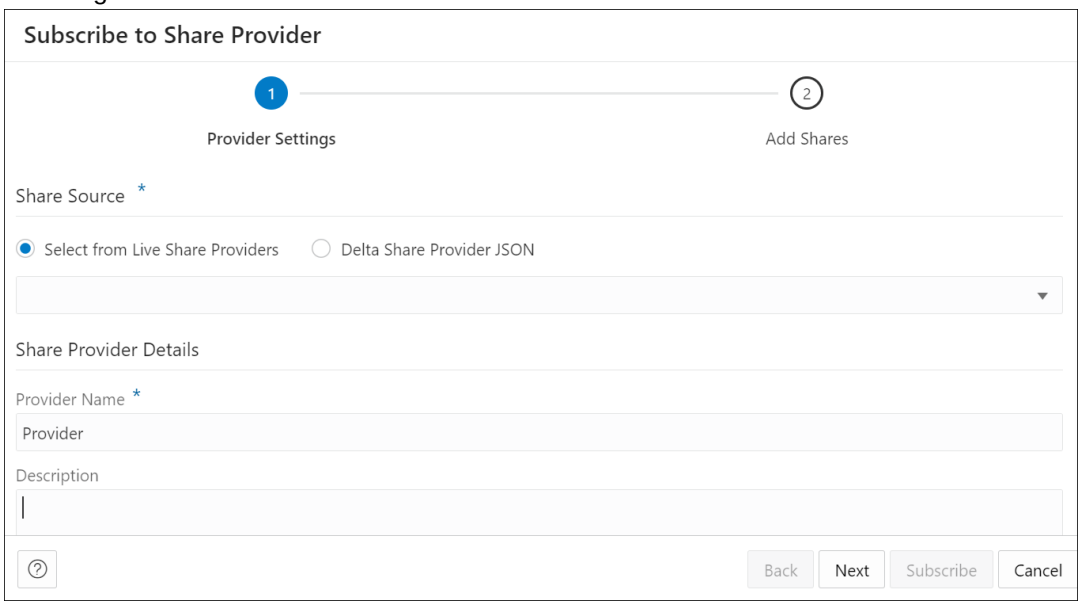


Consume Live Share

Live data shared with you through data sharing is not automatically available for consumption.

To consume live data shares, you need to subscribe to them and create views of tables included in the live share. The views can be queried using SQL scripts.

1. Open the Consume Share page and click **+ Subscribe to Share Provider** to select **Subscribe to Live Share Provider** from the drop-down. This opens the Subscribe to Share Provider dialog box.
2. On the Provider Settings pane of the Subscribe to Share Provider dialog box, specify the following details:



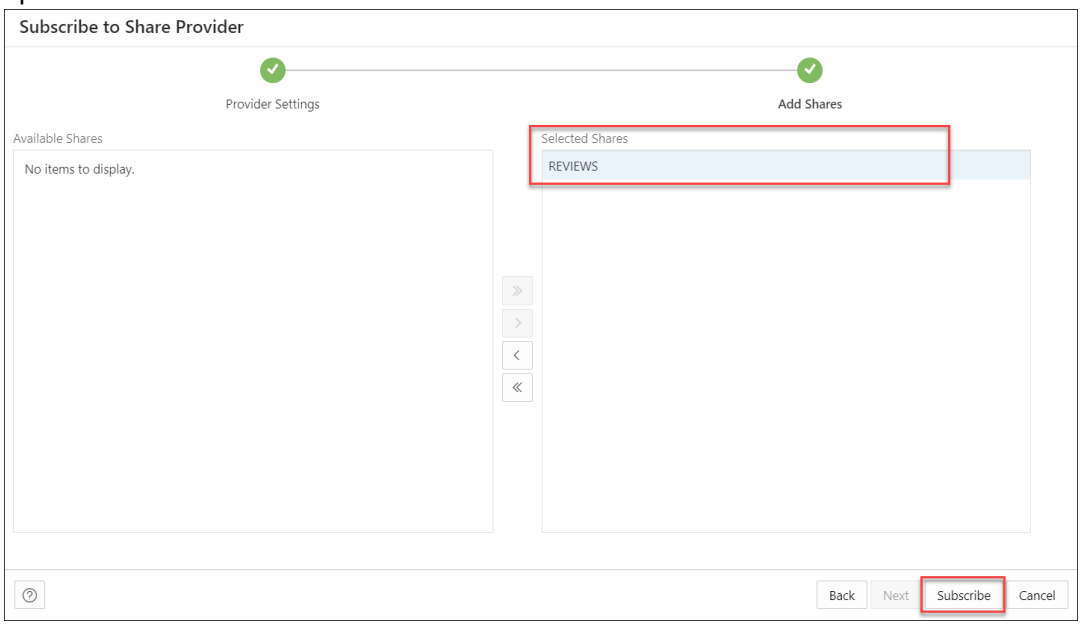
Under Share Source section, choose **Select from Live Share Providers** and select the provider from the drop-down.

Under Share Provider Details field, enter the following:

- **Provider Name**: Specify the name of the provider.
- **Description**: Enter a description of the Provider.

Click **Next** to progress to the Add Shares tab.

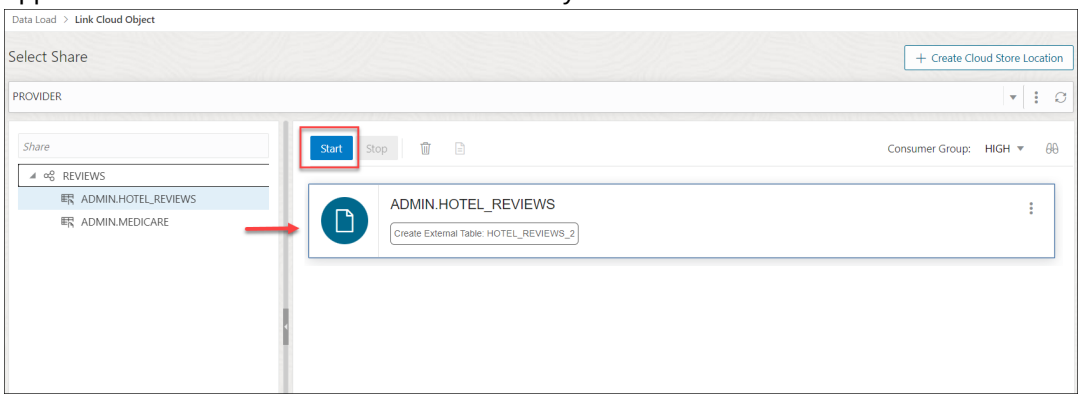
3. On the Add Shares tab of the dialog, you will view the list of available shares. Click the share you wish to consume from the Available Shares and select any of the available options:



- **>**: This option enables you to move the Available Share to Selected Shares.
- **<**: Select this option to remove the selected share from Selected Shares.
- **>>**: This option allows you to move all the shares to the Selected Shares screen.
- **<<**: Select this option to remove all the selected shares from Selected Shares.

Click **Subscribe** to add the share. A confirmation prompt appears when the provider is created successfully. After successful creation of the provider, you will now view the Link Cloud Object screen of the Data Load page.

4. You can view the name of the Share provider in the cloud storage location field. The share appears in the source file location with the files you add to the share.



Expand the Share folder cart, drag and drop the file you share from the source to the Data Link cart.

Select **Start** in the Data link cart to run the data link job.

View Share Provider Entity Details

To view details about the Share Provider entity, click the **Actions** icon at the right of the Share Provider entity entry, then click **View Details**.

For all entities, the details include Lineage and Impact sections.

For a specific Share Provider entity, you can perform the following actions using the **Actions** context menu.

- **View Details:** See View Share Provider Entity Details.
- **Manage Shares:** Opens a Manage Shares for Share provider dialog box. This lists the shares you select to share with the recipient. You can edit the list of shares you wish to share with the recipient. Click **OK** to save any changes or select **Cancel** to discard the process of editing.
- **Rename:** Renames the Provider Name. Select **Yes** to make changes else click **No**.
- **Delete:** Removes the Share Provider Entity.
- **Load Tables:** You are directed to the Load Data page with the Share tab selected.
- **Link Tables:** You are directed to the Link Data page with the Share tab selected to view and run the related link object. Drag and drop the shared data to add it to the data link job.

Data Share Limitations on Autonomous Database

This section summarizes the limitations of Data Share (both Versioned and Live Share) included in the Oracle Autonomous Database.

Oracle Database Release Restriction

Only **Oracle Database 19c** supports both Versioned Share and Live Share. You cannot use Versioned Share and Live Share with any other database version.

Supported Column Types

Only the following list of column types are supported:

- `BINARY_DOUBLE`
- `BINARY_FLOAT`
- `CHAR`
- `CLOB`
- `DATE`
- `FLOAT`
- `INTERVAL DAY TO SECOND`
- `INTERVAL YEAR TO MONTH`
- `NCHAR`
- `NCLOB`
- `NUMBER`
- `NVARCHAR2`
- `TIMESTAMP`
- `VARCHAR2`

Live Share Data Provider

The following are limitations for using the Data Share tool as a Data Share provider:

1. A live share data provider can create shares with a maximum of four objects.

2. If the live share data provider is sharing multiple tables which would normally be joined together, it is recommended that the producer create a view that performs the joins and then only shares the view.

Troubleshooting Errors with the Data Share Tool

This chapter provides instructions on what issues cause slow performance and how to troubleshoot the issues.

With an increase in ECPUs, does the performance of the Live share feature in the Data Share tool improve while running complex queries?

No, it does not. In such cases, you need to configure the **Degree Of Parallelism (DOP)** using Data Studio to improve the performance of complex running queries.

Configure the Degree Of Parallelism (DOP) using Data Studio

When the Live Share feature of the Database Actions instance runs on an increased value of ECPU, by default the performance of the complex running queries does not improve.

You can use the DOP option to improve the performance of the complex running queries. The DOP sets the number of processors employed to run a single statement, for each parallel plan execution. This section describes how to configure the `degree of parallelism` (DOP) configuration option in the Autonomous Database by using the Live Share feature of Data Studio.

In this example, you will use the Live share feature and share tables of an uncompressed size of 1TB.

Before you begin, here are a few things you need to consider:

- The option to configure the value of DOP can only be done by an `admin`.

Prerequisite:

- Log in to your Database Actions instance as an admin and under **Administration**, select **Set Resource Management Rules**.
- Note the default **concurrency value** and **DOP**.

To configure the degree of parallelism using Data Studio to improve the performance of complex running queries, you will:

- Create Share as a provider.
- Consume Share as a consumer and create external queries on top of your shares to record the query execution time.
- Increase the ECPUs to a value greater than the current value as provider and consumer and record the query execution time.
- Configure the `DOP` value for the `MEDIUM` consumer group.
- Consume the same share and record the query execution time of the same table.

> **Note:**
>
> You must configure the `ECPU` count and `DOP` value as a **provider** with `ADMIN` privileges.

On the Database Actions launchpad, click **Selector** and select the **Data Share** tool from the **Data Studio** suite of tools.

1. **Create Share as a provider.**

   a. Click **Provide Shares** and select **+Create Share**.

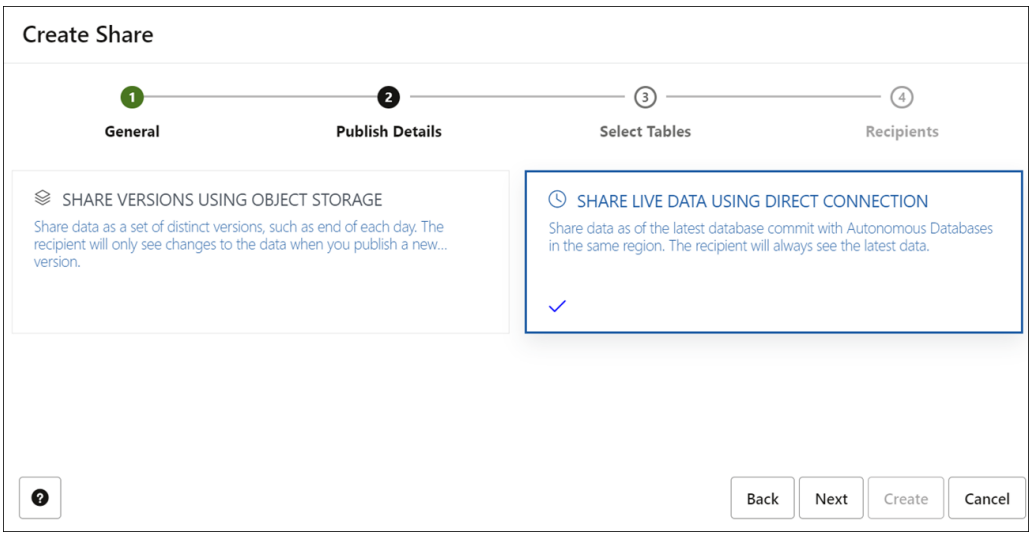
   b. On the Create Share wizard, in the **Name** field of the General tab, enter a name for the Share. For example: *TEST*.



   In the **Description** field, enter a description for the data you share. This field is optional.

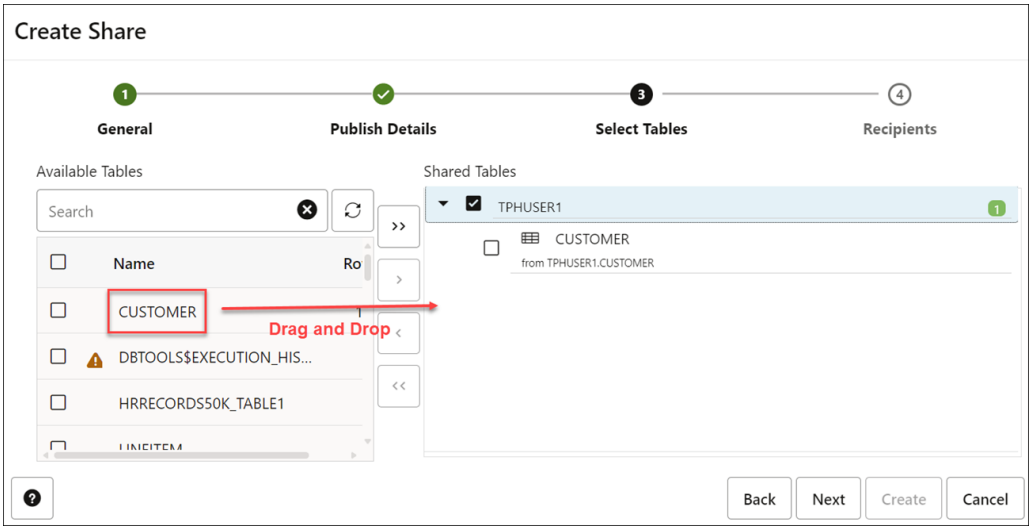
   Select **Next** to progress to the Publish Details tab.

   c. In the Publish Details Tables tab of the wizard, select **SHARE LIVE DATA USING DIRECT CONNECTION**.



   Click **Next**.

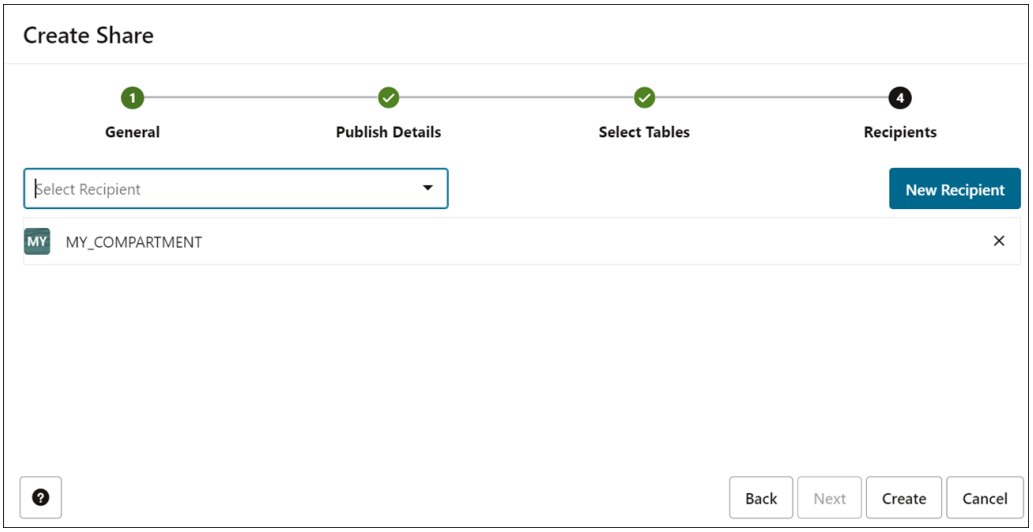
   d. On the **Select Tables** tab of the wizard, select the schema from the drop-down menu, drag and drop the `CUSTOMER` table from the Available Tables.



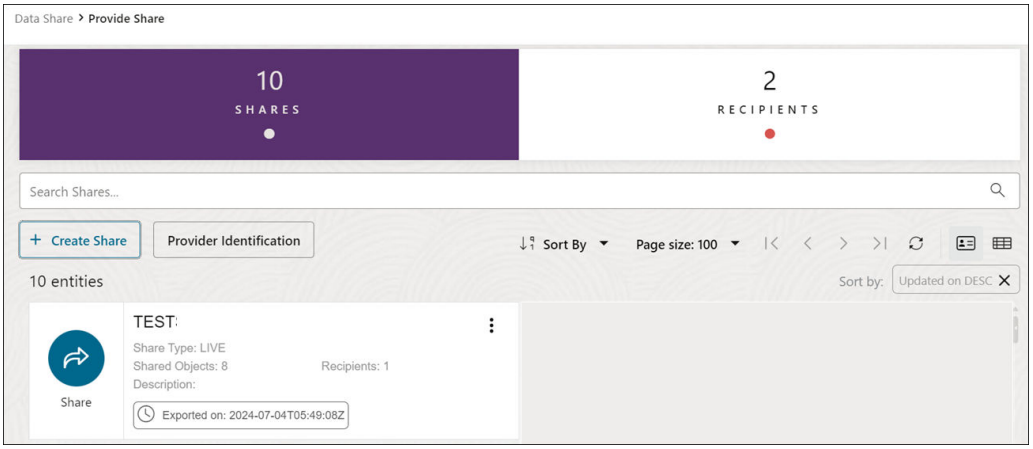
> **Note:**
>
> In this example, we are sharing a single table named `Customer` of large size.

Click **Next** to proceed to the Recipients tab of the Create Share Wizard.

e. On the Recipients tab of the Create Share wizard, select `MY_COMPARTMENT` from the drop-down.



f. Select **Create** to create the share.



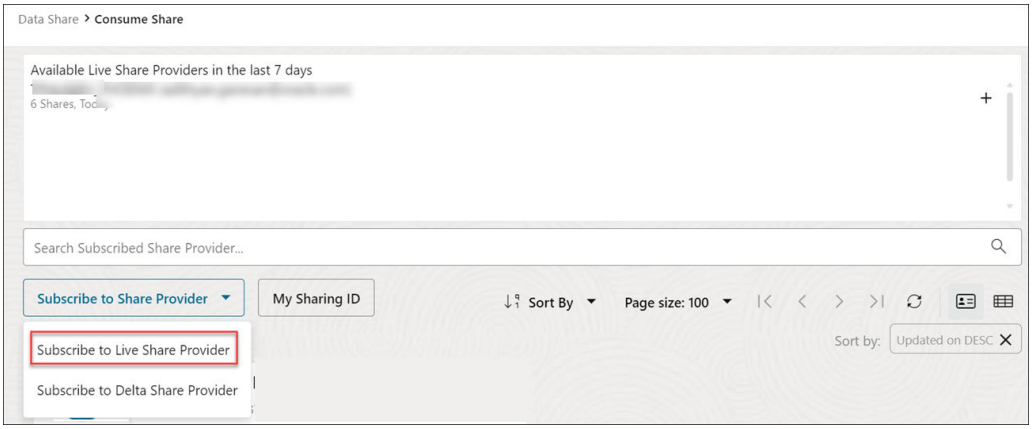
2. **Consume Share as a consumer and create external queries on top of your shares to record the query execution time.**

   a. After you have created the share as a provider, you will consume the data shared by the provider.
   See Consume Share for more details.

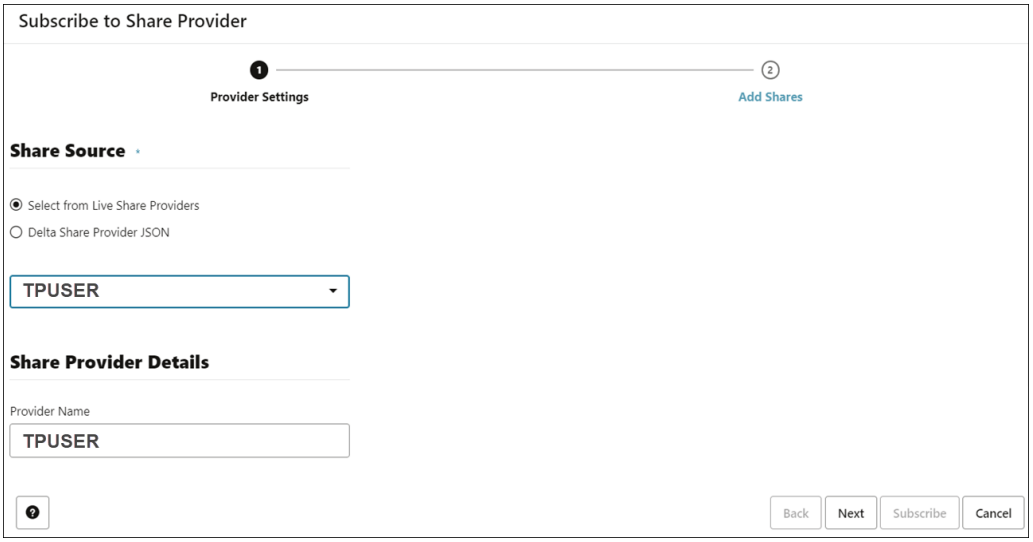
   b. Log in as a **consumer** to subscribe to the data share provider and access the data in your share.

   c. On the **Provider and Consumer** page of the Data Share tool, click the **CONSUME SHARE** tile to display the **Consume Share** page.

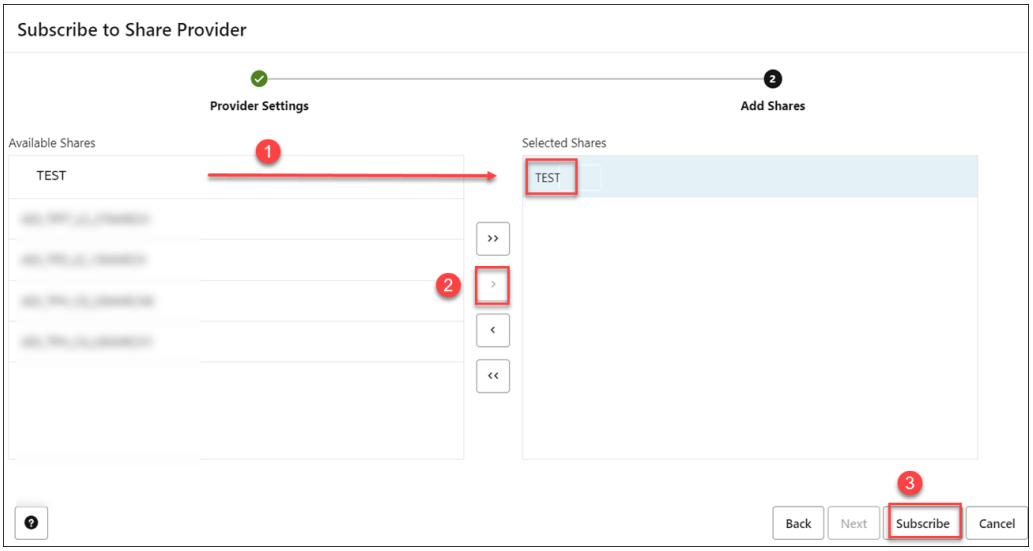
   d. Click the **Subscribe to Share Provider** drop-down list, and then click the **Subscribe to Live Share Provider** option.



   e. Under Share Provider Details, select the **Provider Name**. In this example, `TPUSER` is the name of the provider.

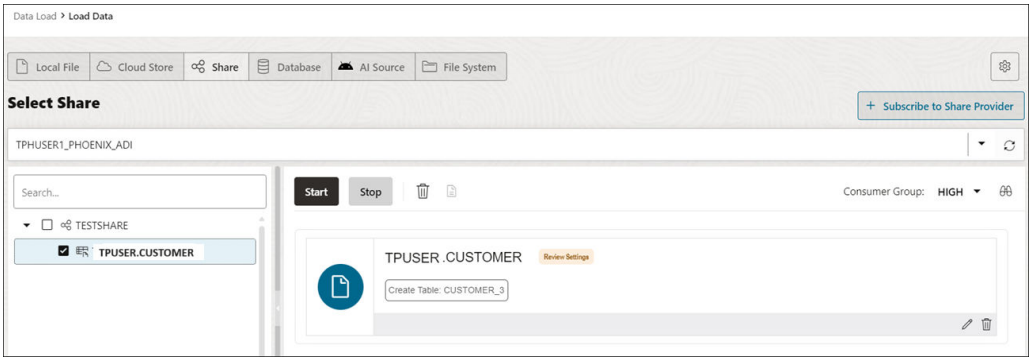


f. On the **Add Shares** panel of the Subscribe to Share **Provider Settings** wizard page, click the **TEST** in the **Available Shares** section, and then click the **(>)** icon.



g. The tool adds the share to the **Selected Shares** section. Click **Subscribe**.

h. A `TPUSER` share provider was created successfully informational message is displayed. You are directed to the **Load Data** page with the **Share** tab selected.

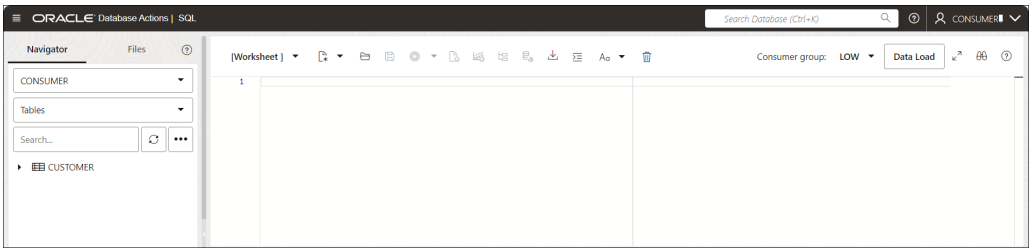


i. Expand the `TESTSHARE` data share node to display the objects in it. Drag and drop the `TPUSER.CUSTOMER` shared table to add it to the data load job.

j. The shared table is added to the data link job. When you run this job, a new `CUSTOMER` view will be created on top of the `CUSTOMER` table that was shared with you.

k. Click **Start**. A **Start Link from Cloud Store** dialog box is displayed. Click **Run**. When the job completes successfully, a link icon is displayed. The `CUSTOMER` view is created.

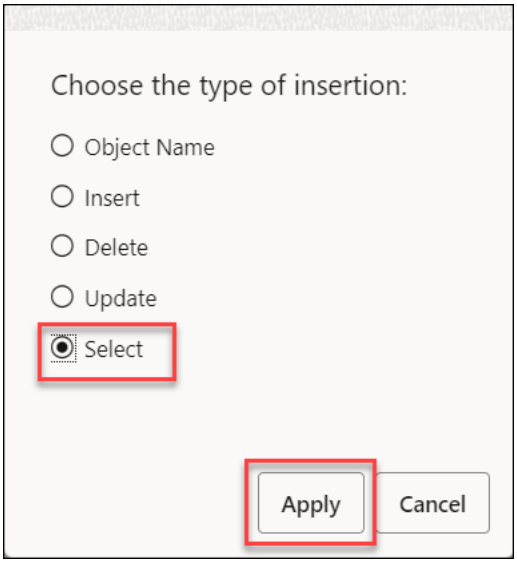
l. Click the **Database Actions | Data Share** banner. On the **Database Actions | Launchpad**, click the **Development** tab, and then click the **SQL** tab to display the SQL Worksheet.

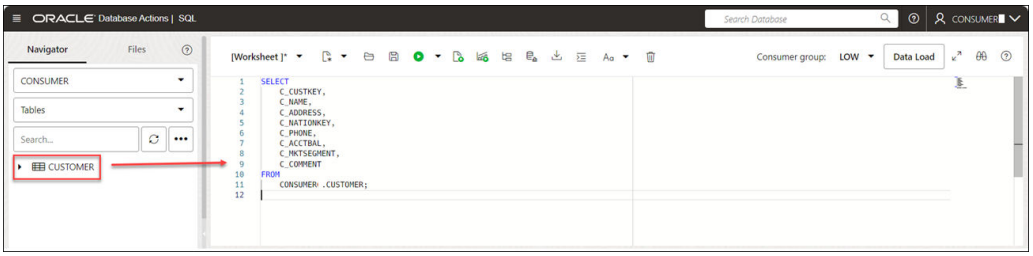


m. Drag and drop the `CUSTOMER` table to the SQL worksheet area.

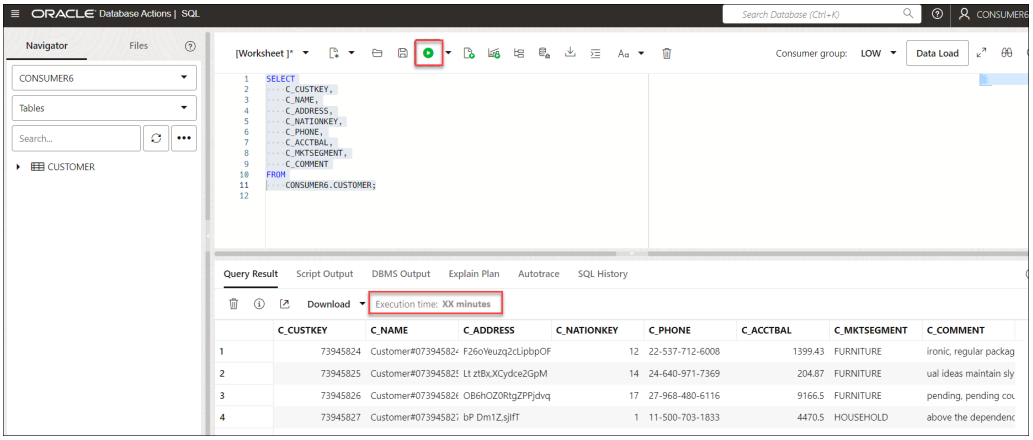
n. A dialog displays offering the types of available insertions.

o. Click **Select**, then **Apply**.





p. Click **Run**. You can view the automatically generated query displayed in the Worksheet.



q. Note the query execution time in the **Execution time** field displayed under Query Result tab.

3. **Increase the ECPUs as provider and consumer and record the query execution time.**

a. Increase the ECPUs of your Autonomous Database instance as a **provider** and a **consumer** from your OCI console. See Add CPU or Storage Resources or Enable Auto Scaling.

> **Note:**
>
> Make sure you have not enabled Auto Scaling.

In this example we have increased the ECPU count to **32**. You can upscale the ECPU count to any value of your choice. For more details on the maximum ECPU count you set and licensing see, View and Update Your License and Oracle Database Edition on Autonomous Database (ECPU Compute Model)

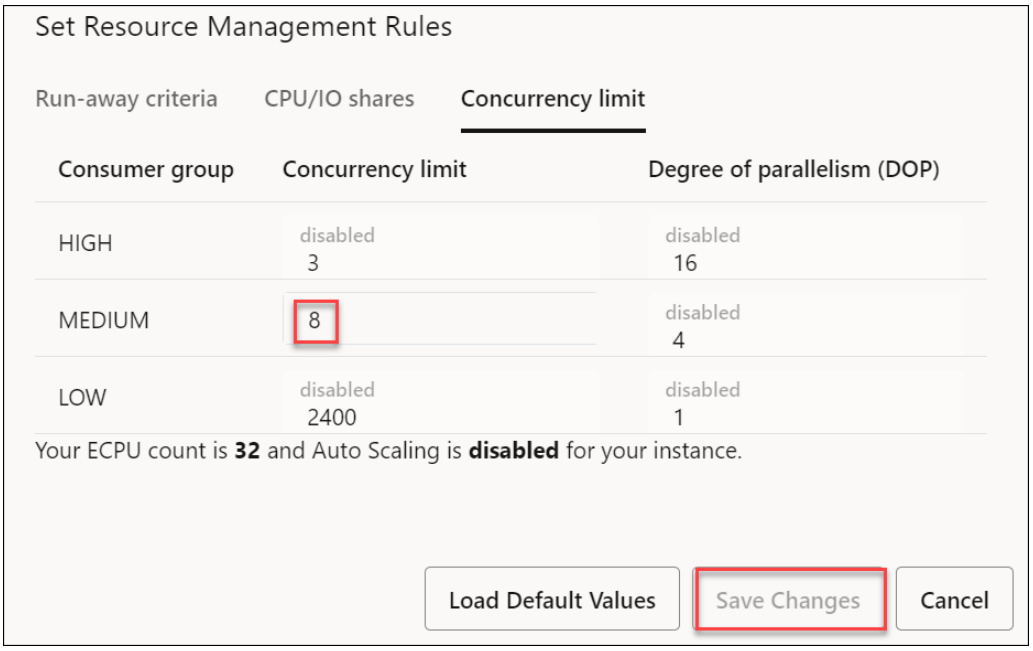
b. Repeat consuming the same `CUSTOMER` share and calculating the query execution time. There will be no change in the time to query the table.

4. **Configure the DOP value for the `MEDIUM` consumer group.**

   a. Log in to your database actions instance as a **consumer** with `ADMIN` privilege.

   b. Click **Selector** and under **Administration**, select **Set Resource Management Rules**.

   c. Decrease the **Concurrency limit** of the `MEDIUM` consumer group to **8** and click **Save Changes**. Decreasing the concurrency limit increases the DOP.



> **Note:**
>
> The concurrency limit you set in this step should be lesser than the default value you noted in the prerequisite.

5. **Follow up step- After you configure the degree of parallelism option, repeat the process of consuming the share, creating external queries on top of your shares to record the query execution time.**
   You will see the value of query execution time has decreased.

The improvement in the performance of the same complex SQL queries happens immediately without restarting the server.

FAQs on the Data Share Tool

The following are frequently asked questions on the Data Share tool.

1. When do recipients receive the Data Definition Language (DDL) updates to the source tables?
   The recipients need to link objects again from the cloud storage location.
   - Navigate to the link cloud Object page from the Data Studio menu.
   - Select the table in the Data Link cart.
   - Click on the **Actions** icon on the table and select Settings.
     This opens the Link Data from Cloud Store Location dialog.
   - Click the **Settings** pane and select "Drop Table and Create New External Table" option from the drop-down.

   You will receive the updated source table.
2. When do recipients receive the Data Manipulation Language (DML) updates to the source tables?
   The provider selects the option to publish versioned data, changes will be reflected if the provider publishes a new version.
3. Does the Data Share tool support object path-based partitioning on the Data Provider end when Oracle Autonomous Database is the consumer?
   No, the Data Share tool does not support object path-based partitioning on the Data Provider end. It supports consumption of only the values that are physically present within a file. The Oracle Autonomous Database consumers will not be able to view the partition values that are based on the file path.

Develop

Describes developer tasks and procedures on Autonomous Database.

- Create Applications with Oracle APEX in Autonomous Database
  You can create applications with Oracle APEX on Autonomous Database.
- Using JSON Documents with Autonomous Database
  Autonomous Database has full support for data represented as JSON documents. In Autonomous Databases, JSON documents can coexist with relational data.
- Developing RESTful Services in Autonomous Database
  You can develop and deploy RESTful Services with native Oracle REST Data Services (ORDS) support on Autonomous Databases. Simple Oracle Document Access (SODA) for REST lets you use a database as a simple JSON document store.
- Using Oracle Database API for MongoDB
  Oracle Database API for MongoDB makes it possible to connect to Oracle Autonomous Database using MongoDB language drivers and tools.
- Send Email and Notifications on Autonomous Database
  There are a number of options for sending email on Autonomous Database. You can also send notifications to a Slack channel, to MSTeams, or to an OCI Notifications topic.
- User Defined Notification Handler for Scheduler Jobs

- Oracle Extensions for IDEs
  Oracle extensions let developers connect to, browse, and manage Autonomous Databases directly from common IDEs.
- Use Oracle Java on Autonomous Database
- Test Workloads with Oracle Real Application Testing
  Oracle Real Application Testing is an extremely cost-effective and easy-to-use proactive performance management solution that enables you to fully assess the outcome of a system change in test or production.
- Manage and Store Files in a Cloud Code Repository with Autonomous Database
  Autonomous Database provides routines to manage and store files in Cloud Code (Git) Repositories. The supported Cloud Code Repositories are: GitHub, AWS CodeCommit, and Azure Repos.
- Invoke User Defined Functions
  You can invoke external functions such as OCI, AWS Lambda, Google Cloud Run functions and Azure remote functions in your Autonomous Database as SQL functions.
- Make External Calls Using a Customer-Managed Wallet
  When your Autonomous Database instance is on a private endpoint you can use a customer-managed wallet with procedures in `UTL_HTTP`, `DBMS_LDAP`, `UTL_SMTP`, or `UTL_TCP`. You can also use a customer-managed wallet when the scheduler sends SMTP email notifications for various events related to scheduler jobs.
- Use Cloud Tables to Store Logging and Diagnostic Information
  You can create Cloud Tables where table data resides on Oracle managed Cloud Storage and the table data does not consume database storage.

Create Applications with Oracle APEX in Autonomous Database

You can create applications with Oracle APEX on Autonomous Database.

- About Oracle APEX
  Oracle APEX is a low-code development platform that enables you to build scalable, secure enterprise applications with world-class features that can be deployed anywhere.
- Access Oracle APEX Administration Services
  Each Autonomous Database instance includes a dedicated instance of Oracle APEX; you can use this instance to create multiple workspaces. A workspace is a shared work area where you can build applications. You create workspaces in Oracle APEX Administration Services.
- Create Oracle APEX Workspaces in Autonomous Database
  An Autonomous Database instance does not have precreated Oracle APEX workspaces. Create a workspace if you have not already done so or use these instructions to create additional workspaces.
- Access Oracle APEX App Builder
  Use App Builder to create and manage Oracle APEX applications and application pages. The App Builder home page displays all installed applications in the current Oracle APEX workspace.
- Create Oracle APEX Developer Accounts
- Load Data from the Cloud into Oracle APEX
  Using Oracle APEX SQL Workshop, you can declaratively load data from Object Store into the database schema associated with your Oracle APEX workspace.

- Use JSON Data with Oracle APEX
  You can use Oracle APEX to create applications with JSON data. You must first create a view to extract the required attributes from the JSON data and maps them into columns of a relational view.
- Use Web Services with Oracle APEX
- Send Email from Oracle APEX
  You can use the `APEX_MAIL` package to send email from Oracle APEX applications deployed in Autonomous Database.
- Control Oracle APEX Upgrades
  Oracle automatically maintains Oracle APEX in your Autonomous Database including the deployment of patches and upgrades. When a new Oracle APEX release becomes available in your region, for example, APEX 24.2.0, you have a choice when to apply this upgrade within the 45-day upgrade window. You can optionally extend this window to 90 days.
- Access Oracle APEX, Oracle REST Data Services, and Developer Tools Using a Vanity URL
  By default you access Oracle APEX apps, REST endpoints, and developer tools on Autonomous Database using the `oraclecloudapps.com` domain name. You can optionally configure a vanity URL or custom domain name that is easy to remember to help promote your brand identity.
- Oracle APEX Limitations on Autonomous Database
  This section summarizes restrictions and limitations of Oracle APEX included in Oracle Autonomous Database and APEX Service.

About Oracle APEX

Oracle APEX is a low-code development platform that enables you to build scalable, secure enterprise applications with world-class features that can be deployed anywhere.

Oracle APEX provides you with an easy-to-use browser-based environment to load data, manage database objects, develop REST interfaces, and build applications which look and run great on both desktop and mobile devices. You can use Oracle APEX to develop a wide variety of solutions: import spreadsheets and develop a single source of truth in minutes, create compelling data visualizations against your existing data, deploy productivity applications to elegantly solve a business need, or build your next mission-critical data management application.

Oracle APEX embraces SQL. Anything you can express with SQL can be easily employed in an Oracle APEX application. Oracle APEX also embodies low code with powerful data management and data visualization components, as well as responsive development out of the box. Instead of writing code by hand, you are able to use intelligent wizards to guide you through the rapid creation of applications and components.

Oracle APEX on Autonomous Database provides a preconfigured, fully managed and secured environment to both build and deploy world-class data-centric applications. There are no limits on the number of developers or end users for your Oracle APEX applications; Autonomous Database can instantly scale compute and storage online as needed, based upon your workload. Additionally, Oracle APEX applications developed on-premise can be easily deployed to Oracle APEX on Autonomous Database, or vice-versa.

Configuration, patching, monitoring, and upgrading of all Oracle APEX components is fully managed by Oracle, leaving you free to focus on developing your solutions and solving your business problems. With Oracle APEX and low code, your organization can be more agile and develop solutions faster, for less cost, and with greater consistency. You can adapt to changing

requirements with ease. And you can empower professional developers and everyone else in your organization to be a part of the solution.

This chapter covers information on Oracle APEX specific to working on Autonomous Database.

For more information on APEX, see the following:

- Oracle APEX Release 24.2
- apex.oracle.com

Access Oracle APEX Administration Services

Each Autonomous Database instance includes a dedicated instance of Oracle APEX; you can use this instance to create multiple workspaces. A workspace is a shared work area where you can build applications. You create workspaces in Oracle APEX Administration Services.

> **Note:**
>
> If your Autonomous Database is configured to use a Private Endpoint, then you can only access Oracle APEX from clients in the same Virtual Cloud Network (VCN). See Configure Network Access with Private Endpoints for more information.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To access Oracle APEX Administration Services you can use the Oracle Cloud Infrastructure Console or Database Actions.

To access Administration Services from Database Actions:

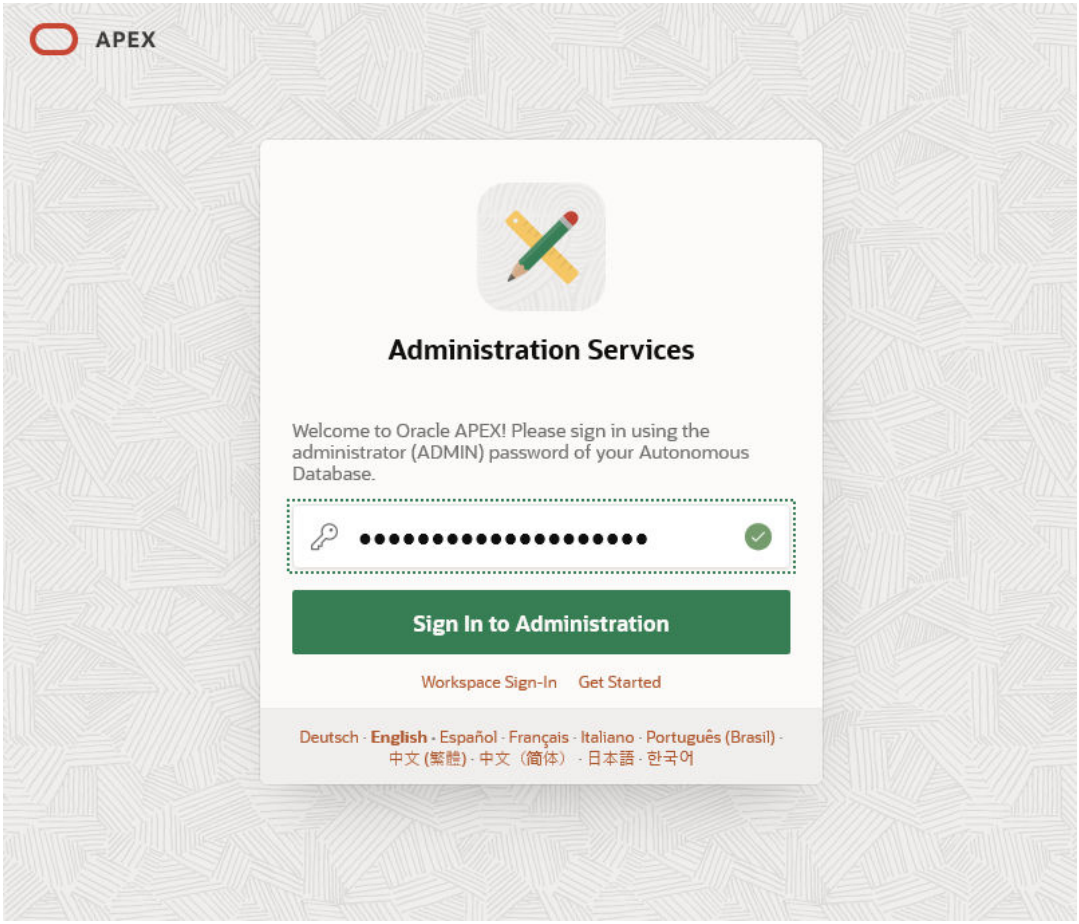
1. On the Autonomous Database details page select **Database Actions** and in the list click **View all database actions**.
2. From the Database Actions Launchpad, under **Development**, click **APEX**.

   The Oracle APEX Administration Services sign-in page appears.

   If you already created a workspace, the APEX workspace sign-in page appears instead. In this case, to open Administration Services, click Administration Services link.
3. In the **Password** field, enter the password for the ADMIN user.

> **Note:**
>
> By default, Administration Services and the Oracle APEX development environment on Autonomous Database use Database Accounts authentication. This authentication method uses the database account user name and password to authenticate users. You can switch to a different authentication method on the Administration Services, Manage Instance, Security page.

See Set the ADMIN Password in Autonomous Database to change the password.



4. Click **Sign In to Administration**.

   When you sign in for the first time, follow the prompts to create an APEX workspace. See Create Oracle APEX Workspaces in Autonomous Database for more information.

To access Administration Services from Oracle Cloud Infrastructure Console:

1. On the Autonomous Database details page click the **Tools** tab.
2. In the **Tools** column, select the Oracle APEX link.

You can also use Administration Services to manage your APEX instance. See Oracle APEX Administration Services in *Oracle APEX Administration Guide* for more information.

Create Oracle APEX Workspaces in Autonomous Database

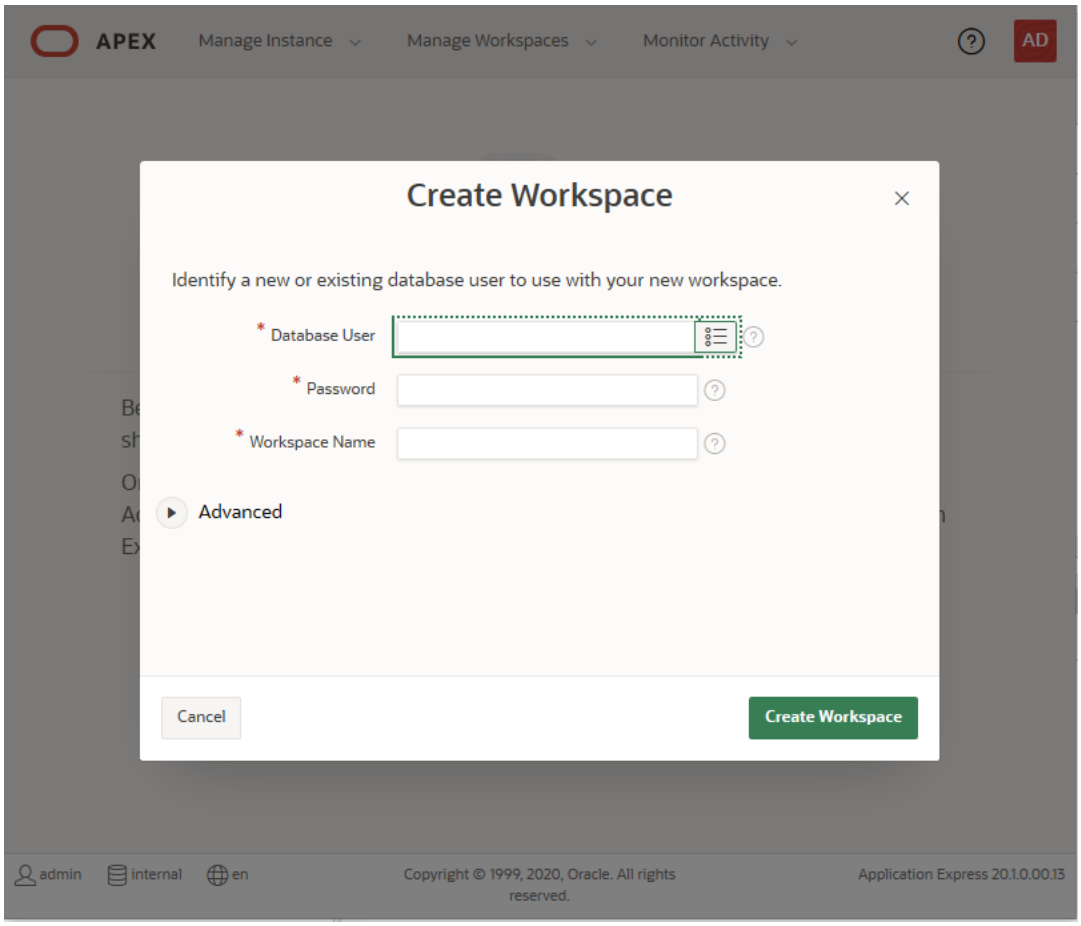
An Autonomous Database instance does not have precreated Oracle APEX workspaces. Create a workspace if you have not already done so or use these instructions to create additional workspaces.

To create an Oracle APEX workspace:

1. Sign in to Oracle APEX Administration Services.

   See Access Oracle APEX Administration Services for more information.

2. Click **Create Workspace**.



3. On the Create Workspace page, in the **Database User** field, enter a new database username or choose an existing user from the list.

   The ADMIN database user cannot be associated with a workspace.

4. In the Password field, provide a strong password if the database user is a new user. If the user is an existing database user you do not enter a password.

   See About User Passwords on Autonomous Database to learn more about the default password complexity rules.

5. (optional) In the **Workspace Name** field, change the name of the workspace that was automatically populated.

6. Click **Create Workspace**.

See Access Oracle APEX App Builder and Create Oracle APEX Developer Accounts to create additional developer accounts.

Access Oracle APEX App Builder

Use App Builder to create and manage Oracle APEX applications and application pages. The App Builder home page displays all installed applications in the current Oracle APEX workspace.

If your Autonomous Database is configured to use a Private Endpoint, you can only access Oracle APEX App Builder from clients in the same Virtual Cloud Network (VCN).

See Configure Network Access with Private Endpoints for more information.

To access Oracle APEX App Builder:

1. Sign in to Oracle APEX using the workspace name, username, and password you specify when you create the workspace.

> **Note:**
>
> Oracle APEX Administration Services and the Oracle APEX development environment on Autonomous Database use Database Accounts authentication. This authentication method uses the database account user name and password to authenticate users.

2. On the Workspace home page, click the App Builder icon.

See Create Oracle APEX Developer Accounts to create developer accounts.

Create Oracle APEX Developer Accounts

Oracle APEX developers need a developer account in each workspace where they wish to build applications.

The initial developer account is created when you create a workspace (this account also has Workspace Administrator privilege). These steps show you how to create additional developer accounts for members of your team or reset their passwords. When you create a developer account, a corresponding database user is automatically created.

To create developer accounts and provide direct access to Oracle APEX:

1. Sign in to Oracle APEX Administration Services.

   See Access Oracle APEX Administration Services for more information.

2. Click **Manage Workspaces**.
3. Under Workspace Actions, click **Manage Developers and Users**.
4. On the Manage Application Developers and Users page, click **Create User**.
5. On the Create/Edit User page, in the **Username** field, enter a username.
6. In the **Email Address** field, enter an email address.
7. (Optional) Use the on-screen and in-line help to fill in additional fields.
8. In the **User is an administrator** field, select **No**.

9. In the **User is a developer** field, select **Yes**.
10. In the **Password** field, enter a strong password.

    See Create Users on Autonomous Database - Connecting with a Client Tool to learn more about the default password complexity rules.
11. In the **Confirm Password** field, confirm the password.
12. At the top of the page, click **Create User**.

    Alternatively, click **Create and Create Another** if you want to create the user and create another user.

To share sign-in details with developers:

1. On the Autonomous Databases page, under the **Display Name** column, select the Autonomous Database instance that contains the user and the user's workspace.
2. On the Autonomous Database Details page select **Database Actions** and click **View all database actions**.
3. In Database Actions Launchpad, under **Development**, right-click **APEX** and choose **Copy URL**.
4. Provide the copied URL, along with the Workspace Name, the Username, and the Password for the developer account you created.

Using this URL developers can access the Oracle APEX environment without having to navigate to the Autonomous Database Oracle Cloud Infrastructure Console.

> **Note:**
>
> Changing the password of Workspace Administrators and Developers through **Manage Developers and Users** page or **Edit Profile** page only affects applications configured with "Oracle APEX Accounts" authentication scheme. To change the password used to access App Builder, use Database Actions or another client to change the password of the corresponding database user.

See Creating User Accounts in *Oracle APEX Administration Guide* for more information.

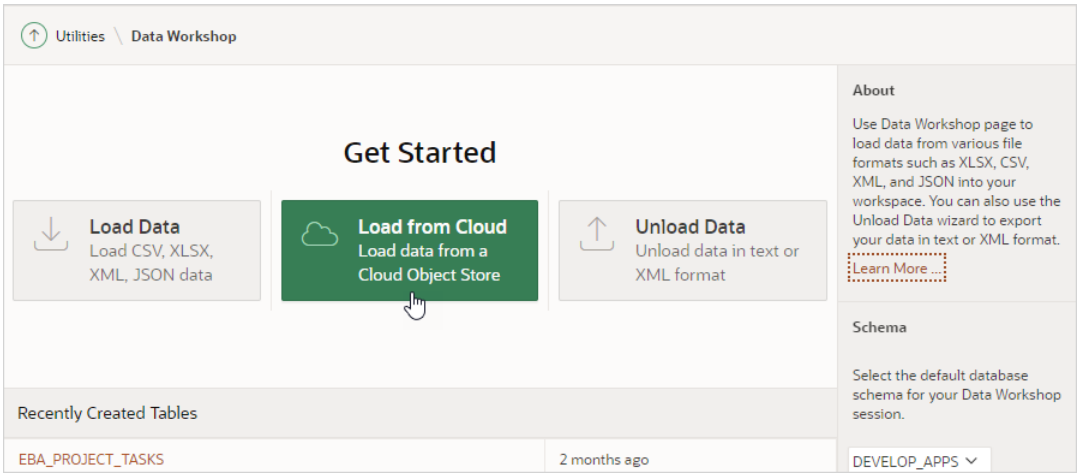
Load Data from the Cloud into Oracle APEX

Using Oracle APEX SQL Workshop, you can declaratively load data from Object Store into the database schema associated with your Oracle APEX workspace.

To load data from the cloud, do the following:

1. Open your Oracle APEX workspace.
2. Select **SQL Workshop**.
3. Select **Utilities**.
4. Select **Data Workshop**.

   This shows the Load from Cloud button.



5. Under Get Started, select **Load from Cloud**.

See Loading Data From the Cloud for more information.

Use JSON Data with Oracle APEX

You can use Oracle APEX to create applications with JSON data. You must first create a view to extract the required attributes from the JSON data and maps them into columns of a relational view.

- Create a View from JSON Data Guide
  Oracle APEX interprets data in relational format. Creating a view extracts required attributes from the JSON data and maps them into columns of a relational view.
- Create a View with JSON_TABLE Function
  You can create views of JSON data using the `json_table` SQL/JSON function.

Create a View from JSON Data Guide

Oracle APEX interprets data in relational format. Creating a view extracts required attributes from the JSON data and maps them into columns of a relational view.

For creating a view of JSON data that is stored in SODA collections, you can use SODA APIs and JSON Data Guide. The following PL/SQL code uses SODA APIs to create a Data Guide view on JSON Data stored in SODA Collections.

Run the following code in Oracle APEX SQL Workshop to create a view named `myview`:

```
-- Fetch the data guide and create a view
DECLARE
    coll  SODA_Collection_T;
    dg    CLOB;
    n     NUMBER;
BEGIN    -- Fetch the data guide from the collection or create one with
hierarchical format
    coll := dbms_soda.open_Collection('mycollection');
    dg := coll.get_Data_Guide;
    dbms_output.put_line(JSON_QUERY(dg, '$' pretty));
    -- User can modify the data guide as needed
    n := coll.create_View_From_DG('myview', dg);
    dbms_output.put_line('Status: ' || n);
```

```
    dbms_lob.freeTemporary(dg);
END;
/
```

Use the following command to check if the view has been created:

```
select count(1) from user_views where view_name = 'myview';
```

Use the following command to see the structure of the view:

```
describe myview;
```

See Create View using JSON Data Guide for more information on creating a view using JSON Data Guide.

Create a View with JSON_TABLE Function

You can create views of JSON data using the `json_table` SQL/JSON function.

The `json_table` SQL/JSON function projects specific JSON data to columns of various SQL data types. You can use the `json_table` function to map parts of a JSON document into the rows and columns of a new, virtual table, which you can also think of as an inline view.

See Create View on JSON Data for more information on creating views over JSON Data.

Use Web Services with Oracle APEX

You can interact with both SOAP and RESTful style web services from Oracle APEX in your Autonomous Database instance.
Web services enable applications to interact with one another over the web in a platform-neutral, language independent environment. In a typical web services scenario, a business application sends a request to a service at a given URL by using the HTTP protocol. The service receives the request, processes it, and returns a response. Web services are typically based on Simple Object Access Protocol (SOAP) or Representational State Transfer (REST) architectures.

Using REST Data Sources, APEX developers can declaratively access data services from a variety of REST endpoints, allowing both read and write operations. In addition to supporting smart caching rules for remote REST data, Oracle APEX also offers the unique ability to directly manipulate the results of REST data sources using industry standard SQL.

The `APEX_WEB_SERVICE` package enables you to integrate other systems with APEX by allowing you to interact with web services anywhere you can use PL/SQL in your application. The package contains procedures and functions to call both SOAP and RESTful style web services, and to simplify implementation of OAuth 2.0 flows.

- Submit an HTTP Request to a Private Host with Oracle APEX
  There are several requirements to submit an HTTP request to a private host from Oracle APEX in your Autonomous Database instance.
- Make External Calls Using a Customer-Managed Wallet with Oracle APEX
  Oracle APEX on Autonomous Database is preconfigured with an Oracle Wallet that will work with most HTTPS endpoints. If you want to consume REST APIs from servers that use less common or self-signed SSL certificates, you can configure a customer-managed wallet.

- Make External Calls Using a Web Proxy with Oracle APEX
  Your Oracle APEX instance does not require an outbound web proxy to make external REST APIs calls. Autonomous Databases configured with a private endpoint can use a web proxy deployed in your VCN if desired.
- Notes for Using Web Services with Oracle APEX
  Provides notes for working with web services in Oracle APEX from an Autonomous Database instance.

Submit an HTTP Request to a Private Host with Oracle APEX

There are several requirements to submit an HTTP request to a private host from Oracle APEX in your Autonomous Database instance.

To submit a request to a target host on a private endpoint or behind on-premises firewalls, the target host must be accessible from the source database's Oracle Cloud Infrastructure VCN:

1. Ensure the endpoint meets the prerequisites to submit an HTTP request, as described here: Submit an HTTP Request to a Private Host with UTL_HTTP.
2. Add the following access control list for the desired host as ADMIN (note the `private_target` parameter):

```
BEGIN
  DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
    host => 'www.example.com',
    ace => XS$ACE_TYPE(
      privilege_list => XS$NAME_LIST('http'),
      principal_name => APEX_APPLICATION.g_flow_schema_owner,
      principal_type => XS_ACL.ptype_db),
    private_target => true);
END;
/
```

> **Note:**
>
> If you set `ROUTE_OUTBOUND_CONNECTIONS` database property to `PRIVATE_ENDPOINT`, you do not need to define access control lists for individual hosts in order to access them from APEX. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

Make External Calls Using a Customer-Managed Wallet with Oracle APEX

Oracle APEX on Autonomous Database is preconfigured with an Oracle Wallet that will work with most HTTPS endpoints. If you want to consume REST APIs from servers that use less common or self-signed SSL certificates, you can configure a customer-managed wallet.

Oracle APEX on Autonomous Database is preconfigured with an Oracle Wallet that contains more than 90 of the most common trusted root and intermediate SSL certificates. The `APEX_WEB_SERVICE` package automatically takes advantage of this Oracle Wallet without additional configuration from application developers.

If you want to consume REST APIs from servers that use less common or self-signed SSL certificates, you can configure a customer-managed wallet. First, follow configuration and deployment instructions in Prerequisites to Use a Customer-Managed Wallet with External

Calls. After the wallet is deployed to your Autonomous Database, grant `READ` on the directory object containing the wallet to the APEX platform schema as `ADMIN`:

```
BEGIN
  execute immediate 'grant READ on directory WALLET_DIR to ' ||
    APEX_APPLICATION.g_flow_schema_owner;
END;
```

Next, when you reference the wallet in `APEX_WEB_SERVICE` API calls, make sure to use the `DIR:` prefix instead of the usual `FILE:` prefix. For example:

```
l_resp := APEX_WEB_SERVICE.MAKE_REST_REQUEST(
  p_url         => 'https://www.example.com/',
  p_http_method => 'GET',
  p_wallet_path => 'DIR:WALLET_DIR'
);
```

You can also configure this wallet as the default for the entire Oracle APEX instance in **APEX Administration Services** by selecting Manage Instance → Instance Settings → Wallet → Wallet Path, or by setting the `WALLET_PATH` instance parameter.

To switch back to the preconfigured wallet, set `WALLET_PATH` to an empty value.

See Access Oracle APEX Administration Services for more information.

Make External Calls Using a Web Proxy with Oracle APEX

Your Oracle APEX instance does not require an outbound web proxy to make external REST APIs calls. Autonomous Databases configured with a private endpoint can use a web proxy deployed in your VCN if desired.

To use a web proxy define the following access control list as `ADMIN`:

```
BEGIN
  DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
    host           => 'www-proxy.example.com',
    lower_port     => 80,
    ace => XS$ACE_TYPE(
      privilege_list => XS$NAME_LIST('HTTP_PROXY'),
      principal_name => APEX_APPLICATION.g_flow_schema_owner,
      principal_type => XS_ACL.ptype_db),
    private_target => true);
END;
/
```

Then, you can use the web proxy in `APEX_WEB_SERVICE` API calls:

```
l_resp := APEX_WEB_SERVICE.MAKE_REST_REQUEST(
  p_url            => 'https://www.example.com/',
  p_http_method    => 'GET',
  p_proxy_override => 'www-proxy.example.com:80'
);
```

You can also configure a web proxy at the application level by setting Proxy Server property on Shared Components → Application Definition → Properties page, or at the instance level in **APEX Administration Services** from Manage Instance → Security → Instance Proxy.

If you attempt to use a web proxy when your Autonomous Database is configured with a public endpoint, an error is returned:

```
ORA-01031: insufficient privileges
```

See Submit an HTTP Request to a Private Host with UTL_HTTP for more information on proxies when using `UTL_HTTP` API.

See Access Oracle APEX Administration Services for more information.

Notes for Using Web Services with Oracle APEX

Provides notes for working with web services in Oracle APEX from an Autonomous Database instance.

Note the following when working with web services in Oracle APEX:

- All web services must be secured. Only HTTPS services are supported on the default port (443). Connections through IP addresses are not allowed.
- Each Autonomous Database instance is preconfigured with a network access control list (ACL) to permit outbound web service calls from Oracle APEX to public endpoints.
- `APEX_WEB_SERVICE` requests do not honor the custom wallet you set with `UTL_HTTP.SET_WALLET` API calls.
- There is a default limit of 50,000 outbound web service requests per APEX workspace in a rolling 24-hour period. If the limit of outbound web service calls is reached, the following SQL exception is raised on the subsequent request and the request is blocked:

  ```
  ORA-20001: You have exceeded the maximum number of web service requests per
  workspace. Please contact your administrator.
  ```

  You can raise or remove the default limit of outbound web service requests by setting a value for the `MAX_WEBSERVICE_REQUESTS` instance parameter or by updating the Maximum Web Service Requests attribute in APEX Administration Services. For example, to change the limit to 250,000, connect to your database as ADMIN using a SQL client and execute the following:

  ```
  BEGIN
    APEX_INSTANCE_ADMIN.SET_PARAMETER('MAX_WEBSERVICE_REQUESTS', '250000');
    COMMIT;
  END;
  /
  ```

To learn more, see:

- APEX_WEB_SERVICE in *Oracle APEX API Reference*
- Managing REST Data Sources in *Oracle APEX App Builder User's Guide*

Send Email from Oracle APEX

You can use the `APEX_MAIL` package to send email from Oracle APEX applications deployed in Autonomous Database.

Before you use `APEX_MAIL` you must configure an SMTP server that will deliver email from Oracle APEX. Oracle APEX on Autonomous Database supports Oracle Cloud Infrastructure Email Delivery service as well as third party email providers.

- Configure OCI Email Delivery
  Your Oracle APEX instance can utilize OCI Email Delivery service to send email from APEX applications.
- Configure Third Party Email Provider
  Your Oracle APEX instance can utilize third party email providers to send email from APEX applications.
- Validate Email Configuration
  Shows the steps to confirm your email provider is ready to use with `APEX_MAIL`.
- Notes for Sending Email from Oracle APEX
  Provides notes for working with email providers in Oracle APEX from an Autonomous Database instance.

Configure OCI Email Delivery

Your Oracle APEX instance can utilize OCI Email Delivery service to send email from APEX applications.

To configure `APEX_MAIL` functionality to use OCI Email Delivery service:

1. Identify the SMTP connection endpoint for Email Delivery. You configure the endpoint as the SMTP Host in your APEX instance in Step 4. You may need to subscribe to additional Oracle Cloud Infrastructure regions if Email Delivery is not available in your current region. See Configure SMTP Connection for more information.
2. Generate SMTP credentials for Email Delivery. Your APEX instance uses credentials to authenticate with Email Delivery servers when you send email. See Creating SMTP Credentials for more information.
3. Create an approved sender for Email Delivery. You need to complete this step for all email addresses you use as the "From" with `APEX_MAIL.SEND` calls, as the Application Email "From" address in your apps, or in the `SMTP_FROM` instance parameter. See Managing Approved Senders for more information.
4. Connect to your Autonomous Database as ADMIN user using a SQL client and configure the following SMTP parameters using `APEX_INSTANCE_ADMIN.SET_PARAMETER`:
   - `SMTP_HOST_ADDRESS`: Specifies the SMTP connection endpoint from Step 1
   - `SMTP_USERNAME`: Specifies the SMTP credential user name from Step 2
   - `SMTP_PASSWORD`: Specifies the SMTP credential password from Step 2
   - `SMTP_HOST_PORT`: Keep the default value for this parameter (587)
   - `SMTP_TLS_MODE`: Keep the default value for this parameter (`STARTTLS`)

For example:

```
BEGIN
  APEX_INSTANCE_ADMIN.SET_PARAMETER('SMTP_HOST_ADDRESS', 'smtp.email.us-
phoenix-1.oci.oraclecloud.com');
  APEX_INSTANCE_ADMIN.SET_PARAMETER('SMTP_USERNAME',
'ocid1.user.oc1.username');
  APEX_INSTANCE_ADMIN.SET_PARAMETER('SMTP_PASSWORD', 'password');
  COMMIT;
END;
/
```

See APEX_INSTANCE_ADMIN for more information.

See Overview of Email Delivery for more information.

Configure Third Party Email Provider

Your Oracle APEX instance can utilize third party email providers to send email from APEX applications.

Before configuring a third party email provider in APEX, ensure the following prerequisites are met:

- Your Autonomous Database must be configured to use a Private Endpoint.
- The email provider's SMTP endpoint (port 25 or 587) must be accessible from the Virtual Cloud Network (VCN) and subnet where your Autonomous Database is provisioned.

See Configure Network Access with Private Endpoints for more information.

After the above prerequisites are met:

1. Add the following access control list for the SMTP endpoint as ADMIN (note the `private_target` parameter):

```
BEGIN
  DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
    host => 'smtp.example.com',
    ace => XS$ACE_TYPE(
      privilege_list => XS$NAME_LIST('SMTP'),
      principal_name => APEX_APPLICATION.g_flow_schema_owner,
      principal_type => XS_ACL.ptype_db),
    private_target => true);
END;
/
```

> **Note:**
>
> If you set `ROUTE_OUTBOUND_CONNECTIONS` database property to `PRIVATE_ENDPOINT`, you do not need to define access control lists for individual SMTP endpoints in order to access them from APEX. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

2. Connect to your Autonomous Database as ADMIN user using a SQL client and configure the following SMTP parameters using `APEX_INSTANCE_ADMIN.SET_PARAMETER`. If your third

party email provider does not require authentication, skip the `SMTP_USERNAME` and `SMTP_PASSWORD` parameters.

- `SMTP_HOST_ADDRESS`: Specifies the SMTP connection endpoint from Step 1
- The valid values for `SMTP_HOST_PORT` parameter are 25 or 587
- Set the value for the `SMTP_TLS_MODE` parameter to `Y`

For example:

```
BEGIN
  APEX_INSTANCE_ADMIN.SET_PARAMETER('SMTP_HOST_ADDRESS',
'smtp.example.com');
  APEX_INSTANCE_ADMIN.SET_PARAMETER('SMTP_HOST_PORT', '587');
  APEX_INSTANCE_ADMIN.SET_PARAMETER('SMTP_TLS_MODE', 'Y');
  APEX_INSTANCE_ADMIN.SET_PARAMETER('SMTP_USERNAME', 'username');
  APEX_INSTANCE_ADMIN.SET_PARAMETER('SMTP_PASSWORD', 'password');
  COMMIT;
END;
/
```

See APEX_INSTANCE_ADMIN for more information.

Validate Email Configuration

Shows the steps to confirm your email provider is ready to use with `APEX_MAIL`.

1. Validate the email configuration settings using a SQL client as ADMIN:

   ```
   BEGIN
     APEX_INSTANCE_ADMIN.VALIDATE_EMAIL_CONFIG;
   END;
   /
   ```

   If any errors are reported (for example, "`ORA-29279: SMTP permanent error: 535 Authentication credentials invalid`"), adjust the SMTP parameters or Virtual Cloud Network (VCN) configuration and repeat the validation step.

   See APEX_INSTANCE_ADMIN for more information.

2. Send a test email using APEX SQL Workshop, If you are using OCI Email Delivery, ensure the `p_from` parameter matches one of the approved senders. For example:

   ```
   BEGIN
     APEX_MAIL.SEND(
       p_from => 'alice@example.com',
       p_to   => 'bob@example.com',
       p_subj => 'Email from Oracle Autonomous Database',
       p_body => 'Sent using APEX_MAIL');
   END;
   /
   ```

   See APEX_MAIL for more information.

3. To monitor email delivery in your APEX instance:

a. Sign in to APEX Administration Services.

b. Open the Manage Instance page.

c. Click the Mail Queue link in the Manage Meta Data section.

Alternatively, query `APEX_MAIL_QUEUE` and `APEX_MAIL_LOG` views using a SQL client.

Notes for Sending Email from Oracle APEX

Provides notes for working with email providers in Oracle APEX from an Autonomous Database instance.

Note the following when working with email providers in Oracle APEX:

- There is a default limit of 1,000 emails per workspace in a 24-hour period. You can update or remove this limit in Oracle APEX Administration Services or by setting the `WORKSPACE_EMAIL_MAXIMUM` instance parameter. If you are using OCI Email Delivery service, the service may impose additional limitations.
- When using OCI Email Delivery service, an approved sender must be set up for all possible “From” addresses or mail will be rejected. See Managing Approved Senders for more information.

Control Oracle APEX Upgrades

Oracle automatically maintains Oracle APEX in your Autonomous Database including the deployment of patches and upgrades. When a new Oracle APEX release becomes available in your region, for example, APEX 24.2.0, you have a choice when to apply this upgrade within the 45-day upgrade window. You can optionally extend this window to 90 days.

> **Note:**
>
> The latest Oracle APEX Patch Set Bundles are automatically applied on Autonomous Database and cannot be deferred.

- Defer Oracle APEX Upgrades
  Describes the steps to defer Oracle APEX upgrades.
- Apply Oracle APEX Upgrades
  When you have deferred upgrades and an upgrade is available, you have the option to apply the upgrade at any time before the specified upgrade date.

Defer Oracle APEX Upgrades

Describes the steps to defer Oracle APEX upgrades.

Deferring Oracle APEX upgrades allows you to do the following:

- Control when Oracle APEX is upgraded on your Autonomous Database. While there should be minimal disruption on a live system, you can still choose to upgrade at a time or on a day with less activity, for example on a weekend day.
- Validate your applications against the new version of Oracle APEX before you upgrade your production system.

  To validate and test before you upgrade your production system, do the following:

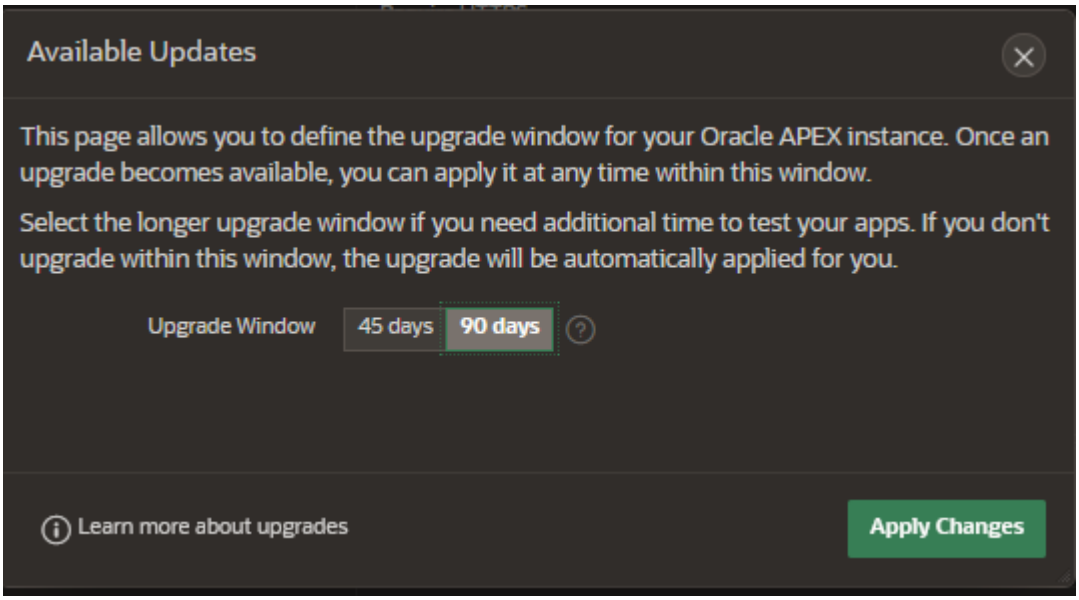
- Clone your production or test environment. See Clone an Autonomous Database Instance for more information.
- Upgrade Oracle APEX on the clone as specified in Apply Oracle APEX Upgrades.
- Perform your testing and validation on the clone in a non-production environment.

To defer Oracle APEX upgrades:

1. Open Oracle APEX Administration Services.

   See Access Oracle APEX Administration Services for more information.
2. Click **APEX** to open the Administration Services page.
3. Open the Available Updates dialog by clicking the  next to **Available Updates**.

   This shows the Available Updates dialog.
4. In the Available Updates dialog select an **Upgrade Window**.



5. Click **Apply Changes**.

Notes for deferred Oracle APEX upgrades:

- If an Upgrade Window is selected and updates are available, then on the date indicated in the Available Updates dialog, Oracle APEX updates are automatically applied to your Autonomous Database instance. You cannot postpone the updates beyond the listed date.
- If an upgrade window is set and updates are available, you cannot schedule the Oracle APEX upgrades for a specific date and time. When you are ready to upgrade, access Oracle APEX Administration Services and click **Upgrade Now**. See Apply Oracle APEX Upgrades for details.
- If your Autonomous Database is stopped at the time when a new Oracle APEX release becomes available in your region, the update is automatically applied the next time you start your database, regardless of the Upgrade Window setting. For example, if the Autonomous Database is stopped when a new Oracle APEX release becomes available, based on the Auto Start/Stop Schedule, the update is automatically applied the next time the database starts.

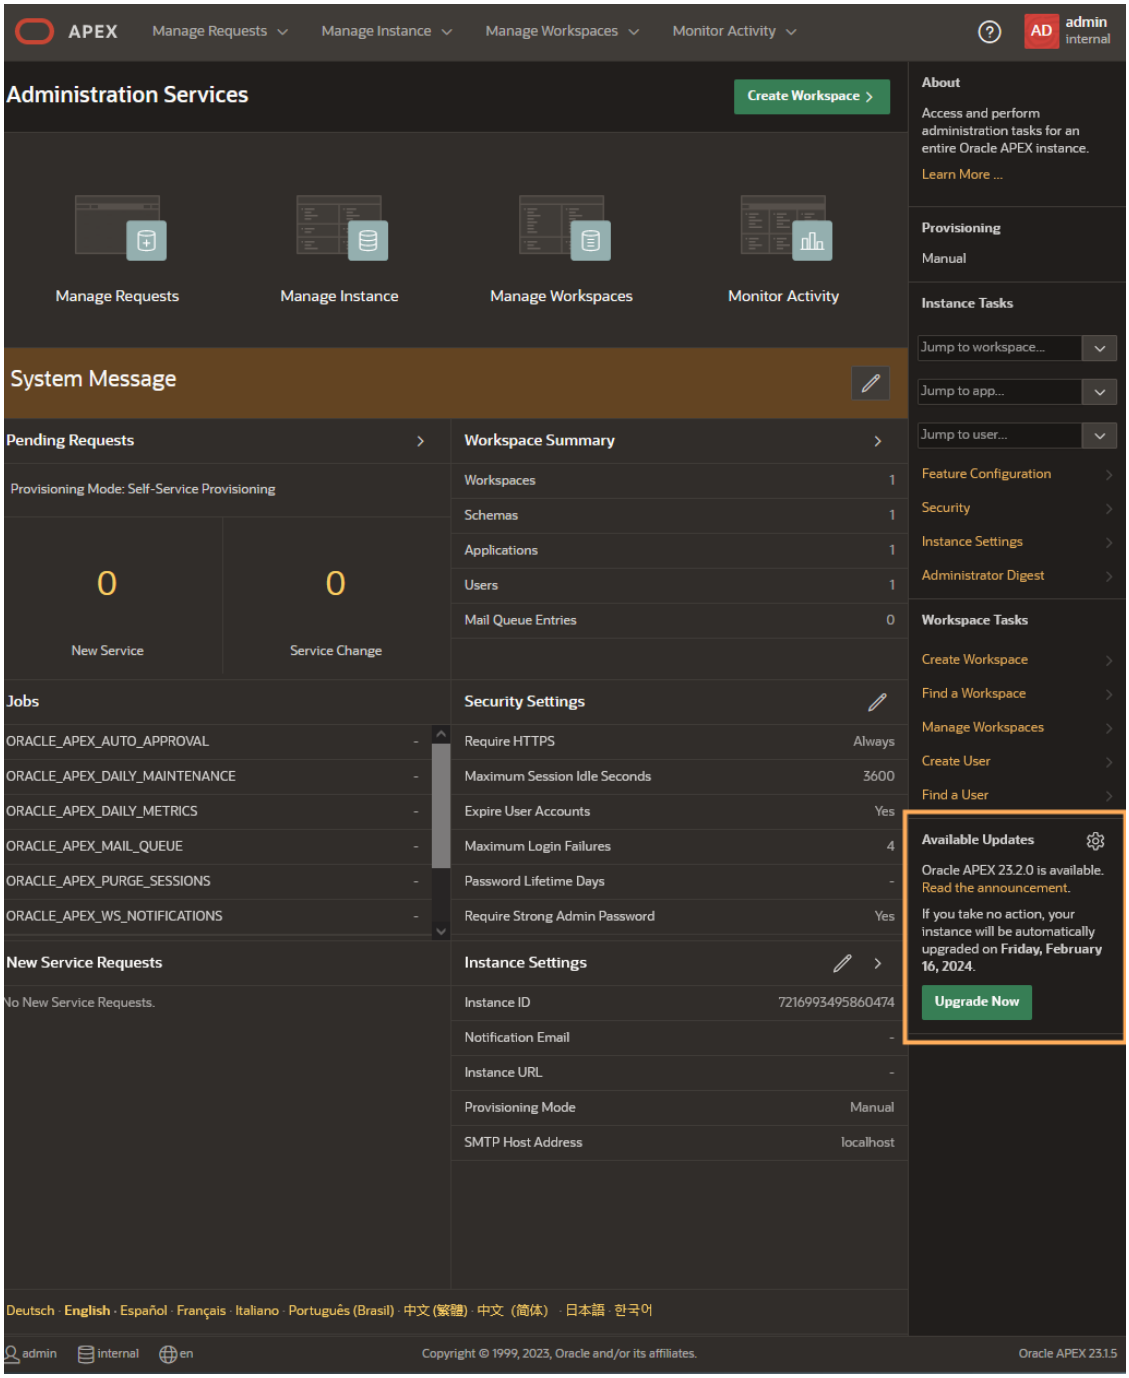
- When you change the Upgrade Window setting, the change applies to future upgrades and does not impact an Oracle APEX upgrade that is already available to your Autonomous Database.
- If your database did not have Oracle APEX activity in the past 6 months, the next Oracle APEX upgrade will be automatically applied regardless of the Upgrade Window setting.

Apply Oracle APEX Upgrades

When you have deferred upgrades and an upgrade is available, you have the option to apply the upgrade at any time before the specified upgrade date.

When Oracle APEX upgrades are not available the Available Updates area shows **System is up-to-date**.

When the Defer Upgrade setting specifies an Upgrade Window, if an Oracle APEX upgrade is available and you take no action, the upgrade is applied automatically on the date specified in the Available Updates area.



If you previously deferred upgrades by specifying an upgrade window, as specified in Defer Oracle APEX Upgrades, then you can upgrade at any time, as follows:

1. Open Oracle APEX Administration Services.
2. Click **APEX** to open the Administration Services page.
3. Open the Available Updates dialog by clicking the  next to **Available Updates**.
4. On the Available Updates dialog click **Upgrade Now**.

This starts the Oracle APEX upgrade process in the background. You can continue using Administration Services and App Builder as well as running apps during most of the upgrade process.

Access Oracle APEX, Oracle REST Data Services, and Developer Tools Using a Vanity URL

By default you access Oracle APEX apps, REST endpoints, and developer tools on Autonomous Database using the `oraclecloudapps.com` domain name. You can optionally configure a vanity URL or custom domain name that is easy to remember to help promote your brand identity.

After you acquire a desired domain name and matching SSL certificate from a vendor of your choice, deploy an Oracle Cloud Infrastructure Load Balancer in your Virtual Cloud Network (VCN) using your Autonomous Database as the backend. Your Autonomous Database instance must be configured with a private endpoint in the same VCN. See Configuring Network Access with Private Endpoints for more information.

To learn more, see the following:

- Introducing Vanity URLs for APEX and ORDS on Oracle Autonomous Database
- Automate Vanity URL Configuration Using Terraform

Oracle APEX Limitations on Autonomous Database

This section summarizes restrictions and limitations of Oracle APEX included in Oracle Autonomous Database and APEX Service.

> **Tip:**
>
> This section does not apply to co-managed databases (for example, OCI Base Database Service) and Oracle Autonomous Database on Dedicated Exadata Infrastructure.

In a fully managed environment (such as Oracle Autonomous Database Serverless), certain limitations are required to protect the security and performance of your Oracle APEX environment.

- Administration Services - The following Oracle APEX Administration Services configuration options are disabled, or have been predefined by Oracle and cannot be altered.
  - Manage Instance, Feature Configuration:
    - Monitoring - Web Service Activity Logging, Enable Application Tracing
  - Manage Instance, Security, Security Settings:
    - Security - Unhandled Errors
    - Authentication Control:
      - HTTP Protocol - Require Outbound HTTPS
      - General - Single Sign-On Logout URL
  - Manage Instance, Manage Logs and Files:
    - SQL Workshop Log
    - Page View Activity Log
    - Developer Activity Log

        - External Click Counting Log
        - Login Access Log
        - Web Service Activity Log
        - REST Synchronization Log
        - Automation Log
    - Manage Instance, Instance Settings - Storage - All tablespace settings
    - Manage Instance, Workspace Purge Settings
    - Manage Workspaces, Existing Workspaces, Edit Workspace Information:
        - Login Control
        - Session Timeout
        - Workspace Isolation
- Only the following `APEX_INSTANCE_ADMIN` procedures and functions are supported:
    - `ADD_AUTHORIZED_URL`
    - `ADD_SCHEMA`
    - `ADD_WORKSPACE`
    - `CREATE_CLOUD_CREDENTIAL`
    - `CREATE_OR_UPDATE_ADMIN_USER`
    - `DISABLE_WORKSPACE`
    - `ENABLE_WORKSPACE`
    - `FREE_WORKSPACE_APP_IDS`
    - `GET_AUTHORIZED_URLS`
    - `GET_PARAMETER`
    - `GET_SCHEMAS`
    - `GRANT_EXTENSION_WORKSPACE`
    - `REMOVE_APPLICATION`
    - `REMOVE_AUTHORIZED_URL`
    - `REMOVE_SAVED_REPORTS`
    - `REMOVE_SAVED_REPORT`
    - `REMOVE_SCHEMA`
    - `REMOVE_SUBSCRIPTION`
    - `REMOVE_WORKSPACE`
    - `RESERVE_WORKSPACE_APP_IDS`
    - `REVOKE_EXTENSION_WORKSPACE`
    - `SET_LOG_SWITCH_INTERVAL`
    - `SET_PARAMETER`
    - `UNLOCK_USER`

  - `VALIDATE_EMAIL_CONFIG`

  See APEX_INSTANCE_ADMIN in *Oracle APEX API Reference*.

- The following application authentication schemes are supported with limitations:
  - HTTP Header Variable: Only after configuring a vanity URL.

    See Access Oracle APEX, Oracle REST Data Services, and Developer Tools Using a Vanity URL for more information.
  - LDAP Directory, including `APEX_LDAP` API: With the same restrictions that apply to the `DBMS_LDAP` package.

    See PL/SQL Package Notes for Autonomous Database for more information.
  - SAML Sign-In: Only when using a customer managed ORDS.

    See About Customer Managed Oracle REST Data Services on Autonomous Database for more information.
- Disabled options in SQL Workshop:

  The ability to create and manage database links in Object Browser is disabled. To create a database link, use the `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` PL/SQL API.

  See Use Database Links with Autonomous Database for more information.
- Runtime environment:

  Oracle APEX is only available as a Full Development environment. Converting into a Runtime environment, which minimizes the installed software footprint and removes UI components such as App Builder and Administration Services, is not supported.

  See About the Differences Between Runtime and Full Development Environments in *Oracle APEX App Builder User's Guide* to learn more.

Using JSON Documents with Autonomous Database

Autonomous Database has full support for data represented as JSON documents. In Autonomous Databases, JSON documents can coexist with relational data.

You can access and manage JSON documents as follows:

- Using Simple Oracle Document Access (SODA) APIs for NoSQL-style access to document collections.
- Using SQL and PL/SQL to access JSON documents in tables.
- Work with Simple Oracle Document Access (SODA) in Autonomous Database
  Simple Oracle Document Access (SODA) is a set of NoSQL-style APIs that let you use collections of JSON documents in Autonomous Database, retrieve them, and query them, without needing to know Structured Query Language (SQL) or how the documents are stored in the database.
- Work with JSON Documents Using SQL and PL/SQL APIs on Autonomous Database
  Autonomous Database supports JavaScript Object Notation (JSON) data natively in the database.
- Load JSON Documents with Autonomous Database
  You can use the PL/SQL procedure `DBMS_CLOUD.COPY_COLLECTION` to load JSON documents into SODA collections or use `DBMS_CLOUD.COPY_DATA` to load JSON documents into an existing table in Autonomous Database.

- Import SODA Collection Data Using Oracle Data Pump Version 19.6 or Later
  Shows the steps to import SODA collections into Autonomous Database with Oracle Data Pump.

Work with Simple Oracle Document Access (SODA) in Autonomous Database

Simple Oracle Document Access (SODA) is a set of NoSQL-style APIs that let you use collections of JSON documents in Autonomous Database, retrieve them, and query them, without needing to know Structured Query Language (SQL) or how the documents are stored in the database.

Autonomous Database supports storing and querying JSON documents natively. SODA document collections are backed by ordinary database tables and views; you can take advantage of database features for use with the content of SODA documents.

SODA drivers are available for several languages and frameworks including: Java, Node.js, Python, C (using Oracle Call Interface), and PL/SQL, and SODA for REST. SODA for REST maps SODA operations to Uniform Resource Locator (URL) patterns, so it can be used with most programming languages.

To get started with SODA, see the following:

- Overview of SODA

Depending on the SODA API you want to work with, see the following:

> **Note:**
>
> If you are using Always Free Autonomous Database with Oracle Database 23ai, then to avoid compatibility problems of SODA drivers, Oracle recommends the following:
>
> - Use the driver versions that are needed for working with `JSON` type as specified in SODA Drivers. See SODA Drivers for more information.
> - For projects that were started using a database release prior to Oracle Database 21c, explicitly specify the metadata for the default collection as specified in the example in SODA Drivers. For projects started using release Oracle Database 21c or later, just use the default metadata. See SODA Drivers for more information.

SODA API	Download and Installation	More Information
SODA for Java	Download SODA for Java SODA for Java Prerequisites **Versions:** SODA for Java, using the latest version is recommended. The minimum supported version is: 1.1.4. Use SODA for Java in conjunction with `ojdbc8.jar` (available at Oracle Database JDBC driver and Companion Jars Downloads, or on Maven Central). Autonomous Database does not support Metadata builder. To customize collection metadata pass collection metadata strings directly to the createCollection method. See SODA Collection Metadata on Autonomous Database for more information.	SODA for Java
SODA for REST	Access RESTful Services and SODA for REST	Use SODA for REST with Autonomous Database
SODA for C	Oracle Instant Client Downloads **Versions:** For SODA for C, Oracle Client libraries must be 19.6 and above. You can obtain Oracle Instant Client from Oracle Instant Client Downloads.	SODA for C
SODA for PL/SQL	No need to download. This is included with Autonomous Database.	SODA for PL/SQL
SODA for Node.js	Multiple downloads described in install instructions: Quick Start node-oracledb Installation **Versions:** SODA support was introduced in version 3.0. Using the latest version is recommended, the minimum recommended version is 4.0. Oracle Client libraries must be 19.6 and above. You can obtain Oracle Instant Client from Oracle Instant Client Downloads.	Node-oracledb SODA Requirements
SODA for Python	Multiple downloads described in install instructions: cx_Oracle 7 Installation **Versions:** SODA support was introduced in version 7.0. Using the latest version is recommended, the minimum recommended version is 7.1. Oracle Client libraries must be 19.6 and above. You can obtain Oracle Instant Client from Oracle Instant Client Downloads.	Introduction to cx_Oracle

- SODA Notes
  When you use SODA with Autonomous Database the following restrictions apply:

SODA Notes

When you use SODA with Autonomous Database the following restrictions apply:

- Automatic indexing is not supported for SQL and PL/SQL code that uses the SQL/JSON function `json_exists`. See SQL/JSON Condition JSON_EXISTS for more information.
- Automatic indexing is not supported for SODA query-by-example (QBE).

Work with JSON Documents Using SQL and PL/SQL APIs on Autonomous Database

Autonomous Database supports JavaScript Object Notation (JSON) data natively in the database.

When you use Autonomous Database to store JSON data you can take advantage of all the features available in your database. You can combine your JSON data with non-JSON data; then, using Autonomous Database features such as Oracle Machine Learning Notebooks, you can analyze your data and create reports. You can access JSON data stored in the database the same way you access other database data, including using Oracle Call Interface (OCI), Microsoft .NET Framework, and Java Database Connectivity (JDBC).

See JSON in Oracle Database to get started.

Load JSON Documents with Autonomous Database

You can use the PL/SQL procedure `DBMS_CLOUD.COPY_COLLECTION` to load JSON documents into SODA collections or use `DBMS_CLOUD.COPY_DATA` to load JSON documents into an existing table in Autonomous Database.

Topics

- About Loading JSON Documents
- Load a JSON File of Line-Delimited Documents into a Collection
- Load an Array of JSON Documents into a Collection
- Monitor and Troubleshoot COPY_COLLECTION Loads
- Create Credentials and Copy JSON Data into an Existing Table

Import SODA Collection Data Using Oracle Data Pump Version 19.6 or Later

Shows the steps to import SODA collections into Autonomous Database with Oracle Data Pump.

You can export and import SODA collections using Oracle Data Pump Utilities starting with version 19.6. Oracle recommends using the latest Oracle Data Pump version for importing data from Data Pump files into your database.
Download the latest version of Oracle Instant Client, which includes Oracle Data Pump, for your platform from Oracle Instant Client Downloads. See the installation instructions on the platform install download page for the installation steps required after you download Oracle Instant Client.

In Oracle Data Pump, if your source files reside on Oracle Cloud Infrastructure Object Storage you can use Oracle Cloud Infrastructure native URIs, Swift URIs, or pre-authenticated URIs. See DBMS_CLOUD Package File URI Formats for details on these file URI formats.

If you are using an Oracle Cloud Infrastructure pre-authenticated URI, you still need to supply a `credential` parameter. However, credentials for a pre-authenticated URL are ignored (and the supplied credentials do not need to be valid). See DBMS_CLOUD Package File URI Formats for information on Oracle Cloud Infrastructure pre-authenticated URIs.

This example shows how to create the SODA collection metadata and import a SODA collection with Data Pump.

1. On the source database, export the SODA collection using the Oracle Data Pump `expdp` command.

   See Export Your Existing Oracle Database to Import into Autonomous Database for more information.

2. Upload the dump file set from Step 1 to Cloud Object Storage.

3. Create a SODA collection with the required SODA collection metadata on your Autonomous Database.

   For example, if you export a collection named *MyCollectionName* from the source database with the following metadata:

   - The content column is a `BLOB` type.
   - The version column uses the `SHA256` method.

   Then on the Autonomous Database where you import the collection, create a new collection:

   - By default on Autonomous Database for a new collection the content column is set to `BLOB` with the `jsonFormat` specified as `OSON`.
   - By default on Autonomous Database for a new collection the `versionColumn.method` is set to `UUID`.

   See SODA Default Collection Metadata on Autonomous Database for details.

   For example:

```
DECLARE
   collection_create SODA_COLLECTION_T;
BEGIN
   collection_create := DBMS_SODA.CREATE_COLLECTION('MyCollectionName');
END;
/
COMMIT;
```

   You can use the PL/SQL function `DBMS_SODA.LIST_COLLECTION_NAMES` to discover existing collections. See LIST_COLLECTION_NAMES Function for more information.

   You can view the metadata for the SODA collections by querying the view `USER_SODA_COLLECTIONS`. See USER_SODA_COLLECTIONS for more information.

4. Store your Cloud Object Storage credential using `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example, to create Oracle Cloud Infrastructure Auth Token credentials:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'adb_user@example.com',
```

```
    password => 'password'
  );
END;
/
```

For more information on Oracle Cloud Infrastructure Auth Token authentication see CREATE_CREDENTIAL Procedure.

For example, to create Oracle Cloud Infrastructure Signing Key based credentials:

```
BEGIN
   DBMS_CLOUD.CREATE_CREDENTIAL (
       credential_name => 'DEF_CRED_NAME',
       user_ocid       =>
'ocid1.user.oc1..aaaaaaaauq54mi7zdyfhw33ozkwuontjceel7fok5nq3bf2vwetkpqsoa'
,
       tenancy_ocid    =>
'ocid1.tenancy.oc1..aabbbbbbaafcue47pqmrf4vigneebgbcmmoy5r7xvoypicjqqge32ew
nrcyx2a',
       private_key     =>
'MIIEogIBAAKCAQEAtUnxbmrekwgVac6FdWeRzoXvIpA9+0r1.....wtnNpESQQQ0QLGPD8NM//
JEBg=',
       fingerprint     =>
'f2:db:f9:18:a4:aa:fc:94:f4:f6:6c:39:96:16:aa:27');
END;
/
```

For more information on Oracle Cloud Infrastructure Signing Key based credentials see CREATE_CREDENTIAL Procedure.

Supported credential types:

- Data Pump Import supports Oracle Cloud Infrastructure Auth Token based credentials and Oracle Cloud Infrastructure Signing Key based credentials.

  For more information on Oracle Cloud Infrastructure Signing Key based credentials see CREATE_CREDENTIAL Procedure.

- Data Pump supports using an Oracle Cloud Infrastructure Object Storage pre-authenticated URL for the `dumpfile` parameter. When you use a pre-authenticated URL, providing the `credential` parameter is required and the `credential` value can be `NULL`. See Using Pre-Authenticated Requests for more information.

5. Run Data Pump Import with the *dumpfile* parameter set to the list of file URLs on your Cloud Object Storage and the *credential* parameter set to the name of the credential you created in the previous step.

> **Note:**
>
> Import the collection data using the option `CONTENT=DATA_ONLY`.

Specify the collection you want to import using the `INCLUDE` parameter. This is useful if a data file set contains the entire schema and the SODA collection you need to import is included as part of the dump file set.

Use `REMAP_DATA` to change any of the columns during import. This example shows using `REMAP_DATA` to change the version column method from `SHA256` to `UUID`.

```
impdp admin/password@db2022adb_high \
     directory=data_pump_dir \
     credential=def_cred_name \
     dumpfile= https://namespace-string.objectstorage.us-ashburn-1.oci.customer-
oci.com/n/namespace-string/b/bucketname/o/export%l.dmp \

     encryption_pwd_prompt=yes \
     SCHEMA=my_schema \
     INCLUDE=TABLE:\"= \'MyCollectionName\'\"  \
     CONTENT=DATA_ONLY \
     REMAP_DATA=my_schema.'\"MyCollectionName\"'.VERSION:SYS.DBMS_SODA.TO_UUID
```

Notes for Data Pump parameters:

- If during the export with `expdp` you used the `encryption_pwd_prompt=yes` parameter then use `encryption_pwd_prompt=yes` and input the same password at the `impdp` prompt that you specified during the export.
- The `dumpfile` parameter supports the `%L` and `%l` wildcards in addition to the legacy `%U` and `%u` wildcards. For example, `dumpfile=export%L.dmp`. Use the `%L` or `%l` wildcard for exports from Oracle Database Release 12.2 and higher. This wildcard expands the dumpfile file name into a 3-digit to 10-digit, variable-width incrementing integer, starting at 100 and ending at 2147483646.

  Use the legacy `%U` or `%u` wildcard for exports from Oracle Database prior to Release 12.2. If you use this option and more than 99 dump files are needed, you must specify multiple dumpfile names, each with the `%U` or `%u` parameter.

For `dumpfile`, this example uses the recommended URI format using OCI Dedicated Endpoints for commercial realm (OC1). The *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Object Storage Dedicated Endpoints, Regions and Availability Domains, and Understanding Object Storage Namespaces for more information.

In Oracle Data Pump version 19.6 and later, the *credential* argument authenticates Oracle Data Pump to the Cloud Object Storage service you are using for your source files. The `credential` parameter cannot be an Azure service principal, Amazon Resource Name (ARN), or a Google service account. See Accessing Cloud Resources by Configuring Policies and Roles for more information on resource principal based authentication.

The *dumpfile* argument is a comma delimited list of URLs for your Data Pump files.

For the best import performance use the `HIGH` database service for your import connection and set the `parallel` parameter to one quarter the number of ECPUs (.25 x **ECPU count**). If you are using OCPU compute model, set the parallel parameter to the number of OCPUs (1 x **OCPU count**).

For information on which database service name to connect to run Data Pump Import, see Manage Concurrency and Priorities on Autonomous Database.

For the dump file URL format for different Cloud Object Storage services, see DBMS_CLOUD Package File URI Formats.

> **Note:**
>
> To perform a full import or to import objects that are owned by other users, you need the `DATAPUMP_CLOUD_IMP` role.

For information on disallowed objects in Autonomous Database, see SQL Commands.

In this import example, the specification for the `REMAP_DATA` parameter uses the function `DBMS_SODA.TO_UUID` to generate UUID values. By default, for on-premise databases, the version column of a SODA collection is computed using SHA-256 hash of the document's content. On Autonomous Database the version column uses UUID generated values, which are independent of the document's content.

In this example the `REMAP_DATA` parameter uses the `DBMS_SODA.TO_UUID` function to replace the source collection version type with UUID versioning. If in the export dump file set that you are importing the `versionColumn.method` is already set to UUID, then the `REMAP_DATA` for this field is not required.

For detailed information on Oracle Data Pump Import parameters see *Oracle Database Utilities*.

The log files for Data Pump Import operations are stored in the directory you specify with the Data Pump Import `DIRECTORY` parameter. See Accessing the Log File for Data Pump Import for more information.

Developing RESTful Services in Autonomous Database

You can develop and deploy RESTful Services with native Oracle REST Data Services (ORDS) support on Autonomous Databases. Simple Oracle Document Access (SODA) for REST lets you use a database as a simple JSON document store.

- About Oracle REST Data Services in Autonomous Database
- Access RESTful Services and SODA for REST
  Each Autonomous Database includes Oracle REST Data Services (ORDS) that provides HTTPS interfaces for working with the contents of your Oracle Database in REST enabled schemas.
- Develop with Oracle REST Data Services on Autonomous Database
  Autonomous Database supports Oracle REST Data Services (ORDS).
- Use SODA for REST with Autonomous Database
  Autonomous Database supports Simple Oracle Document Access (SODA) for REST.
- Access Oracle REST Data Services, Oracle APEX, and Developer Tools Using a Vanity URL
- About Customer Managed Oracle REST Data Services on Autonomous Database

About Oracle REST Data Services in Autonomous Database

Oracle REST Data Services (ORDS) makes it easy to develop REST interfaces for relational data in a database.

ORDS is a mid-tier Java application that maps HTTP(S) verbs, such as GET, POST, PUT, DELETE, and so on, to database transactions, and returns any results as JSON data.

> **Note:**
>
> The Oracle REST Data Services (ORDS) application in Autonomous Database is preconfigured and fully managed. ORDS connects to the database using the `low` predefined database service with a fixed maximum number of connections (the number of connections for ORDS does not change based on the number of ECPUs (OCPUs if your database uses OCPUs). It is not possible to change the default ORDS configuration.
> See About Customer Managed Oracle REST Data Services on Autonomous Database for information on using an additional alternative ORDS deployment that enables flexible configuration options.

See Oracle REST Data Services for information on using Oracle REST Data Services.

See Database Service Names for Autonomous Database for information on the `low` database service.

Access RESTful Services and SODA for REST

Each Autonomous Database includes Oracle REST Data Services (ORDS) that provides HTTPS interfaces for working with the contents of your Oracle Database in REST enabled schemas.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To use Oracle REST Data Services and SODA for REST:

1. From the Autonomous Database details page select **Database Actions** and in the list click **View all database actions**.
2. On the Database Actions launchpad, under **Related Services**, click the **RESTful Services and SODA** card to see the base URL.
3. Click **Copy** to copy the URL.

If you are using Always Free Autonomous Database with Oracle Database 23ai Oracle recommends the following:

- For projects that were started using a database release prior to Oracle Database 21c, explicitly specify the metadata for the default collection as specified in the example in the section SODA Drivers.
- For projects started using release Oracle Database 21c or later, just use the default metadata.

See SODA Drivers for more information.

Develop with Oracle REST Data Services on Autonomous Database

Autonomous Database supports Oracle REST Data Services (ORDS).

Developing RESTful services is easy with the following development interfaces:

- Database Actions (SQL Developer Web): Use Database Actions on Autonomous Database to REST enable users. See Manage Users and User Roles on Autonomous Database - Connecting with Database Actions for more information.
- SQL Developer (desktop): With SQL Developer on your desktop, you can connect to your database and enable REST services access to tables and views, or develop custom RESTful Services based on your SQL and PL/SQL code. See Connect Oracle SQL Developer with a Wallet (mTLS) for more information.
- Oracle APEX: With APEX you can use the RESTful Services development pages to build and maintain your services and REST enabled objects. You can use the APEX SQL Workshop to access your Oracle RESTful Services and REST enabled objects. See How to Access RESTful Services for more information.

The Autonomous Database `ADMIN` account is REST Enabled. This allows for REST Services to be published in the ADMIN schemas and allows you to access Database Actions using the ADMIN database user account. Oracle recommends you create an application schema account for your RESTful Services and REST enabled objects. Services are secured using Database Authentication and your REST enabled schema.

The authenticated database user is only permitted access if the schema is REST enabled and the URL mapping for the request points to their own schema. A user is not authenticated when a request points to any other database schema. For example, the following request authenticated as the REST enabled schema `HR` is accessible:

```
GET /ords/hr/module/service/
```

However, when authenticated as the REST enabled schema `SCOTT`, the same request:

```
GET /ords/hr/module/service/
```

results in an error:

```
401 HTTP Unauthorized response/error
```

Any database user whose credentials are correct and meets these rules is authenticated and granted the ORDS, mid-tier, role: `SQL Developer`. The `SQL Developer` role enables the user to access any endpoint that requires the `SQL Developer` role.

See REST-Enable a Database Table in *Quick Start Guide* for information on how to enable a table for REST access.

Use SODA for REST with Autonomous Database

Autonomous Database supports Simple Oracle Document Access (SODA) for REST.

- Overview of Using SODA for REST

- Load Purchase-Order Sample Data Using SODA for REST
  Oracle provides a substantial set of JSON purchase-order documents, in plain-text file `POList.json`, as a JSON array of objects, where each such object represents a document.
- Use SODA for REST with OAuth Client Credentials
  You can access SODA for REST on Autonomous Database using OAuth authentication. Depending on your application, accessing SODA for REST with OAuth authentication can improve performance and security.

Overview of Using SODA for REST

SODA for REST is a predeployed REST service that can be used to store JSON documents in your database.

SODA enables flexible, NoSQL-style application development without having to use SQL. With SODA, JSON documents are stored in named collections and managed using simple CRUD operations (create, read, update and delete). And while SQL isn't required, JSON stored in SODA collections is still fully accessible from SQL when needed. For example, an operational application may be fully built using SODA (without SQL) but then the data may be later analyzed using SQL from outside of the application. Autonomous Database SODA gives application developers the best of the NoSQL and SQL worlds - fast, flexible, and scalable application development without losing the ability to leverage SQL for analytics and reporting.

SODA for REST is deployed in ORDS under the following URL pattern, where *schema* corresponds to a REST-enabled database schema.

```
/ords/schema/soda/latest/*
```

The following examples use the cURL command line tool (http://curl.haxx.se/) to submit REST requests to the database. However, other 3rd party REST clients and libraries should work as well. The examples use database schema `ADMIN`, which is REST-enabled. You can SODA for REST with cURL commands from the Oracle Cloud Shell.

This command creates a new collection named "fruit" in the `ADMIN` schema:

```
> curl -X PUT -u 'ADMIN:<password>' \
"https://example-db.adb.us-phoenix-1.oraclecloudapps.com/ords/admin/soda/latest/fruit"
```

These commands insert three JSON documents into the fruit collection:

```
> curl -X POST -u 'ADMIN:<password>' \
 -H "Content-Type: application/json" --data '{"name":"orange", "count":42}' \
 "https://example-db.adb.us-phoenix-1.oraclecloudapps.com/ords/admin/soda/latest/fruit"

{"items":[{"id":"6F7E5C60197E4C8A83AC7D7654F2E375"...

> curl -X POST -u 'ADMIN:<password>' \
 -H "Content-Type: application/json" --data '{"name":"pear", "count":5}' \
 "https://example-db.adb.us-phoenix-1.oraclecloudapps.com/ords/admin/soda/latest/fruit"

{"items":[{"id":"83714B1E2BBA41F7BA4FA93B109E1E85"...

> curl -X POST -u 'ADMIN:<password>' \
 -H "Content-Type: application/json" \
 --data '{"name":"apple", "count":12, "color":"red"}' \
 "https://example-db.adb.us-phoenix-1.oraclecloudapps.com/ords/admin/soda/latest/fruit"
```

```
{"items":[{"id":"BAD7EFA9A2AB49359B8F5251F0B28549"...
```

This example retrieves a stored JSON document from the collection:

```
> curl -X POST -u 'ADMIN:<password>' \
 -H "Content-Type: application/json" --data '{"name":"orange"}' \
 "https://example-db.adb.us-phoenix-1.oraclecloudapps.com/ords/admin/soda/latest/fruit?
action=query"

{
  "items": [
    {
      "id":"6F7E5C60197E4C8A83AC7D7654F2E375",
      "etag":"57215643953D7C858A7CB28E14BB48549178BE307D1247860AFAB2A958400E16",
      "lastModified":"2019-07-12T19:00:28.199666Z",
      "created":"2019-07-12T19:00:28.199666Z",
      "value":{"name":"orange", "count":42}
    }
  ],
  "hasMore":false,
  "count":1
}
```

This SQL query accesses the fruit collection:

```
SELECT
     f.json_document.name,
     f.json_document.count,
     f.json_document.color
FROM fruit f;
```

The query returns these three rows:

```
name      count     color
--------- --------- -------
orange    42        null
pear      5         null
apple     12        red
```

> **Note:**
>
> If you are using Always Free Autonomous Database with Oracle Database 23ai, Oracle recommends the following:
> For projects that were started using a database release prior to Oracle Database 21c, explicitly specify the metadata for the default collection as specified in the example in the section SODA Drivers. For projects started using release Oracle Database 21c or later, just use the default metadata. See SODA Drivers for more information.

These examples show a subset of the SODA and SQL/JSON features. See the following for more information:

- SODA for REST for complete information on Simple Oracle Document Access (SODA)
- SODA for REST HTTP Operations for information on the SODA for REST HTTP operations

Load Purchase-Order Sample Data Using SODA for REST

Oracle provides a substantial set of JSON purchase-order documents, in plain-text file `POList.json`, as a JSON array of objects, where each such object represents a document.

The following examples use the cURL command line tool (http://curl.haxx.se/) to submit REST requests to the database. However, other 3rd party REST clients and libraries should work as well. The examples use database schema `ADMIN`, which is REST-enabled. You can use SODA for REST with cURL commands from the Oracle Cloud Shell.

You can load this sample purchase-order data set into a collection `purchaseorder` on your Autonomous Database with SODA for REST, using these curl commands:

```
curl -X GET "https://raw.githubusercontent.com/oracle/db-sample-schemas/
master/order_entry/POList.json" -o POList.json

curl -X PUT -u 'ADMIN:password' \
"https://example-db.adb.us-phoenix-1.oraclecloudapps.com/ords/admin/soda/
latest/purchaseorder"

curl -X POST -H -u 'ADMIN:password' 'Content-type: application/json' -d
@POList.json \
"https://example-db.adb.us-phoenix-1.oraclecloudapps.com/ords/admin/soda/
latest/purchaseorder?action=insert"
```

You can then use this purchase-order data to try out examples in *Oracle Database JSON Developer's Guide*.

For example, the following query selects both the `id` of a JSON document and values from the JSON purchase-order collection stored in column `json_document` of table `purchaseorder`. The values selected are from fields `PONumber`, `Reference`, and `Requestor` of JSON column `json_document`, which are projected from the document as virtual columns (see SQL NESTED Clause Instead of JSON_TABLE for more information).

```
SELECT id, t.*
  FROM purchaseorder
    NESTED json_document COLUMNS(PONumber, Reference, Requestor) t;
```

See the following for more information:

- SODA for REST for complete information on Simple Oracle Document Access (SODA)
- SODA for REST HTTP Operations for information on the SODA for REST HTTP operations

Use SODA for REST with OAuth Client Credentials

You can access SODA for REST on Autonomous Database using OAuth authentication. Depending on your application, accessing SODA for REST with OAuth authentication can improve performance and security.

Perform the following steps to use OAuth authentication to provide limited access to SODA for REST on Autonomous Database:

1. As the ADMIN user, access Database Actions and create a user with the required privileges.
   a. Access Database Actions as ADMIN.

      See Access Database Actions as ADMIN for more information.

   b. In Database Actions, click ☰ to show the available actions.
   c. In Database Actions, under **Administration** select **Database Users**.
   d. Click **Create User**.
   e. In the **Create User** area, on the **User** tab enter **User Name** and a **Password** and confirm the password.
   f. Select **Web Access**.
   g. In the **Create User** area, select the **Granted Roles** tab and grant `DWROLE` to the user.
   h. Click **Create User**.

   See Manage Users and User Roles on Autonomous Database - Connecting with Database Actions for more information.
2. Use a SQL worksheet in Database Actions to grant user privileges required to load data.
   a. Access Database Actions as ADMIN.

      See Access Database Actions as ADMIN for more information.

   b. In Database Actions, click ☰ to show the available actions.
   c. In Database Actions, under **Development** click **SQL** to open a SQL worksheet.
   d. Grant user privileges required to load data to the user from Step 1.

      ```
      GRANT UNLIMITED TABLESPACE TO user_name;
      ```

   See Manage User Privileges on Autonomous Database - Connecting with a Client Tool for more information.
3. Sign out as the ADMIN user.
4. Sign in to Database Actions as the user that is setting up to use OAuth authentication.
5. In Database Actions, use a SQL worksheet to register the OAuth client.
   a. Register the OAuth client.

      For example, enter the following commands into the SQL worksheet, where you supply the appropriate values for your user and your client application.

      ```
      BEGIN
        OAUTH.create_client(
          p_name            => 'my_client',
      ```

```
    p_grant_type      => 'client_credentials',
    p_owner           => 'Example Company',
    p_description     => 'A client for my SODA REST resources',
    p_support_email   => 'user_name@example.com',
    p_privilege_names => 'my_priv'
  );

  OAUTH.grant_client_role(
    p_client_name => 'my_client',
    p_role_name   => 'SQL Developer'
  );

  OAUTH.grant_client_role(
    p_client_name => 'my_client',
    p_role_name   => 'SODA Developer'
  );
  COMMIT;
END;
/
```

b. In the SQL worksheet, click **Run Script** to run the command.

See OAUTH PL/SQL Package Reference for more information.

This registers a client named `my_client` to access the `my_priv` privilege using OAuth client credentials.

6. Obtain the `client_id` and `client_secret` required to generate the access token.

For example, in the SQL worksheet run the following command:

```
SELECT id, name, client_id, client_secret FROM user_ords_clients;
```

7. Obtain the access token. To get an access token you send a `REST GET` request to *`database_ORDS_urluser_name`*`/oauth/token`.

The *`database_ORDS_url`* is available from Database Actions, under **Related Services**, on the **RESTful Services and Soda** card. See Access RESTful Services and SODA for REST for more information.

In the following command, use the `client_id` and the `client_secret` you obtained in Step 6.

The following example uses the `cURL` command line tool (http://curl.haxx.se/) to submit REST requests to Autonomous Database. However, other 3rd party REST clients and libraries should work as well.
You can use the `cURL` command line tool to submit the `REST GET` request. For example:

```
> curl -i -k --user SBA-iO9Xe12cdZHYfryBGQ..:vvUQ1AagTqAqdA2oN7afSg.. --
data "grant_type=client_credentials"https://mqssyowmqvgac1y-
doc.adb.region.oraclecloudapps.com/ords/user_name/oauth/token
HTTP/1.1 200 OK
Date: Mon, 22 Jun 2020 15:17:11 GMT
Content-Type: application/jsonTransfer-Encoding: chunked
Connection: keep-alive
X-Frame-Options: SAMEORIGIN
```

```
{"access_token":"JbOKtAuDgEh2DXx0QhvPGg","token_type":"bearer","expires_in"
:3600}
```

To specify both the `client_id` and the `client_secret` with the curl `--user` argument, enter a colon to separate the `client_id` and the `client_secret`. If you only specify the user name, `client_id`, curl prompts for a password and you can enter the `client_secret` at the prompt.

8. Use the access token to access the protected resource.

   The token obtained in the previous step is passed in the Authorization header. For example:

```
> curl -i -H "Authorization: Bearer JbOKtAuDgEh2DXx0QhvPGg" -X GET https://
database_id.adb.region.oraclecloudapps.com/ords/user_name/soda/latest
HTTP/1.1 200 OK
Date: Mon, 22 Jun 2020 15:20:58 GMT
Content-Type: application/json
Content-Length: 28
Connection: keep-alive
X-Frame-Options: SAMEORIGIN
Cache-Control: private,must-revalidate,max-age=0


{"items":[],"hasMore":false}
```

See Configuring Secure Access to RESTful Services for complete information on secure access to RESTful Services.

Access Oracle REST Data Services, Oracle APEX, and Developer Tools Using a Vanity URL

By default you access Oracle REST Data Services (ORDS) endpoints, Oracle APEX apps, and developer tools on Autonomous Database using the `oraclecloudapps.com` domain name.

To access Oracle REST Data Services (ORDS) endpoints, you can optionally configure a vanity URL or custom domain name that is easy to remember to help promote your brand identity.

After you acquire a desired domain name and matching SSL certificate from a vendor of your choice, deploy an Oracle Cloud Infrastructure Load Balancer in your Virtual Cloud Network (VCN) using your Autonomous Database as the backend. Your Autonomous Database instance must be configured with a private endpoint in the same VCN. See Configuring Network Access with Private Endpoints for more information.

To learn more, see the following:

- Introducing Vanity URLs for APEX and ORDS on Oracle Autonomous Database
- Automate Vanity URL Configuration Using Terraform

About Customer Managed Oracle REST Data Services on Autonomous Database

Optionally, you can configure Autonomous Database to use Oracle REST Data Services (ORDS) running in a customer managed environment.

When you provision an Autonomous Database instance, by default Oracle REST Data Services (ORDS) is preconfigured and available for the instance. With the default Oracle REST Data Services, Oracle performs any required configuration, patching, and maintenance. When you use the default ORDS on Autonomous Database, you cannot modify any of the ORDS configuration options.

Use a customer managed environment if you want manual control of the configuration and management of Oracle REST Data Services. For example, use this option when your applications require larger connection pools or if you need more control over the ORDS configuration options.

When ORDS runs in a customer managed environment, you are responsible for configuration, patching, and maintenance of ORDS. After you configure Autonomous Database to use your customer managed ORDS in addition to the existing autonomously managed ORDS, you can route ORDS HTTPS traffic through your environment. The default Autonomous Database web server and ORDS are still running and ORDS traffic goes to the ORDS running in the customer managed environment. This provides an additional and alternative HTTPS solution for Autonomous Database.

Installing and configuring a customer managed environment for ORDS allows you to run ORDS with configuration options that are not possible using the default Oracle managed ORDS available with Autonomous Database.

See Installing and Configuring Customer Managed ORDS on Autonomous Database for more information.

Installing and configuring a customer managed environment for ORDS is only supported with Autonomous Database Serverless. Your Autonomous Database must be configured for one of the following workload types: Data Warehouse, Transaction Processing, or JSON Database.

> **Note:**
>
> Oracle REST Data Services 19.4.6 or higher is required to use a customer managed environment for ORDS with Autonomous Database. Customer managed environment for ORDS is not supported if your database is configured for APEX workload type.

Using Oracle Database API for MongoDB

Oracle Database API for MongoDB makes it possible to connect to Oracle Autonomous Database using MongoDB language drivers and tools.

Oracle Database API for MongoDB leverages the converged database capabilities of an Autonomous Database to manage multiple data types, including JSON data, within a single database. For example, these converged database capabilities allow you to use SQL to query or update JSON data.

> **Note:**
>
> MongoDB requires that you configure network access to use ACLs or that you define a private endpoint for the Autonomous Database instance.

See Oracle Database API for MongoDB for more information.

See About Autonomous JSON Database for more information.

See About Autonomous Database Workload Types for more information.

Topics

- Configure Access for MongoDB and Enable MongoDB
  Oracle Database API for MongoDB enables you to use an Oracle Autonomous Database as the data store.
- User Management for MongoDB
  Oracle Database API for MongoDB enables you to use an Oracle Autonomous Database as the data store. If you want or need to use an existing Autonomous Database for this purpose, here is the workflow.
- Create a Test Autonomous Database User for MongoDB
- Connect MongoDB Applications to Autonomous Database
  Connecting your MongoDB application to Autonomous Database includes several steps, depending upon your requirements.

Configure Access for MongoDB and Enable MongoDB

Oracle Database API for MongoDB enables you to use an Oracle Autonomous Database as the data store.

To use the MongoDB API you can create and configure a new Autonomous Database or modify the configuration of an existing Autonomous Database.

> **Note:**
>
> MongoDB requires that you configure network access to use ACLs or that you define a private endpoint for the Autonomous Database instance. In addition to configuring the network access, you must enable MongoDB API on the Autonomous Database instance.

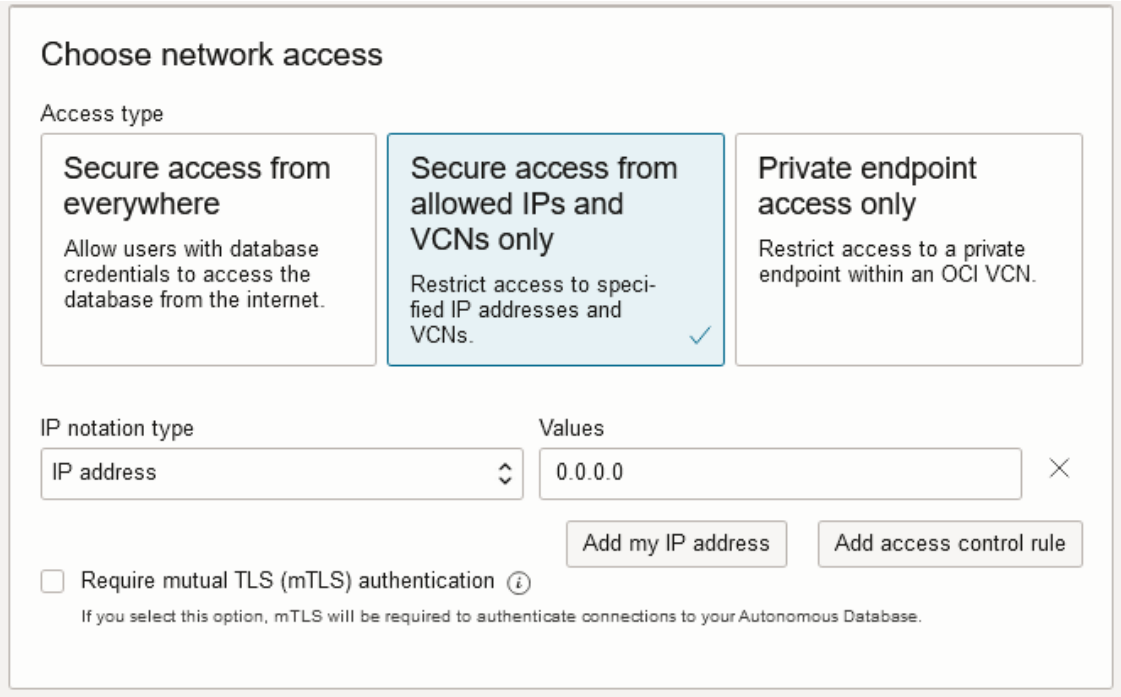
- Configure Access for MongoDB
  To use the MongoDB API, you can create and configure a new Autonomous Database or modify the configuration of an existing Autonomous Database by configuring ACLs or by defining a private endpoint.
- Enable MongoDB API on Autonomous Database
  After you configure the network access for the Autonomous Database instance, enable the MongDB API.

Configure Access for MongoDB

To use the MongoDB API, you can create and configure a new Autonomous Database or modify the configuration of an existing Autonomous Database by configuring ACLs or by defining a private endpoint.

Configure a New Autonomous Database for MongoDB

Follow the steps in Provision or Clone an Autonomous Database, up to the point where you select your Network Access Type.



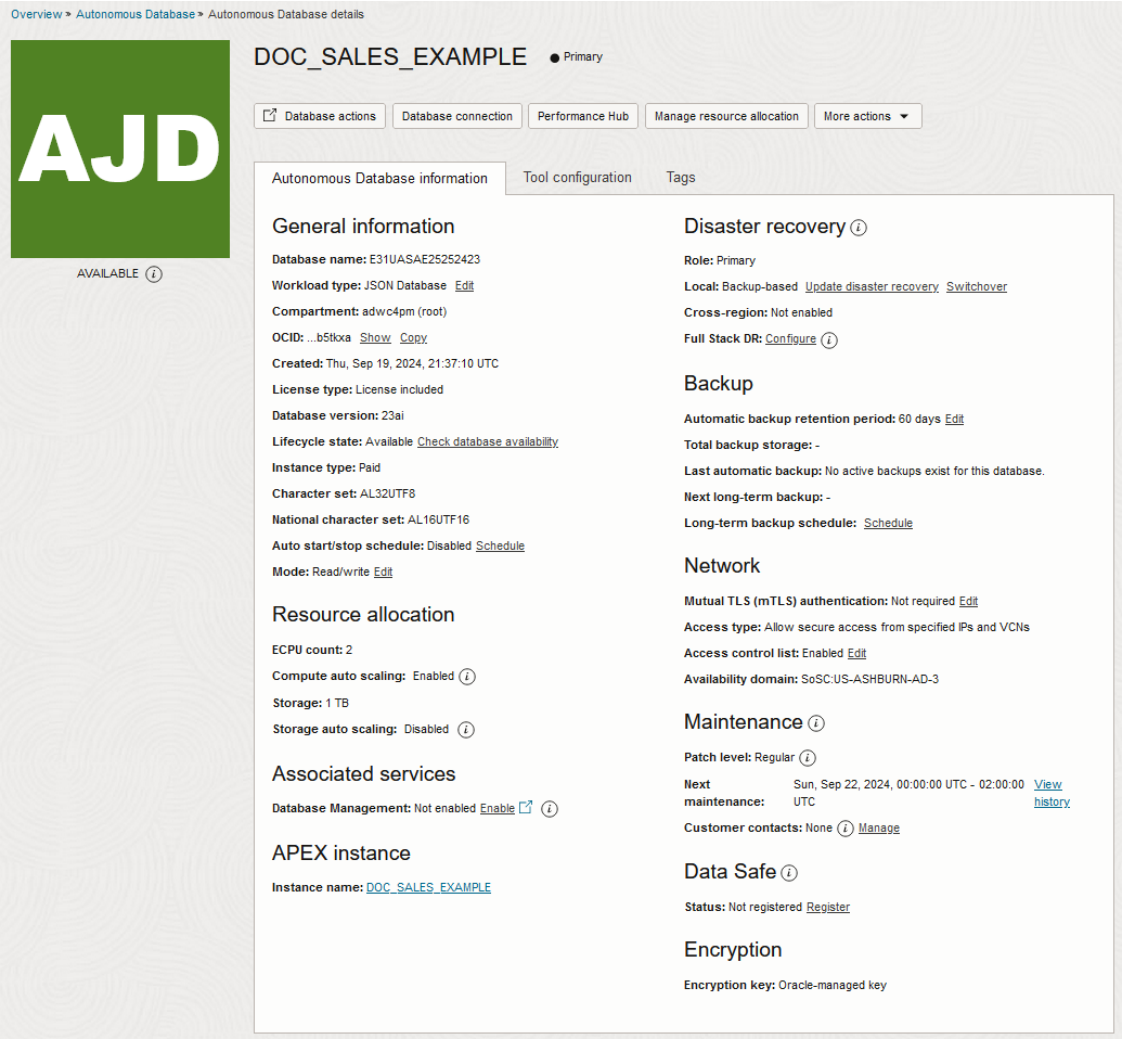
At this point, to use Oracle Database API for MongoDB configure secure access by selecting and configuring one of these network access types:

- **Secure access from allowed IPs and VCNs only**
- **Private endpoint access only**

  See Configure Network Access with Private Endpoints for information on configuring an Autonomous Database instance with a private endpoint.

Configure an Existing Autonomous Database for MongoDB

Open the Oracle Cloud Infrastructure Console for your Autonomous Database instance.



> **Note:**
>
> To use Oracle Database API for MongoDB the Network must be configured and the **Access type** must be either: **Secure access from allowed IPs and VCNs only** or **Private endpoint access only**.

Access Control List (ACL) Setup

See Configure Access Control Lists for an Existing Autonomous Database Instance for more information.

To configure your ACL for an IP address, you will need to obtain the public IP address. There are several ways to show your public IP address:

- In the choose network access area, click **Add my IP address**. This copies your IP address into the **Values** field.
- After disabling any VPN, use the WhatIsMyIP website.
- After disabling any VPN, use the curl command: `curl -s https://ifconfig.me`.

> **Note:**
>
> Public IP addresses may change. Any change to your public IP address will require a change in the ACL. If you are unable to access your database, your ACL should be something you check.

ACLs Types and Use Cases

ACL Type	Use Case	Comment
IP Address	Local development laptops sharing the same public IP address	Easiest way to get started. Any laptop connected on this LAN will have access to the database with the database credentials.
CIDR Block	Local development laptop	Using IPv4/32 notation
IP Addresses separated by commas	Small number of local development laptops connected on distinct LANs (having distinct public IP addresses)	Can be tedious to manage with 10+ laptops.
CIDR Block	Local development laptops connected on the same subnet exposed to Internet (each laptop has its own public IP Address)	Rely on CIDR Block notation. See calculator here for more information. Example: `89.84.109.0/24` gives 256 possible IP addresses from `89.84.109.0` to `89.84.109.255`
VCN with CIDR Block	For testing, production, or CI/CD pipeline hosted on OCI having their own VCN and Compute instances	Assign OCI compartment per environment type.
Mixing IP Address and VCN with CIDR Block	Local development laptop accessing a test Autonomous Database with connections from the testing environment or CI/CD pipeline	A common configuration option for on-going development work.

Enable MongoDB API on Autonomous Database

After you configure the network access for the Autonomous Database instance, enable the MongDB API.

To enable the MongoDB API for an existing instance:

1. On the Autonomous Database details page, select the **Tool configuration** tab.
2. Click **Edit tool configuration**.
3. In the MongoDB API row, select in the **Enable tool** column to show **Enabled**.
4. Click **Apply**.

The **Lifecycle state** changes to updating until MongoDB is enabled.

You can also enable the MongoDB API when you provision or clone an instance by selecting **Show advanced options** and selecting the **Tools** tab.

See Manage Autonomous Database Built-in Tools for more information.

User Management for MongoDB

Oracle Database API for MongoDB enables you to use an Oracle Autonomous Database as the data store. If you want or need to use an existing Autonomous Database for this purpose, here is the workflow.

Oracle Database API for MongoDB enables the mapping of Autonomous Database objects to MongoDB objects as follows:

MongoDB Object	Oracle Autonomous Database Object
database	schema
collection	table
document	document (in a column)

For example, you could create a collection using the Oracle Database API for MongoDB as follows:

```
use scott;
db.createCollection('fruit');
```

A table named **FRUIT** is created in the schema **SCOTT**.

When you connect to the Oracle Database API for MongoDB, you authenticate using an Autonomous Database username and password. This authenticated connection then accesses collections within the corresponding schema. This user must meet the following requirements:

- The user's schema must be ORDS-enabled, which is sometimes referred to as enabled for Web Access. See Basic Setup to Enable ORDS Database API for more information.
- The user must have the following roles and privileges: `SODA_APP`, `CREATE TABLE`, and `CREATE SESSION`. See Manage User Roles and Privileges on Autonomous Database for more information.
- The user has a quota on tablespace **DATA**. See Create Users on Autonomous Database for more information.

> **Note:**
>
> The role `DWROLE` in the Autonomous Database contains these roles, among others.

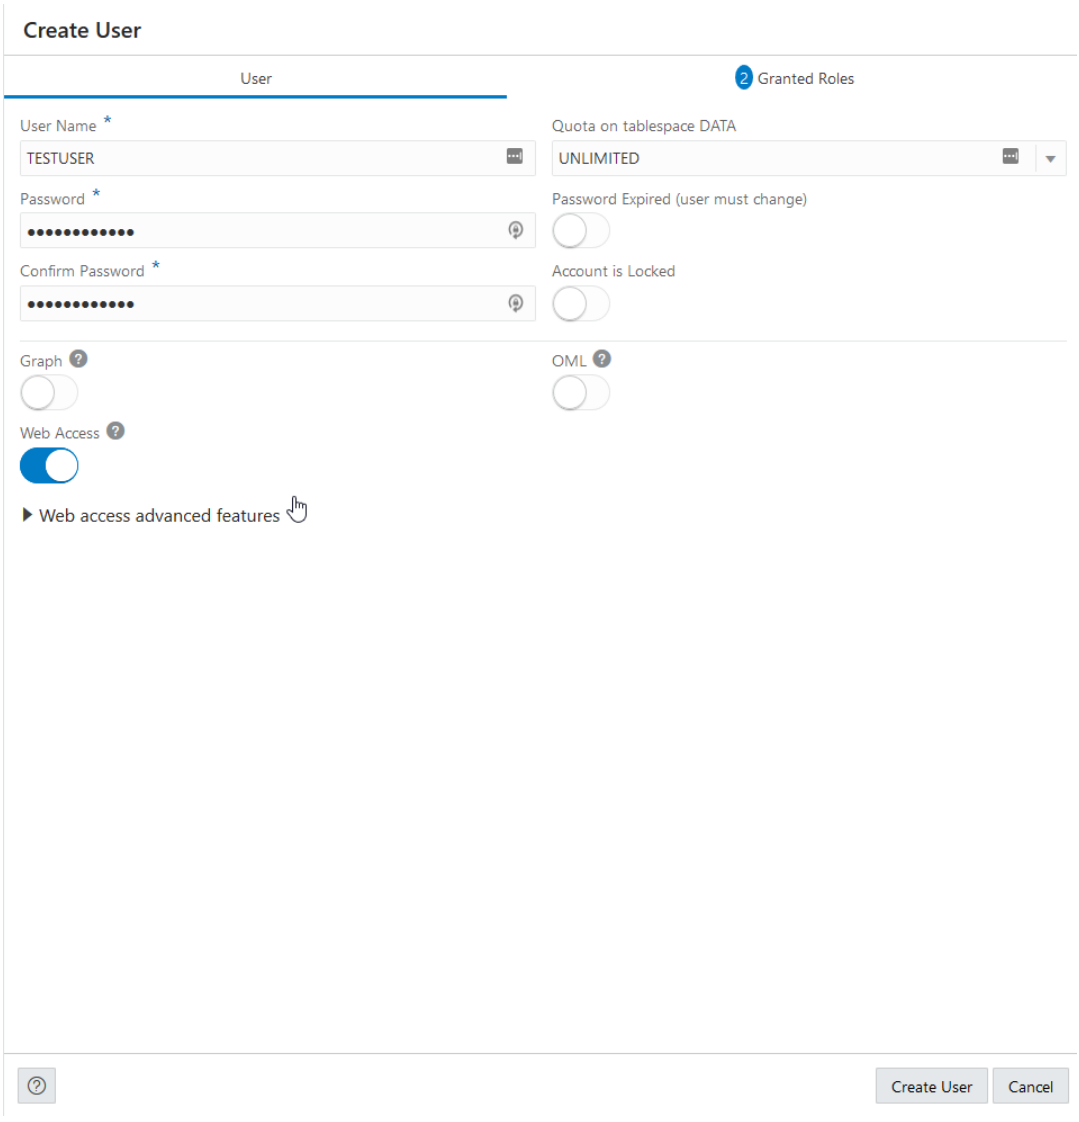
Access to schemas not granted to the user is prohibited. For example, the user **SCOTT** can only access collections in the schema **SCOTT**. There is one exception. If the authenticated user has the Autonomous Database privileges `CREATE USER`, `ALTER USER` and `DROP USER`, that user can access any ORDS-enabled schema.

Additionally, a user with these privileges can implicitly create schemas. That is, when the user creates a collection in a database that does not exist, the schema will be created automatically. See Oracle Database API for MongoDB for more information.

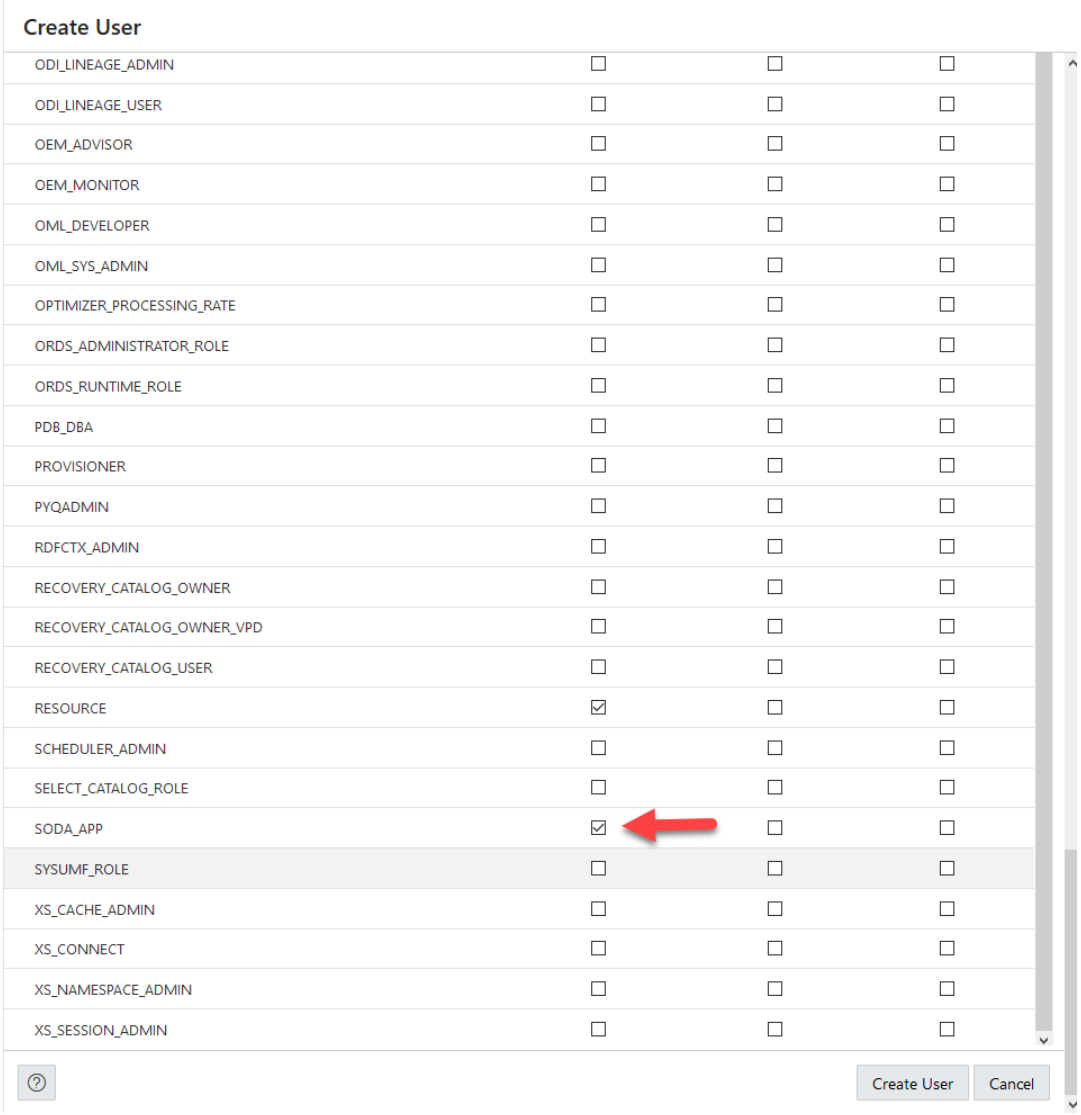
Create a Test Autonomous Database User for MongoDB

The steps that follow use Database Actions to create a test user with the proper roles. You can also use Database Actions SQL or other SQL command line utilities to execute the SQL statements directly. See Create Users on Autonomous Database - Connecting with a Client Tool for more information.

1. Open the Oracle Cloud Infrastructure Console for your Autonomous Database.
2. Select **Database Actions**.
3. From the Database Actions Launchpad, select **Administration** > **Database Users**.
4. Select **+ Create User** on the Database Users page.
5. Enter a **User Name** and **Password** for your test user. Select **Web Access**, and set the **Quota on tablespace DATA**.



6. Select the **Granted Roles** tab.
7. In addition to the default roles, select and add the `SODA_APP` role for the user.



8. Select the **Create User** button.

You can user this user, or a similarly authenticated user, for testing purposes. See Test Connection Using the Command Line for more information.

Connect MongoDB Applications to Autonomous Database

Connecting your MongoDB application to Autonomous Database includes several steps, depending upon your requirements.

Topics

- Retrieve the Autonomous Database Connection String from Database Actions
- Test Connection Using the Command Line
- Test Connection Using a Node.js Application
- Retrieve the Autonomous Database MongoDB Connection String
- Test Connection Using the Command Line
- Test Connection Using a Node.js Application

Retrieve the Autonomous Database MongoDB Connection String

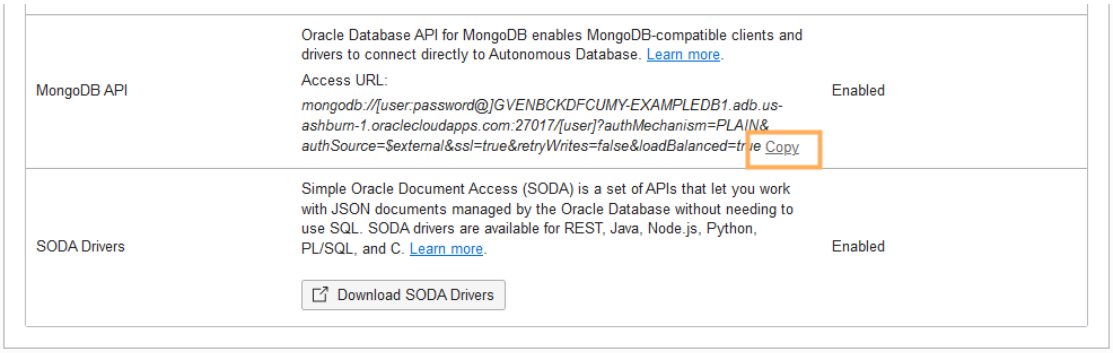
You can retrieve the MongoDB API connection string from the Oracle Cloud Infrastructure Console.

After you obtain the MongoDB API connection string you can use the MongoDB Shell, which is a command-line utility, to connect and query your data.

First you must configure network access and enable the MongoDB API. See Configure Access for MongoDB for more information.

To retrieve the MongoDB API connection string:

1. On the Autonomous Database details page, select the **Tool configuration** tab.
2. In the MongoDB API row, under Access URL click **Copy**.



- Retrieve the Autonomous Database Connection String from Database Actions

Retrieve the Autonomous Database Connection String from Database Actions

Use the MongoDB Shell which is a command-line utility used to connect and query your data.

1. You can retrieve the connection string for your Autonomous Database instance with Database Actions.

   See Access Database Actions as ADMIN for more information.
2. On the Database Actions Launchpad, under **Related Services**, click **Oracle Database API for MongoDB**.
3. On the **Oracle Database API for MongoDB** page click **Copy**.

Test Connection Using the Command Line

1. Login as the test user. See Create a Test Autonomous Database User for MongoDB for more information.

```
$ mongosh --tls --tlsAllowInvalidCertificates 'mongodb://
TESTUSER:<PASSWORD>@<database URL>.
<OCI region>.oraclecloudapps.com:27017/admin?
authMechanism=PLAIN&authSource=$external&ssl=true&loadBalanced=false'
Current Mongosh Log ID: 614c9e2a01e3575c8c0b2ec7
Connecting to:          mongodb://TESTUSER:<PASSWORD>@<database
URL>.<OCIregion>.oraclecloudapps.com:27017/admin?
authMechanism=PLAIN&authSource=$external&tls=true&loadBalanced=false
```

```
Using MongoDB:                  3.6.2
Using Mongosh:                  1.0.7
For mongosh info see: https://docs.mongodb.com/mongodb-shell/admin
> show dbs
testuser    0 B
>
```

Note:

Use URI percent-encoding to replace any reserved characters in your connection-string URI — in particular, characters in your username and password. These are the reserved characters and their percent encodings:

!	#	$	%	&	'	(	)	*	+
%21	%23	%24	%25	%26	%27	%28	%29	%2A	%2B

,	/	:	;	=	?	@	[	]
%2C	%2F	%3A	%3B	%3D	%3F	%40	%5B	%5D

For example, if your username is `RUTH` and your password is `@least1/2#?` then your MongoDB connection string to server *<server>* might look like this:

```
'mongodb://RUTH:%40least1%2F2%23%3F@<server>:27017/ruth/ ...'
```

Depending on the tools or drivers you use, you might be able to provide a username and password as separate parameters, instead of as part of a URI connection string. In that case you likely won't need to encode any reserved characters they contain.

See also:

- Percent Encoding - Reserved Characters
- Uniform Resource Identifier (URI): Generic Syntax

2. Create a collection and insert documents into your collection.

```
testuser> show collections

testuser> db.createCollection( 'fruit' )
{ ok: 1 }
testuser> show collections
fruit
testuser> db.fruit.insertOne( {name:"orange", count:42} )
{
  acknowledged: true,
  insertedId: ObjectId("614ca31fdab254f63e4c6b47")
}
testuser> db.fruit.insertOne( {name:"apple", count:12, color: "red"} )
{
  acknowledged: true,
  insertedId: ObjectId("614ca340dab254f63e4c6b48")
```

```
}
testuser> db.fruit.insertOne( {name:"pear", count:5} )
{
  acknowledged: true,
  insertedId: ObjectId("614ca351dab254f63e4c6b49")
}
testuser>
```

3. Query the collection using a SQL client such as Database Actions.

```
SELECT
     f.json_document.name,
     f.json_document.count,
     f.json_document.color
FROM fruit f;
```

Test Connection Using a Node.js Application

1. Download Node.js. If you have already downloaded or installed Node.js for your environment, you can skip this step.

```
$ wget https://nodejs.org/dist/latest-v14.x/node-v14.17.5-linux-x64.tar.xz
```

2. Extract the contents of the Node,js archive. If you have already downloaded or installed Node.js for your environment, you can skip this step.

```
$ tar -xf node-v14.17.5-linux-x64.tar.xz
```

3. Configure the `PATH` environment variable. If you have already downloaded or installed Node.js for your environment, you can skip this step.

```
$ export PATH="`pwd`"/node-v14.17.5-linux-x64/bin:$PATH
```

4. Test your connection with a Javascript example.
   a. Create a new directory.

```
$ mkdir autonomous_mongodb
$ cd autonomous_mongodb
$ npm init -y
```

   b. Install mongodb dependency.

```
$ npm install mongodb
```

   c. Create a JavaScript application named `connect.js`.

```
const { MongoClient } = require("mongodb");
const uri =
   "mongodb://TESTUSER:<PASSWORD>@<Database URI>.<OCI
region>.oraclecloudapps.com:27017/admin?
authMechanism=PLAIN&authSource=$external&ssl=true&loadBalanced=false";

const client = new MongoClient(uri);
```

```
async function run() {
  try {
        await client.connect();

        const database = client.db('admin');
        const movies = database.collection('movies');

        // Insert a movie
        const doc = { title: 'Back to the Future',
                      year: 1985, genres: ['Adventure', 'Comedy', 'Sci-
Fi'] }
        const result = await movies.insertOne(doc);

        // Query for a movie that has the title 'Back to the Future' :)
        const query = { title: 'Back to the Future' };
        const movie = await movies.findOne(query);

        console.log(movie);
  } finally {
     // Ensures that the client will close when you finish/error
     await client.close();
  }
}
run().catch(console.dir);
```

> **Note:**
>
> Use URI percent-encoding to replace any reserved characters in your connection-string URI — in particular, characters in your username and password. These are the reserved characters and their percent encodings:
>
> | ! | # | $ | % | & | ' | (|) | * | + |
> |---|---|---|---|---|---|---|---|---|---|
> | %21 | %23 | %24 | %25 | %26 | %27 | %28 | %29 | %2A | %2B |
>
> | , | / | : | ; | = | ? | @ | [|] |
> |---|---|---|---|---|---|---|---|---|
> | %2C | %2F | %3A | %3B | %3D | %3F | %40 | %5B | %5D |
>
> For example, if your username is `RUTH` and your password is `@least1/2#?` then your MongoDB connection string to server *<server>* might look like this:
>
> ```
> 'mongodb://RUTH:%40least1%2F2%23%3F@<server>:27017/ruth/ ...'
> ```
>
> Depending on the tools or drivers you use, you might be able to provide a username and password as separate parameters, instead of as part of a URI connection string. In that case you likely won't need to encode any reserved characters they contain.
>
> See also:
>
> - Percent Encoding - Reserved Characters
> - Uniform Resource Identifier (URI): Generic Syntax

d. Run the example. The output should look similar to the following.

```
$ node connect
{
  _id: new ObjectId("611e3266005202371acf27c1"),
  title: 'Back to the Future',
  year: 1985,
  genres: [ 'Adventure', 'Comedy', 'Sci-Fi' ]
}
```

Send Email and Notifications on Autonomous Database

There are a number of options for sending email on Autonomous Database. You can also send notifications to a Slack channel, to MSTeams, or to an OCI Notifications topic.

- Send Email on Autonomous Database
  There are a number of options for sending email on Autonomous Database. You can also send text messages or the output of an SQL query to a Slack or MSTeams channel.
- Send Slack Notifications from Autonomous Database
  Describes how to configure Slack so that you can send messages, alerts, or output of a query from Autonomous Database to a Slack Channel. Also describes the procedures you use to send Slack notifications.

- Send Messages to OCI Notifications Service from Autonomous Database
  Describes how to send messages from Autonomous Database to an OCI Notifications service topic.
- Send Microsoft Teams Notifications from Autonomous Database
  Describes how to configure Microsoft Teams so that you can send messages, alerts, or output of a query from Autonomous Database to a Microsoft Teams Channel. Also describes the procedures you use to send Microsoft Teams notifications.

Send Email on Autonomous Database

There are a number of options for sending email on Autonomous Database. You can also send text messages or the output of an SQL query to a Slack or MSTeams channel.

- Send Email with Email Delivery Service on Autonomous Database
  Describes the steps to send email using `UTL_SMTP` on Autonomous Database.
- Send Email with an Email Provider on a Private Endpoint
  Describes the steps to send email with an email provider that is on Private Endpoint.
- Use Credential Objects to set SMTP Authentication
  Describes how to pass a credential objects to `UTL_SMTP.SET_CREDENTIAL` APIs.
- Send Email from Autonomous Database Using DBMS_CLOUD_NOTIFICATION
  Use the `DBMS_CLOUD_NOTIFICATION` package to send messages and query results as email.

Send Email with Email Delivery Service on Autonomous Database

Describes the steps to send email using `UTL_SMTP` on Autonomous Database.

To send email with Oracle Cloud Infrastructure Email Delivery Service:

1. Identify your SMTP connection endpoint for Email Delivery. You may need to subscribe to additional Oracle Cloud Infrastructure regions if Email Delivery is not available in your current region.

   For example, select one of the following for the SMTP connection endpoint:

   - smtp.us-phoenix-1.oraclecloud.com
   - smtp.us-ashburn-1.oraclecloud.com
   - smtp.email.uk-london-1.oci.oraclecloud.com
   - smtp.email.eu-frankfurt-1.oci.oraclecloud.com

   > **Note:**
   >
   > If the SMTP connection endpoint is in a different region from your Autonomous Database, then you must create a request at Oracle Cloud Support to have Cloud Operations enable that region's Email Delivery service for your Autonomous Database.

   See Configure SMTP Connection for more information.
2. Generate SMTP credentials for Email Delivery. `UTL_SMTP` uses credentials to authenticate with Email Delivery servers when you send email.

   See Generate SMTP Credentials for a User for more information.

3. Create an approved sender for Email Delivery. Complete this step for all email addresses you use as the "`From`" with `UTL_SMTP.MAIL`.

   See Managing Approved Senders for more information.

4. Allow SMTP access for ADMIN user by appending an Access Control Entry (ACE).

   For example:

```
BEGIN
  -- Allow SMTP access for user ADMIN
  DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
    host => 'www.us.example.com',
    lower_port => 587,
    upper_port => 587,
    ace => xs$ace_type(privilege_list => xs$name_list('SMTP'),
                       principal_name => 'ADMIN',
                       principal_type => xs_acl.ptype_db));
END;
/
```

5. Create a PL/SQL procedure to send email.

   For example, see the sample code shown in SMTP Send Email Sample Code.

6. Send a test email using the PL/SQL procedure you created in step 5.

   For example:

```
execute send_mail('taylor@example.com', 'Email from Oracle Autonomous
Database', 'Sent using UTL_SMTP');
```

See UTL_SMTP for information on `UTL_SMTP`.

See PL/SQL Package Notes for Autonomous Database for `UTL_SMTP` restrictions with Autonomous Database.

- SMTP Send Email Sample Code
  Shows sample code for sending email with `UTL_SMTP` on Autonomous Database.

SMTP Send Email Sample Code

Shows sample code for sending email with `UTL_SMTP` on Autonomous Database.

```
CREATE OR REPLACE PROCEDURE SEND_MAIL (
  msg_to varchar2,
  msg_subject varchar2,
  msg_text varchar2 )
IS

  mail_conn utl_smtp.connection;
  username varchar2(1000):= 'ocid1.user.oc1.username';
  passwd varchar2(50):= 'password';
  msg_from varchar2(50) := 'adam@example.com';
  mailhost VARCHAR2(50) := 'smtp.us-ashburn-1.oraclecloud.com';

BEGIN
  mail_conn := UTL_smtp.open_connection(mailhost, 587);
  utl_smtp.starttls(mail_conn);
```

```
  UTL_SMTP.AUTH(mail_conn, username, passwd, schemes => 'PLAIN');

  utl_smtp.mail(mail_conn, msg_from);
  utl_smtp.rcpt(mail_conn, msg_to);

  UTL_smtp.open_data(mail_conn);

  UTL_SMTP.write_data(mail_conn, 'Date: ' || TO_CHAR(SYSDATE, 'DD-MON-YYYY
HH24:MI:SS') || UTL_TCP.crlf);
  UTL_SMTP.write_data(mail_conn, 'To: ' || msg_to || UTL_TCP.crlf);
  UTL_SMTP.write_data(mail_conn, 'From: ' || msg_from || UTL_TCP.crlf);
  UTL_SMTP.write_data(mail_conn, 'Subject: ' || msg_subject || UTL_TCP.crlf);
  UTL_SMTP.write_data(mail_conn, 'Reply-To: ' || msg_to || UTL_TCP.crlf ||
UTL_TCP.crlf);
  UTL_SMTP.write_data(mail_conn, msg_text || UTL_TCP.crlf || UTL_TCP.crlf);

  UTL_smtp.close_data(mail_conn);
  UTL_smtp.quit(mail_conn);

EXCEPTION
  WHEN UTL_smtp.transient_error OR UTL_smtp.permanent_error THEN
    UTL_smtp.quit(mail_conn);
    dbms_output.put_line(sqlerrm);
  WHEN OTHERS THEN
    UTL_smtp.quit(mail_conn);
    dbms_output.put_line(sqlerrm);
END;
/
```

Where:

- *mailhost*: specifies the SMTP Connection Endpoint from Step 1 in Send Email with Email Delivery Service on Autonomous Database.
- *username*: specifies the SMTP credential username from Step 2 in Send Email with Email Delivery Service on Autonomous Database.
- *passwd*: specifies the SMTP credential password from Step 2 in Send Email with Email Delivery Service on Autonomous Database.
- *msg_from*: specifies one of the approved senders from Step 3 in Send Email with Email Delivery Service on Autonomous Database.

Send Email with an Email Provider on a Private Endpoint

Describes the steps to send email with an email provider that is on Private Endpoint.

To send email from Autonomous Database using a email provider on a private endpoint, the email provider must be accessible from the Oracle Cloud Infrastructure VCN (the Autonomous Database instance's private endpoint). For example, you can access an email provider when:

- Both the source Autonomous Database instance and the email provider are in the same Oracle Cloud Infrastructure VCN.
- The source Autonomous Database instance and the email provider are in different Oracle Cloud Infrastructure VCNs that are paired.

- The email provider is on an on-premises network that is connected to the source Autonomous Database instance's Oracle Cloud Infrastructure VCN using FastConnect or VPN.

As a prerequisite, to send email using an email provider, define the following ingress and egress rules:

- Define an egress rule in the source database's subnet security list or network security group such that the traffic to the target host is allowed on port 587 or port 25 (depending on which port you are using).
- Define an ingress rule in the target host's subnet security list or network security group such that the traffic from the source Autonomous Database instance's IP address to port 587 or port 25 is allowed (depending on which port you are using).

To send email from an email provider on private endpoint:

1. Allow SMTP access for ADMIN user by appending an Access Control Entry (ACE).

   For example:

```
-- Create an Access Control List for the host
BEGIN
   DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
         host => 'www.example.com',
         lower_port => 587,
         upper_port => 587,
         ace => xs$ace_type(privilege_list => xs$name_list('SMTP'),
                            principal_name => 'ADMIN',
                            principal_type => xs_acl.ptype_db),
         private_target => TRUE);
END;
/
```

   > **Note:**
   >
   > `DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE` only supports a single hostname for the `host` parameter (on a private endpoint, using an IP address, a SCAN IP, or a SCAN hostname is not supported).

   If you set `ROUTE_OUTBOUND_CONNECTIONS` to `PRIVATE_ENDPOINT`, setting the `private_target` parameter to `TRUE` is not required in this API. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

2. Create a PL/SQL procedure to send email.
3. Send a test email using the PL/SQL procedure you created in step 2.

   For example:

```
execute send_mail('taylor@example.com', 'Email from Oracle Autonomous
Database', 'Sent using private email provider');
```

Use Credential Objects to set SMTP Authentication

Describes how to pass a credential objects to `UTL_SMTP.SET_CREDENTIAL` APIs.

The `SET_CREDENTIAL` subprogram sends the AUTH command to authenticate to the SMTP server.

The `UTL_SMTP.SET_CREDENTIAL` subprogram enables you to pass credential objects to set SMTP authentication. Credential objects are schema objects, hence they can be accessed only by privileged users and enable you to configure schema-level privileges to access control the credentials. Passing credential objects is a appropriate and secure way to store and manage username/password/keys for authentication.

The `UTL_SMTP.SET_CREDENTIAL` subprogram is a secure and convenient alternative to `UTL_SMTP.AUTH` subprogram.

Example

```
...
UTL_SMTP.AUTH (l_mail_conn, 'ocid1.user.oc1.username', 'xxxxxxxxxxxx',
schemes => 'PLAIN');
...
```

As shown in the example above, when you invoke `AUTH` subprogram, you must pass the username/password in clear text as part of PL/SQL formal parameters. You might need to embed the username/password into various PL/SQL automation or cron scripts. Passing clear text passwords is a compliance issue that is addressed in `UTL_SMTP.SET_CREDENTIAL` subprogram.

See AUTH Function and Procedure for more information.

`UTL_SMTP.SET_CREDENTIAL` Syntax

```
PROCEDURE UTL_SMTP.SET_CREDENTIAL (
    c          IN OUT NOCOPY connection,
    credential IN            VARCHAR2,
    schemes    IN            VARCHAR2 DEFAULT NON_CLEARTEXT_PASSWORD_SCHEMES
);

FUNCTION UTL_SMTP.SET_CREDENTIAL (
    c          IN OUT NOCOPY connection,
    credential IN            VARCHAR2,
    schemes    IN            VARCHAR2 DEFAULT NON_CLEARTEXT_PASSWORD_SCHEMES)
    RETURN reply;
```

Example

- Create a credential object:

  ```
  BEGIN DBMS_CLOUD.CREATE_CREDENTIAL (
      credential_name => 'HTTP_CRED',
      username        => 'web_app_user',
      password        => '<password>' );
  END;
  ```

  This creates a credential object which creates a stored username/password pair.

  See CREATE_CREDENTIAL Procedure for more information.

  See Specifying Scheduler Job Credentials for more information.

- Execute `UTL_SMTP.SET_CREDENTIAL` procedure:

```
DECLARE
      l_mail_conn UTL_SMTP.CONNECTION;
    BEGIN
      l_mail_conn := UTL_SMTP.OPEN_CONNECTION('smtp.example.com', 587);
      UTL_SMTP.SET_CREDENTIAL(l_mail_conn, 'SMTP_CRED', SCHEMES =>
'PLAIN');
      ...
END;
```

  This example sends the command to authenticate to the SMTP server. The Web server needs this information to authorize the request. The value `l_mail_conn` is the SMTP connection, `SMTP_CRED` is the credentials name and `PLAIN` is the SMTP authentication scheme.

See UTL_SMTP for more information.

See PL/SQL Package Notes for Autonomous Database for information on restrictions for `UTL_SMTP` on Autonomous Database.

Send Email from Autonomous Database Using DBMS_CLOUD_NOTIFICATION

Use the `DBMS_CLOUD_NOTIFICATION` package to send messages and query results as email.

- Send Messages as Email from Autonomous Database
- Send Query Results as Email from Autonomous Database

Send Messages as Email from Autonomous Database

You can use the `DBMS_CLOUD_NOTIFICATION` to send messages as an email.

1. Identify your SMTP connection endpoint for Email Delivery. You may need to subscribe to additional Oracle Cloud Infrastructure regions if Email Delivery is not available in your current region.

   For example, select one of the following for the SMTP connection endpoint:

   - smtp.us-phoenix-1.oraclecloud.com
   - smtp.us-ashburn-1.oraclecloud.com
   - smtp.email.uk-london-1.oci.oraclecloud.com
   - smtp.email.eu-frankfurt-1.oci.oraclecloud.com

   See Configure SMTP Connection for more information.

2. Generate SMTP credentials for Email Delivery. `UTL_SMTP` uses credentials to authenticate with Email Delivery servers when you send email.

   See Generate SMTP Credentials for a User for more information.

3. Create an approved sender for Email Delivery. Complete this step for all email addresses you use as the "`From`" with `UTL_SMTP.MAIL`.

   See Managing Approved Senders for more information.

4. Allow SMTP access for ADMIN user by appending an Access Control Entry (ACE).

For example:

```
BEGIN
  -- Allow SMTP access for user ADMIN
  DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
    host => 'www.us.example.com',
    lower_port => 587,
    upper_port => 587,
    ace => xs$ace_type(privilege_list => xs$name_list('SMTP'),
                       principal_name => 'ADMIN',
                       principal_type => xs_acl.ptype_db));
END;
/
```

5. Create a credential object and use `DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE` to send a message as an email.

```
BEGIN
   DBMS_CLOUD.CREATE_CREDENTIAL(
        credential_name => 'EMAIL_CRED',
        username        => 'username',
        password        => 'password'
   );
   END;
/
BEGIN
   DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE(
        provider        => 'email',
        credential_name => 'EMAIL_CRED',
        message         => 'Subject content',
        params          => json_object('recipient' value
'mark@example.com, suresh@example.com',
                                       'to_cc'  value 'nicole@example.com,
jordan@example.com',
                                       'to_bcc' value
'manisha@example.com',
                                       'subject' value 'Test subject',
                                       'smtp_host' value
'smtp.email.example.com',
                                       'sender'    value
'approver_sender@example.com' )
   );
   END;
/
```

Use the `params` parameter to specify the sender, smtp_host, subject, recipient, and recipients of a CC or BCC in `string` values.

- *sender*: specifies the Email ID of the approved sender from Step 3.
- *smtp_host*: specifies the SMTP host name from step 2.
- *subject*: specifies the subject of the email.
- *recipient*: specifies the email IDs of recipients. Use a comma between email IDs when there are multiple recipients.

- *to_cc*: specifies the email IDs that are receiving a CC of the email. Use a comma between email IDs when there are multiple CC recipients.
- *to_bcc*: specifies the email IDs that are receiving a BCC of the email. Use a comma between email IDs when there are multiple BCC recipients.

See SEND_MESSAGE Procedure for more information.

Send Query Results as Email from Autonomous Database

You can use the `DBMS_CLOUD_NOTIFICATION` package to send the results of a query as an email.

To use `DBMS_CLOUD_NOTIFICATION` to send mail:

1. Identify your SMTP connection endpoint for Email Delivery. You may need to subscribe to additional Oracle Cloud Infrastructure regions if Email Delivery is not available in your current region.

   For example, select one of the following for the SMTP connection endpoint:

   - smtp.us-phoenix-1.oraclecloud.com
   - smtp.us-ashburn-1.oraclecloud.com
   - smtp.email.uk-london-1.oci.oraclecloud.com
   - smtp.email.eu-frankfurt-1.oci.oraclecloud.com

   See Configure SMTP Connection for more information.

2. Generate SMTP credentials for Email Delivery. `UTL_SMTP` uses credentials to authenticate with Email Delivery servers when you send email.

   See Generate SMTP Credentials for a User for more information.

3. Create an approved sender for Email Delivery. Complete this step for all email addresses you use as the "`From`" with `UTL_SMTP.MAIL`.

   See Managing Approved Senders for more information.

4. Allow SMTP access for ADMIN user by appending an Access Control Entry (ACE).

   For example:

   ```
   BEGIN
     -- Allow SMTP access for user ADMIN
     DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
       host => 'www.us.example.com',
       lower_port => 587,
       upper_port => 587,
       ace => xs$ace_type(privilege_list => xs$name_list('SMTP'),
                          principal_name => 'ADMIN',
                          principal_type => xs_acl.ptype_db));
   END;
   /
   ```

5. Create a credential object and use `DBMS_CLOUD_NOTIFICATION.SEND_DATA` to send the output of a query as an email.

   ```
   BEGIN
      DBMS_CLOUD.CREATE_CREDENTIAL(
           credential_name => 'EMAIL_CRED',
   ```

```
        username        => 'username',
        password        => 'password'
   );
   END;
/
BEGIN
   DBMS_CLOUD_NOTIFICATION.SEND_DATA(
        provider        => 'email',
        credential_name => 'EMAIL_CRED',
        query           => 'SELECT tablespace_name FROM dba_tablespaces',
        params          => json_object('recipient' value
'mark@example.com, suresh@example.com',
                                       'to_cc'  value
'nicole@example.com1, jordan@example.com',
                                       'to_bcc' value
'manisha@example.com',
                                       'subject' value 'Test subject',
                                       'type' value 'json',
                                       'title' value 'mytitle',
                                       'message' value 'This is the
message',
                                       'smtp_host' value
'smtp.email.example.com',
                                       'sender'    value
'approver_sender@example.com' )
   );
   END;
/
```

Use the `params` parameter to specify the sender, smtp_host, subject, recipient, recipients of a CC or BCC, the message, data type, and the title in `String` values.

- *sender*: specifies the Email ID of the approved sender from Step 3.
- *smtp_host*: specifies the SMTP host name from step 2.
- *subject*: specifies the subject of the email. The maximum size is 100 characters.
- *recipient*: This specifies the email IDs of recipients. Use a comma between email IDs when there are multiple recipients.
- *to_cc*: specifies the email IDs that are receiving a CC of the email. Use a comma between email IDs when there are multiple CC recipients.
- *to_bcc*: specifies the email IDs that are receiving a BCC of the email. Use a comma between email IDs when there are multiple BCC recipients.
- *message*: specifies the message text.
- *type*: specifies the output format as either CSV or JSON.
- *title*: specifies the title of the attachment of SQL output. The title should only contain letters, digits, underscores, hyphens, or dots as characters in its value due to it being used to generate a file name.

The maximum message size for use with `DBMS_CLOUD_NOTIFICATION.SEND_DATA` for mail notification is 32k bytes.

See SEND_DATA Procedure for more information.

Send Slack Notifications from Autonomous Database

Describes how to configure Slack so that you can send messages, alerts, or output of a query from Autonomous Database to a Slack Channel. Also describes the procedures you use to send Slack notifications.

- Prepare to Send Slack Notifications from Autonomous Database
  To send Slack notifications, you must configure your Slack application to receive messages from Autonomous Database. Next, create a credential to use with the `DBMS_CLOUD_NOTIFICATION` procedures to send Slack notifications from Autonomous Database.
- Send Messages to a Slack Channel
- Send Query Results to a Slack Channel

Prepare to Send Slack Notifications from Autonomous Database

To send Slack notifications, you must configure your Slack application to receive messages from Autonomous Database. Next, create a credential to use with the `DBMS_CLOUD_NOTIFICATION` procedures to send Slack notifications from Autonomous Database.

To use Slack with `DBMS_CLOUD_NOTIFICATION` procedures:

1. Create your Slack app.

   The Slack app is installed in a Slack Workspace, which in turn has Channels where messages can be sent. The bot token of the Slack app must have the following permission scopes defined:

   ```
   channels:read
   chat:write
   files:write
   ```

   See Creating an app for information on setting up a Slack app.
2. Have your Slack admin add the Slack app to the channels to which `DBMS_CLOUD_NOTIFICATION` can send message through the "Integrations" option in the channel.

   See Basic app setup for more information.
3. Locate the app's bot token available in the Slack app specific page, under `https://app.slack.com`.

   See Basic app setup for more information.
4. Create a credential object to access the Slack app from Autonomous Database.

   > **Tip:**
   >
   > If you can not use the `CREATE_CREDENTIAL` procedure successfully, consult with the ADMIN user to grant execute access on `DBMS_CLOUD` packages.

   The credential's username is `SLACK_TOKEN` and the password is the bot token.

For example:

```
BEGIN
   DBMS_CLOUD.CREATE_CREDENTIAL(
     credential_name => 'SLACK_CRED',
     username    => 'SLACK_TOKEN',
     password    => 'xoxb-34....96-34....52-zW....cy');
END;
/
```

See CREATE_CREDENTIAL Procedure for more information.

5. Configure access control to allow user access to external network services (Slack).

```
BEGIN
 DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE (
  host         => 'slack.com',
  lower_port   => 443,
  upper_port   => 443,
  ace          => xs$ace_type(
      privilege_list => xs$name_list('http'),
      principal_name => example_invoking_user,
      principal_type => xs_acl.ptype_db));
END;
```

See Configuring Access Control for External Network Services for more information.

6. If you are sending SQL, configure access control to allow sending of output with `DBMS_CLOUD_NOTIFICATION.SEND_DATA`.

```
BEGIN
 DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE (
  host         => 'files.slack.com',
  lower_port   => 443,
  upper_port   => 443,
  ace          => xs$ace_type(
      privilege_list => xs$name_list('http'),
      principal_name => example_invoking_user,
      principal_type => xs_acl.ptype_db));
END;
```

See Configuring Access Control for External Network Services for more information.

Send Messages to a Slack Channel

After creating the Slack credential object as described in Prepare to Send Slack Notifications from Autonomous Database, you can use the `DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE` procedure to send a message to a Slack channel.

Example:

```
BEGIN
   DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE(
      provider          => 'slack',
      credential_name   => 'SLACK_CRED',
```

```
      message           => 'Alert from Autonomous Database...',
      params            => json_object('channel' value 'C0....08'));
END;
/
```

Use the `params` parameter to specify the Slack channel.

- `channel`: Specifies the Channel ID.

  The Channel ID is a unique ID for a channel and is different from the channel name. In Slack, when you view channel details you can find the Channel ID on the “About” tab. See How to Find your Slack Team ID and Slack Channel ID for more information.

See SEND_MESSAGE Procedure for more information.

Send Query Results to a Slack Channel

After creating the Slack credential object as described in Prepare to Send Slack Notifications from Autonomous Database, you can use the `DBMS_CLOUD_NOTIFICATION.SEND_DATA` procedure to send the output of a query to a Slack channel.

Example:

```
BEGIN
  DBMS_CLOUD_NOTIFICATION.SEND_DATA(
      provider => 'slack',
      credential_name => 'SLACK_CRED',
      query => 'SELECT username, account_status, expiry_date FROM USER_USERS
WHERE rownum < 5',
      params => json_object('channel' value 'C0....08',
                            'type' value 'csv'));
END;
/
```

Use the `params` parameter to specify the Slack channel and the data type:

- `channel`: Specifies the Channel ID.

  The Channel ID is a unique ID for a channel and is different from the channel name. In Slack, when you view channel details you can find the Channel ID on the “About” tab. See How to Find your Slack Team ID and Slack Channel ID for more information.

- `type`: Specifies the output type. Valid values are: '`csv`' or '`json`'.

See SEND_DATA Procedure for more information.

Related Topics

- SEND_DATA Procedure
  The `SEND_DATA` procedure sends the results of the specified query to a provider.

Send Messages to OCI Notifications Service from Autonomous Database

Describes how to send messages from Autonomous Database to an OCI Notifications service topic.

The subscriptions you configure with the OCI Notifications service determine endpoints for a topic. Published messages are sent to each subscription for a topic. There are a number of

supported subscription protocols, including: email, function, and Slack. See Overview of Notifications for more information.

When you use the OCI Notifications service with `DBMS_CLOUD_NOTIFICATION` package, only the `DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE` procedure is supported (the `DBMS_CLOUD_NOTIFICATION.SEND_DATA` procedure is not supported).

- Prepare to Send OCI Notifications from Autonomous Database
  To send notifications to topics in OCI Notifications service, first create topics and subscriptions in OCI Notifications.
- Send Messages to an OCI Notifications Topic
  Describes the steps to send messages to an OCI Notifications service topic.
- Limitations for Notifications with OCI Notifications Service
  There are limitations when you use `DBMS_CLOUD_NOTIFICATION` to send messages and you use the `oci` provider (OCI Notifications service).

Prepare to Send OCI Notifications from Autonomous Database

To send notifications to topics in OCI Notifications service, first create topics and subscriptions in OCI Notifications.

As a prerequisite, before you send messages with OCI Notifications, do the following:

1. Create an OCI Notifications topic.

   See Creating a Topic for details on creating a topic.
2. Create OCI Notifications subscriptions.

   See Creating a Subscription for details on creating a subscription.

   The subscriptions you configure determine endpoints for a topic. Published messages are sent to each subscription for a topic. There are a number of supported subscription protocols, including: email, function, and Slack.

   See Overview of Notifications for more information.
3. After you create the OCI Notifications subscriptions, confirm the subscriptions.

   You can only publish messages to confirmed subscriptions. See Confirming a Subscription for details on confirming subscriptions.

Send Messages to an OCI Notifications Topic

Describes the steps to send messages to an OCI Notifications service topic.

1. Create a topic and create subscriptions.

   See Prepare to Send OCI Notifications from Autonomous Database for details.
2. Create a credential object to access OCI Notifications service from Autonomous Database.

   You can use either Oracle Cloud Infrastructure signing key credentials or resource principal credentials.

   - Create Oracle Cloud Infrastructure signing key credentials. For example:

```
BEGIN
   DBMS_CLOUD.CREATE_CREDENTIAL (
       credential_name => 'OCI_KEY_CRED',
       user_ocid       =>
'ocid1.user.oc1..aaaaaaaauq54mi7zdyfhw33ozkwuontjceel7fok5nq3bf2vwetkpqs
```

```
oa',
       tenancy_ocid    =>
'ocid1.tenancy.oc1..aabbbbbbaafcue47pqmrf4vigneebgbcmmoy5r7xvoypicjqqge3
2ewnrcyx2a',
       private_key     =>
'MIIEogIBAAKCAQEAtUnxbmrekwgVac6FdWeRzoXvIpA9+0r1.....wtnNpESQQQ0QLGPD8N
M//JEBg=',
       fingerprint     =>
'f2:db:f9:18:a4:aa:fc:94:f4:f6:6c:39:96:16:aa:27');
END;
/
```

See CREATE_CREDENTIAL Procedure for more information.

- Create resource principal credentials and policies.

  See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

3. Send a message.

   - Send a message using Oracle Cloud Infrastructure signing key credentials:

     ```
     BEGIN
         DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE(
             provider         => 'oci',
             credential_name  => 'OCI_KEY_CRED',
             message          => 'Text message',
             params           => json_object('topic_ocid' value
     'oci********pa',
                                             'title'value 'Title for
     message'));
     END;
     /
     ```

   - Send a message with resource principal credentials:

     ```
     BEGIN
         DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE(
             provider         => 'oci',
             credential_name  => 'OCI$RESOURCE_PRINCIPAL',
             message          => 'Text message',
             params           => json_object('topic_ocid' value
     'oci********pa',
                                             'title'      value 'Title for
     message'));
     END;
     /
     ```

   With the `oci` provider there are different possible endpoints for a message, based on the subscriptions. The `title` parameter is treated differently for different subscription protocols:

   - email: The `title` specifies the subject line of the message.
   - Slack: The `title` value is not used. If it is included the value is ignored.

See SEND_MESSAGE Procedure for more information.

Limitations for Notifications with OCI Notifications Service

There are limitations when you use `DBMS_CLOUD_NOTIFICATION` to send messages and you use the `oci` provider (OCI Notifications service).

Note the following limitations:

- Message Size Limit: The message size limit per request is 64KB. Messages exceeding this size cannot be sent or processed within a single request.
- Message Delivery Rate Limit: There are different message delivery rate limits depending on the protocol used.
  - `email` protocol: the limit is 10 messages per minute per endpoint.

  This restriction may impact the speed at which messages can be transmitted and received.
- Transactions Per Minute (TPM) per-tenancy Limit: There is a TPM limit of 60 per topic for this operation. This limit represents the maximum number of messages per minute that can be processed or handled for a given topic. Exceeding this limit may result in messages not being processed or processed at a slower rate.
- The the Autonomous Database instance, the credential, if you are using user credentials, and the subscription topic must all be located in the same region.

See Overview of Notifications for more information on OCI Notifications.

Send Microsoft Teams Notifications from Autonomous Database

Describes how to configure Microsoft Teams so that you can send messages, alerts, or output of a query from Autonomous Database to a Microsoft Teams Channel. Also describes the procedures you use to send Microsoft Teams notifications.

- Prepare to Send Microsoft Teams Notifications from Autonomous Database
  Get started by configuring a bot in your Microsoft Teams app. Next, create a credential to use with the `DBMS_CLOUD_NOTIFICATION` procedures to send Microsoft Teams notifications from Autonomous Database.
- Send Messages to a Microsoft Teams Channel
- Send Query Results to a Microsoft Teams Channel

Prepare to Send Microsoft Teams Notifications from Autonomous Database

Get started by configuring a bot in your Microsoft Teams app. Next, create a credential to use with the `DBMS_CLOUD_NOTIFICATION` procedures to send Microsoft Teams notifications from Autonomous Database.

To configure Microsoft Teams notifications:

1. Create your Microsoft Teams app and add a bot in it. See Developer Portal for Teams for information on setting up an app.
2. From the **Bot Management** section, ensure that the bot has a secret key, scope set to *Team* and permission to send notifications.
3. Publish the app to your org to make it available to the people in your org.
4. After your IT admin approves the app from the admin center, install the app from the **Apps** section in Teams.

5. Request the `Files.ReadWrite.All` and `ChannelSettings.Read.All` permissions to the app for Graph API from the Azure Portal using the following instructions:
   a. Log into your Azure Portal, navigate to **Azure Active Directory** using the left panel, and select the **Application** option.
   b. The Application page displays the apps that you own along with the bots added to those apps. Click the bot to view its details.
   c. Copy the *directory/tenant id* from the bot overview page for later use.
   d. Then, go to **API permissions** in the left panel. Under **API permissions**, click **Add permission**, select **Microsoft graph** and then **Application permission**.
   e. Search for `Files.ReadWrite.All` and `ChannelSettings.Read.All` permissions and add them.
6. Have your IT admin approve the above-requested permissions from the Azure portal by following the steps below:
   a. Log into your Azure Portal, navigate to **Azure Active Directory** using the left panel, and select the **Application** option.
   b. Select **All applications** from the **Application** page.
   c. Search the *application/bot* by its name, go to **API permissions**, and grant the ADMIN consent for the requested permissions: `Files.ReadWrite.All` and `ChannelSettings.Read.All`.

> **Tip:**
>
> After your app is approved by the IT Admin, you can provide the bot id and secret key to other users to install the app within Teams in the org.

7. After the app is approved by the IT admin and the above requested permissions are granted, you can use the app's bot ID and secret key are used to create the credential object and generate a bot token.
8. To send a query result to a Microsoft Teams channel, obtain the `team id` and the `tenant id`.

> **Tip:**
>
> The `team id` is located in the team link between /team/ and /conversations. The `tenant id` is located after "tenantId=" at the end of the team link. This link is found by clicking the three dots next to the team name and selecting **Get link to team**.

   For example:

   `https://teams.microsoft.com/l/team/`*`teamID`*`/conversations?groupId=`*`groupid`*`%tenantId=`*`tenantid`*

9. Obtain the `channelID`.

> **Tip:**
>
> `channelID` is located in the channel link between /team/ and the channel name. This link is found by clicking the three dots next to the channel name and selecting **Get link to channel**.

For example:

```
https://teams.microsoft.com/l/channel/channelID/channel_name?
groupId=groupid&tenantId=tenantid
```

10. Create a credential object to access the Microsoft Teams app from Autonomous Database.

> **Tip:**
>
> If you can not use the `CREATE_CREDENTIAL` procedure successfully, consult with the ADMIN user to grant execute access on `DBMS_CLOUD` packages.

The credential's username is the bot_id and the password is the bot key.

For example:

```
BEGIN
    DBMS_CLOUD.CREATE_CREDENTIAL(credential_name => 'TEAMS_CRED',
        username        => 'bot_id',
        password        => 'bot_secret');
END;
/
```

See CREATE_CREDENTIAL Procedure for more information.

Send Messages to a Microsoft Teams Channel

After creating the Microsoft Teams credential object as described in Prepare to Send Microsoft Teams Notifications from Autonomous Database, you can use the `DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE` procedure to send a message to a Microsoft Teams channel.

Example:

```
BEGIN
     DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE(
        provider        => 'msteams',
        credential_name => 'TEAMS_CRED',
        message         => 'text from new teams api',
        params          => json_object('channel' value 'channelID'));
END;
/
```

Use the `params` parameter to specify the channel.

- *channel*: specifies the Channel ID obtained from Step 10 in Prepare to Send Microsoft Teams Notifications from Autonomous Database.

See SEND_MESSAGE Procedure for more information.

Send Query Results to a Microsoft Teams Channel

After creating the Microsoft Teams credential object as described in Prepare to Send Microsoft Teams Notifications from Autonomous Database, you can use the `DBMS_CLOUD_NOTIFICATION.SEND_DATA` procedure to send the output of a query to a Microsoft Teams channel.
Example:

```
BEGIN
     DBMS_CLOUD_NOTIFICATION.SEND_DATA(provider => 'msteams',
        credential_name => 'TEAMS_CRED',
        query           => 'SELECT tablespace_name FROM dba_tablespaces',
        params          => json_object('tenant'value '5b743bc******c0286',
                                       'team'value '0ae401*********5d2bd',
                                       'channel'value
'19%3a94be023*****%40thread.tacv2',
                                       'title'value 'today',
                                       'type'value 'csv'));
END;
/
```

Use the `params` parameter to specify the tenant, team, channel, title, and data type in `string` values.

- *tenant*: specifies the tenant ID obtained from Step 8 in Prepare to Send Microsoft Teams Notifications from Autonomous Database.
- *team*: specifies the team ID obtained from Step 8 in Prepare to Send Microsoft Teams Notifications from Autonomous Database.
- *channel*: specifies the channel ID obtained from Step 9 in Prepare to Send Microsoft Teams Notifications from Autonomous Database.
- *title*: specifies the title of the file. The title can only contain alphabets, digits, underscores, and hyphens. The file name that appears in Microsoft Teams will be a concatenation of the title parameter and the timestamp to ensure uniqueness. The maximum title size is 50 characters.

  For example: '*title*'_'*timestamp*'.'*format*'
- *type*: This specifies the output format. Valid values are CSV or JSON.

> **Note:**
>
> The maximum file size supported when using `DBMS_CLOUD_NOTIFICATION.SEND_DATA` for Microsoft Teams is 4MB.

See SEND_DATA Procedure for more information.

User Defined Notification Handler for Scheduler Jobs

Database Scheduler provides an email notification mechanism to track the status of periodically running or automated jobs. In addition to this, the Database Scheduler also supports user-defined PL/SQL Scheduler job notification handler procedure.

Adding a scheduler job notification handler procedure allows you to monitor scheduled or automated jobs running in your Autonomous Database.

- About User Defined Notification Handler for Scheduler Jobs
  The Database Scheduler supports job notification handler procedure that can make use of custom code to call HTTP or REST endpoints for improved monitoring of scheduler jobs in an Autonomous Database instance.
- Create a Job Notification Handler Procedure
  Provides steps to create a job notification handler.
- Register the Job Handler Notification Procedure
  Use `DBMS_SCHEDULER.SET_SCHEDULER_ATTRIBUTE` procedure to set `JOB_NOTIFICATION_HANDLER` attribute value to register the job handler notification procedure.
- Trigger the Job Handler Notification Procedure
  You must call the `DBMS_SCHEDULER.ADD_JOB_EMAIL_NOTIFICATION` procedure to trigger the user defined job notification handler procedure.
- De-Register the Job Handler Notification Procedure
  Use `DBMS_SCHEDULER.SET_SCHEDULER_ATTRIBUTE` to de-register the job handler notification procedure.

About User Defined Notification Handler for Scheduler Jobs

The Database Scheduler supports job notification handler procedure that can make use of custom code to call HTTP or REST endpoints for improved monitoring of scheduler jobs in an Autonomous Database instance.

The handler procedure receives all pertinent information regarding the job, such as the job owner's name, class name, event type, and timestamp in JSON format. Based on the information, the handler procedure then takes the required action.

See DBMS_SCHEDULER for more information on Oracle Scheduler.

Configuring user defined notification handler for scheduler jobs consists of these steps:

- Create a job notification handler procedure as described in: Create a Job Notification Handler Procedure.
- Register the job notification handler procedure for the database as described in: Register the Job Handler Notification Procedure.
- Trigger the job notification handler procedure as described in: Trigger the Job Handler Notification Procedure
- De-Register the job notification handler procedure for the database as described in: De-Register the Job Handler Notification Procedure.

Create a Job Notification Handler Procedure

Provides steps to create a job notification handler.

1. Create a credential object.

   See CREATE_CREDENTIAL Procedure for more information.

   See Specifying Scheduler Job Credentials for more information.

2. Create a Job notification handler:

   Example to create a job notification handler procedure to send a message to a Slack channel:

```
CREATE OR REPLACE PROCEDURE ADMIN.SEND_NOTIFICATION(data_in CLOB) AS
BEGIN
  DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE (
    provider          => 'slack',
    credential_name   => 'SLACK_CRED',
    message           => data_in,
    params            => JSON_OBJECT('channel' value 'adw-job-nfy');
end;
/
```

   This example creates the `SEND_NOTIFICATION` procedure.

   This user defined procedure sends the provided input data as a message to the specified Slack channel.

   See Send Slack Notifications from Autonomous Database for more information.

   See SEND_MESSAGE Procedure for more information.

   Example to create a job notification handler procedure to insert the message into a table:

```
CREATE TABLE ADMIN.JOB_STATUS(jnfy_data CLOB);

CREATE OR REPLACE PROCEDURE ADMIN.INSERT_JOB_STATUS(data_in CLOB) AS
  l_sessuser VARCHAR2(128) := SYS_CONTEXT('userenv','session_user');
BEGIN
  INSERT INTO ADMIN.JOB_STATUS (jnfy_data) VALUES (data_in || TO_CLOB(' :
Sent By Session User : ' || l_sessuser));
  COMMIT;
END;
/
```

   This example creates the `JOB_STATUS` table and `INSERT_JOB_STATUS` procedure to insert the session-specific values into the table.

   You must be logged in as the ADMIN user or have `CREATE ANY PROCEDURE` system privilege to create a job notification handler procedure.

> **Note:**
>
> An `ORA-27405` is returned when you specify an invalid owner or object name as the job notification handler procedure.

Existing `DBMS_SCHEDULER` procedures `ADD_JOB_EMAIL_NOTIFICATION` and `REMOVE_JOB_EMAIL_NOTIFICATION` are enhanced to support the job notification handler procedure.

See ADD_JOB_EMAIL_NOTIFICATION Procedure and REMOVE_JOB_EMAIL_NOTIFICATION Procedure for more information.

Use DBA_SCHEDULER_NOTIFICATIONS dictionary view to query the list of notifications for a scheduler job. See DBA_SCHEDULER_NOTIFICATIONS for more information.

Register the Job Handler Notification Procedure

Use DBMS_SCHEDULER.SET_SCHEDULER_ATTRIBUTE procedure to set JOB_NOTIFICATION_HANDLER attribute value to register the job handler notification procedure.

The JOB_NOTIFICATION_HANDLER attribute specifies the job handler notification procedure that you want to use.

To register the job handler notification procedure you must:

- Be logged in as the ADMIN user or have MANAGE SCHEDULER privilege.
- Have EXECUTE privilege on the handler procedure or EXECUTE ANY PROCEDURE system privilege.

The JOB_NOTIFICATION_HANDLER attribute and EMAIL_SERVER attribute are mutually exclusive. The ATTRIBUTE parameter of the SET_SCHEDULER_ATTRIBUTE procedure can have either the JOB_NOTIFICATION_HANDLER or the EMAIL_SERVER value at a time. You are allowed to either configure email notifications or create your notification handler for your scheduler jobs.

An ORA-27488 error is raised when you attempt to set both EMAIL_SERVER and JOB_NOTIFICATION_HANDLER global attributes.

Execute DBMS_SCHEDULER.SET_SCHEDULER_ATTRIBUTE procedure to register the job handler notification procedure:

```
BEGIN

DBMS_SCHEDULER.SET_SCHEDULER_ATTRIBUTE('job_notification_handler','ADMIN.SEND_
NOTIFICATION');
END;
/
```

This example registers the ADMIN.SEND_NOTIFICATION procedure as the job handler notification procedure for your database.

See SET_SCHEDULER_ATTRIBUTE Procedure for more information.

Execute this command to verify the job notification handler:

```
SELECT value FROM dba_scheduler_global_attribute WHERE
attribute_name='JOB_NOTIFICATION_HANDLER';

VALUE
---------------
"ADMIN"."SEND_NOTIFICATION"
```

See DBA_SCHEDULER_GLOBAL_ATTRIBUTE for more information.

You must assign `EXECUTE` privilege to allow other users to use the job notification handler. For example:

```
GRANT EXECUTE ON ADMIN.SEND_NOTIFICATION To DWUSER;
```

`ORA-27476 ("\"%s\".\"%s\" does not exist")` or `ORA-27486 ("insufficient privileges")` error is thrown if you do not have privilege on the job handler notification procedure.

Trigger the Job Handler Notification Procedure

You must call the `DBMS_SCHEDULER.ADD_JOB_EMAIL_NOTIFICATION` procedure to trigger the user defined job notification handler procedure.

The overloaded form of the `DBMS_SCHEDULER.ADD_JOB_EMAIL_NOTIFICATION` enables you to trigger the job notification handler procedure and send a notification. However, these notifications are not sent out in the form of an email when you have registered the job notification handler procedure. Hence, the parameters `SUBJECT`, and `BODY` are optional. The `RECIPIENT` parameter is still mandatory. Since this overloaded form of procedure is not sending email notifications so, you can provide any string value for the `RECIPIENT` parameter.

For example, the following steps create a scheduler job, add notification for the job, enable the job, verify the notification entries, show the data received by the job notification handler procedure, and remove notification for the job.

1. Create a scheduler job in the `DWUSER` schema:

```
BEGIN
 DBMS_SCHEDULER.CREATE_JOB(
   job_name   => 'DWUSER.MY_JOB',
   job_type   => 'PLSQL_BLOCK',
   job_action => 'BEGIN null; END;',
   enabled    => FALSE,
   auto_drop  => FALSE);
END;
/
```

   This creates a job `DWUSER.MY_JOB` with the specified attributes.

   See CREATE_JOB Procedure for more information.

2. Configure the job to send notifications at specified events.

```
BEGIN
  DBMS_SCHEDULER.ADD_JOB_EMAIL_NOTIFICATION(
    job_name   => 'DWUSER.MY_JOB',
    recipients => 'PLACEHOLDER_STRING',
    subject    => 'Job Notification-%job_owner%.%job_name%-%event_type%',
    body       => '%event_type% occurred at %event_timestamp%.
%error_message%',
    events     => 'job_started, job_succeeded, job_completed');
END;
/
```

This procedure adds notifications for the DWUSER.MY_JOB job. The notifications are sent whenever any of the specified job state events is raised.

See ADD_JOB_EMAIL_NOTIFICATION Procedure for more information.

3. Enable the scheduler job.

```
EXEC DBMS_SCHEDULER.ENABLE('DWUSER.MY_JOB');
```

See ENABLE Procedure for more information.

When you enable the DWUSER.MY_JOB job, the USER_SCHEDULER_NOTIFICATIONS view is populated with the job notification entries. To verify you can query the USER_SCHEDULER_NOTIFICATIONS view. For example:

```
SELECT  job_name, recipient, event, subject, body
FROM user_scheduler_notifications
ORDER BY notification_owner, owner, job_name;

JOB_NAME    RECIPIENT            EVENT          SUBJECT         BODY
-------    ------------------   -------------  ---------------
---------------
MY_JOB     placeholder_string   JOB_STARTED    Job Notificatio
%event_type% oc
                                               n-%job_owner%.% curred at
%even
                                               job_name%-%even
t_timestamp%. %
                                               t_type%
error_message%

MY_JOB     placeholder_string  JOB_SUCCEEDED    Job Notificatio
%event_type% oc
                                                n-%job_owner%.% curred at
%even
                                                job_name%-%even
t_timestamp%. %
                                                t_type%
error_message%

MY_JOB    placeholder_string  JOB_COMPLETED    Job Notificatio
%event_type% oc
                                               n-%job_owner%.% curred at
%even
                                               job_name%-%even
t_timestamp%. %
                                               t_type%
error_message%
```

See USER_SCHEDULER_NOTIFICATIONS View for more information.

When the job DWUSER.MY_JOB is executed and any of the specified job state events is raised, the job notification handler procedure is triggered and receives the specified

information as input. For example, the `ADMIN.SEND_NOTIFICATION` job notification handler procedure receives the following:

```
{"job_owner":"DWUSER","job_name":"MY_JOB","job_class_name":"DEFAULT_JOB_CLA
SS","event_type":"JOB_STARTED","event_timestamp":"12-JAN-23 08.13.46.193306
 PM UTC","error_code":0,"error_msg":null,"sender":null,"recipient":"data_lo
ad_pipeline","subject":"Job Notification-DWUSER.MY_JOB-JOB_STARTED","msg_te
xt":"JOB_STARTED occurred at 12-JAN-23 08.13.46.193306 PM UTC. ","comments"
:"User defined job notification handler"}

{"job_owner":"DWUSER","job_name":"MY_JOB","job_class_name":"DEFAULT_JOB_CLA
SS","event_type":"JOB_SUCCEEDED","event_timestamp":"12-JAN-23 08.13.46.2863
44 PM UTC","error_code":0,"error_msg":null,"sender":null,"recipient":"data_
load_pipeline","subject":"Job Notification-DWUSER.MY_JOB-JOB_SUCCEEDED","ms
g_text":"JOB_SUCCEEDED occurred at 12-JAN-23 08.13.46.286344 PM UTC. ","com
ments":"User defined job notification handler"}
```

4. Remove the job notification. For example:

```
EXEC DBMS_SCHEDULER.REMOVE_JOB_EMAIL_NOTIFICATION('DWUSER.MY_JOB');
```

See REMOVE_JOB_EMAIL_NOTIFICATION Procedure for more information.

De-Register the Job Handler Notification Procedure

Use `DBMS_SCHEDULER.SET_SCHEDULER_ATTRIBUTE` to de-register the job handler notification procedure.

To de-register the job handler notification procedure you must be logged in as the ADMIN user or have `MANAGE SCHEDULER` privilege.

Example to de-register the job handler notification procedure:

```
BEGIN
  DBMS_SCHEDULER.SET_SCHEDULER_ATTRIBUTE ('job_notification_handler','');
END;
/
```

Oracle Extensions for IDEs

Oracle extensions let developers connect to, browse, and manage Autonomous Databases directly from common IDEs.

- Use Oracle Cloud Infrastructure Toolkit for Eclipse
  Oracle Cloud Infrastructure Toolkit for Eclipse is a plugin that enables Java developers to easily connect to Oracle Autonomous Database through their IDE. The plugin is free and is available for Linux, UNIX, Microsoft Windows, and Apple Mac OS.
- Use Oracle Developer Tools for Visual Studio
  Oracle Developer Tools for Visual Studio is a tightly integrated extension for Microsoft Visual Studio and Oracle Autonomous Database. The extension is free and supports Visual Studio 2019 and Visual Studio 2017 on Microsoft Windows.

- Use Oracle Developer Tools for VS Code
  Oracle Developer Tools for VS Code is a tightly integrated extension for Microsoft Visual Studio Code (VS Code) and Oracle Autonomous Database. The extension is free and is available for Linux, Microsoft Windows, and Apple Mac OS.

Use Oracle Cloud Infrastructure Toolkit for Eclipse

Oracle Cloud Infrastructure Toolkit for Eclipse is a plugin that enables Java developers to easily connect to Oracle Autonomous Database through their IDE. The plugin is free and is available for Linux, UNIX, Microsoft Windows, and Apple Mac OS.

You can use the plugin to perform cloud and database operations right from Eclipse, such as creating Autonomous Databases, stopping and starting, scaling up and down, and so on. You can also use the plugin to easily connect to the databases to browse the schema, access tables, execute SQL statements, and perform other development tasks.

Users with permissions to manage the databases can perform a number of actions, including those listed below. For detailed information about permissions, see Toolkit for Eclipse in the Oracle Cloud Infrastructure documentation. You can:

- Create Autonomous Databases
- Start, stop, terminate, clone, and restore Autonomous Databases
- Scale up and down
- Download the client credentials zip file (database wallet)
- Connect to Autonomous Databases
- Browse the schema
- Choose compartments and regions
- Change the administrator password, and so on

Download the latest version of the plugin from GitHub (`com.oracle.oci.eclipse-`*`version`*`.zip`, where *`version`* is the latest version, for instance `1.2.0`):

`https://github.com/oracle/oci-toolkit-eclipse/releases`

Then follow the installation instructions and details about how to get started in this step-by-step walkthrough:

New Eclipse Plugin for Accessing Autonomous Database (ATP/ADW)

Use Oracle Developer Tools for Visual Studio

Oracle Developer Tools for Visual Studio is a tightly integrated extension for Microsoft Visual Studio and Oracle Autonomous Database. The extension is free and supports Visual Studio 2019 and Visual Studio 2017 on Microsoft Windows.

You can use the extension to perform database management operations right from Visual Studio, such as creating Autonomous Databases, stopping and starting, scaling up and down, and so on. You can also use the extension to easily connect to the databases and perform development tasks, such as browsing your Oracle schema and launching integrated Oracle designers and wizards to create and alter schema objects.

Users with permissions to manage the databases can perform a number of actions, including the following:

- Sign up for Oracle Cloud

- Connect to a cloud account using a simple auto-generated config file and key file
- Create new or clone existing Always Free Autonomous Database, Autonomous Database Dedicated, and Autonomous Database Serverless databases
- Automatically download credentials files (including wallets) and quickly connect, browse, and operate on Autonomous Database schemas
- Change compartments and regions without reconnecting
- Start, stop, or terminate Autonomous Database
- Scale up/down Autonomous Database resources
- Restore from backup
- Update instance credentials
- Rotate wallets
- Convert Always Free Autonomous Database to paid:

> **Note:**
>
> Promotion of Always Free to a paid Autonomous Database is supported only if the database version for the Always Free Autonomous Database is Oracle Database 19c or Oracle Database 23ai.

Download the extension from Visual Studio Marketplace:

- Oracle Developer Tools for Visual Studio 2019
- Oracle Developer Tools for Visual Studio 2017

You'll find lots of information about the extension on those Marketplace pages.

Then follow the installation instructions and details about how to get started in this step-by-step walkthrough:

New Release: Visual Studio Integration with Oracle Autonomous Database

For detailed information about how to use the extension, see the online documentation that's optionally installed with Oracle Developer Tools for Visual Studio. Press the F1 key to display the context-sensitive help for each dialog.

Use Oracle Developer Tools for VS Code

Oracle Developer Tools for VS Code is a tightly integrated extension for Microsoft Visual Studio Code (VS Code) and Oracle Autonomous Database. The extension is free and is available for Linux, Microsoft Windows, and Apple Mac OS.

You can use the extension to connect to Autonomous Databases right from Visual Studio Code and easily explore database schema, view table data, and edit and execute SQL and PL/SQL.

Download the extension from Visual Studio Marketplace:

Oracle Developer Tools for VS Code

Installation instructions and information about how to get started can be found in this quick start guide:

Getting Started Using Oracle Developer Tools for VS Code

Use Oracle Java on Autonomous Database

Autonomous Database supports Oracle JVM. Oracle JVM is a standard, Java-compatible environment that runs any pure Java application.

Oracle JVM is compatible with the standard JLS and the JVM specifications. It supports the standard Java binary format and the standard Java APIs. In addition, Oracle Database adheres to standard Java language semantics, including dynamic class loading at run time.

See About Using Java in Oracle Database for information on Oracle Java.

- Enable Oracle Java
  Use `DBMS_CLOUD_ADMIN.ENABLE_FEATURE` to enable Oracle Java on Autonomous Database.
- Check Oracle Java Version
  You can check the Oracle Java version and the component registry for information on Oracle Java in the Autonomous Database instance.
- Load Java classes and JAR Files into Autonomous Database
  You can use the client-side `loadjava` option to load Java classes and JAR files into Oracle JVM on an Autonomous Database instance.
- Notes for Oracle Java on Autonomous Database
  Provides notes for using Oracle Java on Autonomous Database.

Enable Oracle Java

Use `DBMS_CLOUD_ADMIN.ENABLE_FEATURE` to enable Oracle Java on Autonomous Database.

1. Run `DBMS_CLOUD_ADMIN.ENABLE_FEATURE`.

   ```
   BEGIN
      DBMS_CLOUD_ADMIN.ENABLE_FEATURE(
          feature_name => 'JAVAVM' );
   END;
   /
   ```

   This initiates the request to install Oracle Java on the Autonomous Database instance.

   See ENABLE_FEATURE Procedure for more information.

2. Restart the Autonomous Database instance.

   See Restart Autonomous Database for more information.

After you restart the Autonomous Database instance, the Java installation proceeds for some time after the database starts. See Check Oracle Java Version to verify that Oracle JVM is enabled and the installation is complete.

Check Oracle Java Version

You can check the Oracle Java version and the component registry for information on Oracle Java in the Autonomous Database instance.

1. Check the component registry status and version for Oracle Java.

   ```
   SELECT status, version FROM DBA_REGISTRY
      WHERE comp_id = 'JAVAVM';
   ```

```
STATUS VERSION
------ ----------
VALID  19.0.0.0.0
```

If Oracle Java is not installed, this query shows no rows.

If Oracle Java is loading this query shows status `LOADING`. For example:

```
STATUS  VERSION
------- -------
LOADING null
```

2. Check the Oracle Java JDK version.

```
SELECT dbms_java.get_jdk_version FROM DUAL;

GET_JDK_VERSION
---------------
1.8.0_411
```

If Oracle Java is not installed, this query shows an error message.

Load Java classes and JAR Files into Autonomous Database

You can use the client-side `loadjava` option to load Java classes and JAR files into Oracle JVM on an Autonomous Database instance.

The client-side `loadjava` option is supported as follows:

1. The JDK opens the JAR file.
2. The client-side opens a communication channel with the Autonomous Database.
3. The Java code is submitted to the Oracle JVM running on the Autonomous Database for loading.

This process is similar to creating a dynamic Java stored procedure from SQL code, where no file access is involved, but the code gets created.

Note the following:

- The Java code you load cannot invoke any operating system or network calls.
- Use of server-side `loadjava`, that is calls to the procedure `dbms_java.loadjava` is not supported. The procedure `dbms_java.loadjava` accesses the file system of the database server where the Oracle JVM runs. This is not allowed on Autonomous Database and calls to the procedure `dbms_java.loadjava` from within your Java application are not supported.

Notes for Oracle Java on Autonomous Database

Provides notes for using Oracle Java on Autonomous Database.

- You cannot disable Oracle Java after it is enabled on the Autonomous Database instance.
- Autonomous Database performs Oracle Java patching, as required, during the regular Autonomous Database maintenance window.

During Oracle Java patching, Java is not available and users could get an error similar to the following:

```
ERROR at line 1:
ORA-29548: Java system class reported: release of Java system classes in
the
database (19.0.0.0.220118 1.8) does not match that of the oracle executable
(19.0.0.0.220419 1.8).
```

During the maintenance window, when the Java patching phase is active there is no response for Java session calls or you see the `ORA-29548` error. After the maintenance window completes, Java usage is restored.

You can use the events `NewMaintenanceSchedule` and `ScheduledMaintenanceWarning` to be notified of Oracle Java patching. See Information Events on Autonomous Database for more information.

See About Scheduled Maintenance and Patching for more information.

Test Workloads with Oracle Real Application Testing

Oracle Real Application Testing is an extremely cost-effective and easy-to-use proactive performance management solution that enables you to fully assess the outcome of a system change in test or production.

- About Oracle Real Application Testing
  You can use Oracle Real Application Testing to capture a workload on a production system and replay it on a test system with the exact timing, concurrency, and transaction characteristics of the original workload.
- Capture-Replay Workloads between Autonomous Databases
  You can Capture and Replay from an Autonomous Database instance into another Autonomous Database instance.
- Capture-Replay Workloads between non-Autonomous and Autonomous Databases
  You can Capture and Replay from a non-Autonomous Database instance into an Autonomous Database.
- Test Workloads Against an Upcoming Patch
  Using the workload auto replay feature you can automatically capture a workload from a production database that is on the regular patch level and replay the workload on a target refreshable clone that is on the early patch level.

About Oracle Real Application Testing

You can use Oracle Real Application Testing to capture a workload on a production system and replay it on a test system with the exact timing, concurrency, and transaction characteristics of the original workload.

Oracle Real Application Testing provides an accurate method to test the impact of a variety of system changes and enables you to perform the following tasks:

- You can test the effects of a system change on a workload without affecting the production system.
- You can capture a workload on a production system and simulate the same workload on a test system.

You can use Oracle Database Replay to capture a workload from either an Autonomous Database instance or an on-premises database or any other cloud service database, and replay the workload on Autonomous Database. This enables you to compare how a workload runs on an Autonomous Database with another Autonomous Database, an on-premises database, or some other cloud service database.

Real Application Testing enables you to do any of the following Capture-Replay actions:

Capture Replay Option	Description
Capture-Replay Workloads between Autonomous Databases.	See Capture-Replay Workloads between Autonomous Databases for more information.
Capture Workload from Autonomous Database Oracle Database 19c and Replay on Autonomous Database with Oracle Database 23ai	This Capture-Replay allows you to capture a workload on an Autonomous Database with Oracle Database 19c and replay on an Autonomous Database with Oracle Database 23ai. See Test 19c Workloads on a on 23ai Refreshable Clone with Capture-Replay for more information.
Capture a Workload from a non-Autonomous Database and Replay on Autonomous Database.	See Capture-Replay Workloads between non-Autonomous and Autonomous Databases for more information.
Capture a workload from a production Autonomous Database and replay it on a target Autonomous Database at a different patch level (after a patch is applied to the target Autonomous Database).	See Test Workloads Against an Upcoming Patch for more information.

Capture-Replay Workloads between Autonomous Databases

You can Capture and Replay from an Autonomous Database instance into another Autonomous Database instance.

This enables you to compare workloads across different Autonomous Database instances. These Autonomous Database instances may vary at patch levels, database versions, or regions.

The Capture-Replay workflow between Autonomous Databases consists of the following steps (you either cancel or finish a workload capture, not both):

- (Optional) Subscribe to Information Events to be Notified of Capture and Replay Details
  Subscribe to `com.oraclecloud.databaseservice.autonomous.database.information` Information events to be notified at the start and completion of a capture and replay.
- Capture a Workload on an Autonomous Database Instance
  The first step in using Database Replay is to capture a production workload.
- Cancel a Workload Capture on an Autonomous Database Instance
  Run `DBMS_CLOUD_ADMIN.CANCEL_WORKLOAD_CAPTURE` to cancel the current workload capture on your Autonomous Database instance.
- Finish a Workload Capture on Autonomous Database Instance
  Run `DBMS_CLOUD_ADMIN.FINISH_WORKLOAD_CAPTURE` to complete a workload capture on your Autonomous Database instance.
- Prepare a Refreshable Clone for Workload Replay
  Provides steps to prepare a refreshable clone for a workload replay.

- Replay a Workload on an Autonomous Database Instance
  After you complete a workload capture you can replay it on a test system. Oracle replays the actions recorded during workload capture with the same timing, concurrency, and transaction dependencies of the production system.
- Test 19c Workloads on a on 23ai Refreshable Clone with Capture-Replay
  You can use Oracle Real Application Testing to capture a workload from Autonomous Database on Oracle Database 19c and replay on a refreshable clone on Oracle Database 23ai.

(Optional) Subscribe to Information Events to be Notified of Capture and Replay Details

Subscribe to `com.oraclecloud.databaseservice.autonomous.database.information` Information events to be notified at the start and completion of a capture and replay.

> **Note:**
>
> This step is optional. You can also find status and historical information for a workload capture in the `DBA_CAPTURE_REPLAY_STATUS` and `DBA_CAPTURE_REPLAY_HISTORY` views.
> See DBA_CAPTURE_REPLAY_STATUS View and DBA_CAPTURE_REPLAY_HISTORY View for more information.

Information events provide notifications about begin and end times of capture and replay and contain a PAR URL to access the capture and replay reports.

Autonomous Database Information events include the following:

- `WorkloadCaptureBegin`: This event is triggered when a workload capture is initiated.
- `WorkloadCaptureEnd`: This event is triggered when a workload capture completes successfully and generates a pre-authenticated (PAR) URL to download the capture file.
- `WorkloadReplayBegin`: This event is triggered when a workload replay is initiated.
- `WorkloadReplayEnd`: This event is triggered when a workload replay completes successfully and generates a pre-authenticated (PAR) URL to download the replay reports.

See Information Events on Autonomous Database for more information.

Capture a Workload on an Autonomous Database Instance

The first step in using Database Replay is to capture a production workload.

> **Note:**
>
> You can capture a workload in an Autonomous Database instance and replay it in another Autonomous Database instance. You can replay the captured workload on a full clone or on a refreshable clone. The capture and replay targets must be in a consistent logical state. So, you must provision a refreshable clone or a full clone of the Autonomous Database instance on which you want to capture the workload.
> See Clone, Move, or Upgrade an Autonomous Database Instance for more information.

When you begin workload capture on a production system, all requests from external clients directed to Oracle Database are tracked and stored in binary files called capture files.

A workload capture results in the creation of two subdirectories, `cap` and `capfiles`, which contain the capture files. The capture files provide all pertinent information about the client request, including transaction details, bind values, and SQL text. These capture files are platform independent and can be transported to another system.

Run `DBMS_CLOUD_ADMIN.START_WORKLOAD_CAPTURE` to initiate workload capture on an Autonomous Database instance.

See Clone, Move, or Upgrade an Autonomous Database Instance for more information.

To initiate a workload capture on your Autonomous Database instance you must be logged in as the `ADMIN` user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.

Example to initiate a workload capture:

```
BEGIN
   DBMS_CLOUD_ADMIN.START_WORKLOAD_CAPTURE(
        capture_name => 'test',
        duration     => 60);
   END;
/
```

This starts the workload capture on your Autonomous Database instance.

The parameters are:

- `capture_name`: is the name of the workload capture.
- `duration`: is the duration (in minutes) for which you need to capture the workload. This parameter is optional.

See START_WORKLOAD_CAPTURE Procedure for more information.

Workload Capture Events

You can subscribe to the Information event `com.oraclecloud.databaseservice.autonomous.database.information` to be notified at the start of `START_WORKLOAD_CAPTURE`. See (Optional) Subscribe to Information Events to be Notified of Capture and Replay Details for more information.

Workload Capture and Replay Views

You can find information about workload capture and replay in the `DBA_CAPTURE_REPLAY_STATUS` and `DBA_CAPTURE_REPLAY_HISTORY` views. See DBA_CAPTURE_REPLAY_STATUS View and DBA_CAPTURE_REPLAY_HISTORY View for more information.

Cancel a Workload Capture on an Autonomous Database Instance

Run `DBMS_CLOUD_ADMIN.CANCEL_WORKLOAD_CAPTURE` to cancel the current workload capture on your Autonomous Database instance.

To cancel a workload capture, you must be logged in as the `ADMIN` user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.

Example:

```
BEGIN
    DBMS_CLOUD_ADMIN.CANCEL_WORKLOAD_CAPTURE;
END;
/
```

This cancels the current workload capture and performs a refresh on the refreshable clone.

You can query the `DBA_CAPTURE_REPLAY_STATUS` view to check the cancel workload status.

See DBA_CAPTURE_REPLAY_STATUS View for more information.

See CANCEL_WORKLOAD_CAPTURE Procedure for more information.

Finish a Workload Capture on Autonomous Database Instance

Run `DBMS_CLOUD_ADMIN.FINISH_WORKLOAD_CAPTURE` to complete a workload capture on your Autonomous Database instance.

Example to finish a workload capture on your Autonomous Database instance:

```
BEGIN
    DBMS_CLOUD_ADMIN.FINISH_WORKLOAD_CAPTURE;
END;
/
```

To run this procedure you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`. When you run this procedure a workload capture file is uploaded to Object Store as a zip file.

See FINISH_WORKLOAD_CAPTURE Procedure for more information.

Workload Capture Events

You can subscribe to the Information event `com.oraclecloud.databaseservice.autonomous.database.information` to be notified about a workload capture, that includes:

- The completion of `FINISH_WORKLOAD_CAPTURE`.
- The `captureDownloadURL` field that contains the PAR URL to access the capture and reports in Object Store. The capture and reports are valid for seven (7) days from the date when the PAR URL is generated.

See (Optional) Subscribe to Information Events to be Notified of Capture and Replay Details for more information.

Workload Capture and Replay Views

You can query the `DBA_CAPTURE_REPLAY_STATUS` view to check the status of a completed workload capture. See DBA_CAPTURE_REPLAY_STATUS View for more information.

You can find information about workload capture and replay in the `DBA_CAPTURE_REPLAY_HISTORY` view. See DBA_CAPTURE_REPLAY_HISTORY View for more information.

You can query the `ID`, `NAME`, `START_TIME`, and `END_TIME` columns of the `DBA_WORKLOAD_CAPTURES` view to retrieve the details of your workload capture. See DBA_WORKLOAD_CAPTURES for more information.

Prepare a Refreshable Clone for Workload Replay

Provides steps to prepare a refreshable clone for a workload replay.

> **Note:**
>
> This step is not applicable when you are replaying a workload on a full clone.

You have two options to prepare a refreshable clone to replay a workload capture. You can run `DBMS_CLOUD_ADMIN.PREPARE_REPLAY` to automatically prepare a refreshable clone for a workload replay. This procedure refreshes the refreshable clone to the start time of the capture and disconnects the refreshable clone. You also have the option to manually prepare a refreshable clone to replay a workload capture.

Automatically Prepare a Refreshable Clone for Workload Replay

Example to automatically prepare a refreshable clone for a workload replay:

```
BEGIN
 DBMS_CLOUD_ADMIN.PREPARE_REPLAY (
    capture_name    'test'
END;
/
```

To run this procedure you must be logged in as the `ADMIN` user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.

`DBMS_CLOUD_ADMIN.PREPARE_REPLAY` does the following:

- Refreshes the refreshable clone to the capture start timestamp.
- Disconnects the refreshable clone.

Optionally, at this point before you replay a capture, you can make changes to the a refreshable clone. For example, changing parameter values and turning certain features on/off to see the impact on the replay.

Manually Prepare a Refreshable Clone for Workload Replay

These manual refreshable clone steps are not required when you automatically prepare a refreshable clone by running `DBMS_CLOUD_ADMIN.PREPARE_REPLAY`.

Perform the following steps to manually prepare for a workload replay:

1. Find the capture start timestamp by querying the `DBA_WORKLOAD_CAPTURES` view. See DBA_WORKLOAD_CAPTURES for more information.
2. Refresh the refreshable clone to the capture start timestamp. See Refresh a Refreshable Clone on Autonomous Database for more information.
3. Manually disconnect the refreshable clone. See Disconnect a Refreshable Clone from the Source Database for more information.

4. Optionally, before you replay a capture, you can make changes to the refreshable clone. For example, changing parameter values, turning certain features on/off to see the impact on the replay.

Replay a Workload on an Autonomous Database Instance

After you complete a workload capture you can replay it on a test system. Oracle replays the actions recorded during workload capture with the same timing, concurrency, and transaction dependencies of the production system.

Run the procedure `DBMS_CLOUD_ADMIN.REPLAY_WORKLOAD` to initiate workload replay on your database. You must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN` to run `DBMS_CLOUD_ADMIN.REPLAY_WORKLOAD`.

You can replay a captured workload on a refreshable clone or on a full clone of the Autonomous Database instance from which the workload was captured. The capture and replay targets must be in a consistent logical state.

Replay Workload On a Refreshable Clone

The following example downloads the capture files from Object Storage, replays the captured workload, and uploads a replay report to Object Storage.

```
BEGIN
  DBMS_CLOUD_ADMIN.REPLAY_WORKLOAD(
      capture_name => 'CAP_TEST1');
END;
/
```

The `CAPTURE_NAME` parameter specifies the name of the workload capture. This parameter is mandatory.

Replay Workload On a Full Clone

This following example downloads capture files from the Object Storage, replays the captured workload on the clone, and uploads a replay report to Object Storage.

```
BEGIN
  DBMS_CLOUD_ADMIN.REPLAY_WORKLOAD(
       capture_name                => 'CAP_TEST1',
       capture_source_tenancy_ocid => 'OCID1.TENANCY.REGION1..ID1',
       capture_source_db_name      => 'ADWFINANCE');
END;
/
```

> **Note:**
>
> If there are multiple captures with the same capture name, the `REPLAY_WORKLOAD` procedure uses the latest capture. Oracle recommends that you use a unique capture name for each capture to prevent confusion on which capture you are replaying.

The `CAPTURE_NAME` parameter specifies the name of the workload capture. This parameter is mandatory.

The `CAPTURE_SOURCE_TENANCY_OCID` parameter specifies the source tenancy OCID of the workload capture. This parameter is mandatory when running the workload capture in a full clone.

The `CAPTURE_SOURCE_DB_NAME` parameter specifies the source database name of the workload capture. This parameter is mandatory when running the workload capture in a full clone.

See REPLAY_WORKLOAD Procedure for more information.

Workload Replay Events

Subscribe to the Information event `com.oraclecloud.databaseservice.autonomous.database.information` to be notified about the following

- The start and completion of a `REPLAY_WORKLOAD`.
- The Object Store link to download replay reports. The event provides a PAR URL to access the reports in the `replayDownloadURL` field. The reports are valid for seven (7) days from the date when the PAR URL is generated.

See Information Events on Autonomous Database for more information.

Workload Capture and Replay Views

You can query the `DBA_CAPTURE_REPLAY_STATUS` view to check the workload replay status.

See DBA_CAPTURE_REPLAY_STATUS View for more information.

You can find information about workload capture and replay in the `DBA_CAPTURE_REPLAY_HISTORY` view. See DBA_CAPTURE_REPLAY_HISTORY View for more information.

Test 19c Workloads on a on 23ai Refreshable Clone with Capture-Replay

You can use Oracle Real Application Testing to capture a workload from Autonomous Database on Oracle Database 19c and replay on a refreshable clone on Oracle Database 23ai.

This Capture-Replay option is especially useful when you want test your workloads running on Oracle Database 19c before you upgrade to Oracle Database 23ai.

Follow these steps to replay your workloads on a different database version:

1. (Optional) Subscribe to Information Events to be Notified of Capture and Replay Details
2. Capture your workloads on an Autonomous Database with Oracle Database 19c. See Capture a Workload on an Autonomous Database Instance for more information.
3. Run `DBMS_CLOUD_ADMIN.FINISH_WORKLOAD_CAPTURE` to complete the workload capture on your Oracle Database 19c sourceAutonomous Database instance.
4. Prepare your target Oracle Database 23ai Refreshable Clone. See Prepare a Refreshable Clone for Workload Replay for more information.
5. After you complete the workload capture and prepare the Oracle Database 23ai target Refreshable Clone, replay the workload on the target. See Replay a Workload on an Autonomous Database Instance for more information.

Capture-Replay Workloads between non-Autonomous and Autonomous Databases

You can Capture and Replay from a non-Autonomous Database instance into an Autonomous Database.

This enables you to compare workloads between an on-prem database or other cloud service database and an Autonomous Database instance.

Topics

- Capture a Workload
  The first step in using Database Replay is to capture the production workload.
- Replay a Workload on an Autonomous Database Instance
  After you complete a workload capture, you can replay it on a test system. Oracle replays on the test system the actions recorded during workload capture, with the same timing , concurrency, and transaction dependencies of the production system.

Capture a Workload

The first step in using Database Replay is to capture the production workload.

When you begin workload capture on the production system, all requests from external clients directed to Oracle Database are tracked and stored in binary files called capture files.

A workload capture results in the creation of two subdirectories, `cap` and `capfiles`, which contain the capture files.

The capture files provide all pertinent information about the client request, including transaction details, bind values, and SQL text.

These capture files are platform independent and can be transported to another system.

See Workload Capture to capture a workload on an on-premises database.

Replay a Workload on an Autonomous Database Instance

After you complete a workload capture, you can replay it on a test system. Oracle replays on the test system the actions recorded during workload capture, with the same timing , concurrency, and transaction dependencies of the production system.

Run `DBMS_CLOUD_ADMIN.REPLAY_WORKLOAD` to initiate workload replay on your database. You must be logged in as the `ADMIN` user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN` to run `REPLAY_WORKLOAD`.

Example to replay on an Autonomous Database instance a workload captured from an on-premises database:

```
BEGIN
   DBMS_CLOUD_ADMIN.REPLAY_WORKLOAD(
      location_uri    => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o',
      credential_name => 'CRED_TEST',
      synchronization => TRUE,
      process_capture => TRUE);
END;
/
```

This downloads the capture files contained in the Object Storage location specified in the `location_uri` parameter, and replays the workload capture from the capture files. The replay generates and uploads the replay and Automatic Workload Repository reports to the Object Storage location specified in the `location_uri` parameter.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

See Upload Files to Your Oracle Cloud Infrastructure Object Store Bucket for more information on uploading files to Object Storage.

See Navigate to Oracle Cloud Infrastructure Object Storage and Create Bucket for more information on Object Storage.

The `credential_name` parameter specifies the credential to access the object storage bucket. The credential that you supply must have the write privileges to write into the Object Storage bucket. The write privileges are required to upload the replay report to the bucket.

If you do not specify a `credential_name` value, then `DEFAULT_CREDENTIAL` is used.

You don't need to create a credential to access Oracle Cloud Infrastructure Object Store if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

The `synchronization` parameter specifies the synchronization method used during workload replay. A `TRUE` value indicates that the synchronization is SCN based.

The `process_capture` specifies whether you need to specify `process_capture` value or not. A `TRUE` value indicates that the replay includes `process_capture`.

> **Note:**
>
> You must maintain the same logical state of the capture and replay databases at the start of the capture time.

See REPLAY_WORKLOAD Procedure for more information.

Workload Replay Events

You can subscribe to the Information event `com.oraclecloud.databaseservice.autonomous.database.information` to be notified about the `WorkloadReplayBegin` and `WorkloadReplayEnd` events. These events provide information about:

- Start and completion of the `REPLAY_WORKLOAD`.
- The `replayDownloadURL` field that contains the PAR URL to access the reports in Object Store. The PAR URL is valid for 7 days from the date of generation.

See Information Events on Autonomous Database for more information.

Workload Capture and Replay Views

You can query the `DBA_CAPTURE_REPLAY_STATUS` and `DBA_CAPTURE_REPLAY_HISTORY` views to check the workload replay status.

See DBA_CAPTURE_REPLAY_STATUS View and DBA_CAPTURE_REPLAY_HISTORY View for more information.

Test Workloads Against an Upcoming Patch

Using the workload auto replay feature you can automatically capture a workload from a production database that is on the regular patch level and replay the workload on a target refreshable clone that is on the early patch level.

This feature allows you to test an upcoming patch by running an existing workload that is in production against a patch, before the patch reaches production.

- About Testing Workloads Against an Upcoming Patch
  Using the workload auto replay feature you can automate the process of capture-replay to capture a workload that runs on a production database and automatically replay the workload on a target refreshable clone after an upcoming patch is applied on the target.
- Enable Workload Auto Replay
  The `WORKLOAD_AUTO_REPLAY` feature allows you run a workload from your production database and monitor for any divergence on an instance that is patched one week in advance. This feature allows you to test an upcoming patch by running an existing workload that is in production against a patch before the patch reaches production.
- Disable Workload Auto Replay
  Run `DBMS_CLOUD_ADMIN.DISABLE_FEATURE` to disable `WORKLOAD_AUTO_REPLAY`.

About Testing Workloads Against an Upcoming Patch

Using the workload auto replay feature you can automate the process of capture-replay to capture a workload that runs on a production database and automatically replay the workload on a target refreshable clone after an upcoming patch is applied on the target.

Autonomous Database provides the ability to provision an instance or create a refreshable clone with the **Early** patch level option. On instances running at the **Early** patch level, Autonomous Database applies upcoming maintenance patches a week before the patches are applied to production databases (databases that are provisioned at the **Regular** patch level). Using the `WORKLOAD_AUTO_REPLAY` feature you can assure that an upcoming patch is tested against your workload before the patch goes to production. This allows you to verify that the patch either fixes a known issue or does not introduce an issue that affects your workload.

To find information about captures and replays, subscribe to Information events. Information events provide notification for workload capture and reply events and include a PAR URL where you can download the capture file and replay report. See (Optional) Subscribe to Information Events to be Notified of Capture and Replay Details for more information.

When `WORKLOAD_AUTO_REPLAY` is enabled the source database captures a workload by running for a specified number of minutes. By default the workload capture starts when you enable `WORKLOAD_AUTO_REPLAY` (optionally you can use parameters to set the capture start day and time). Next, Autonomous Database checks the target database to verify the patching status. After the upcoming weekly patch is applied, Autonomous Database replays the workload on the target database. This capture-replay cycle continues automatically each week with Autonomous Database capturing the workload on the source database, waiting for the upcoming patch to be applied, and replaying the workload on the refreshable clone.

Note the following for enabling `WORKLOAD_AUTO_REPLAY`:

- The source database must use the **Regular** patch level.
- The target database must use the **Early** patch level.

- The target database must be a refreshable clone of the source database, and must be created before you enable `WORKLOAD_AUTO_REPLAY`.
- A source database can enable `WORKLOAD_AUTO_REPLAY` for no more than one refreshable clone (you can enable this feature for a maximum of one refreshable clone, even if you create multiple refreshable clones from the same source database).
- After you enable `WORKLOAD_AUTO_REPLAY`, the capture-replay cycle continues every week. Autonomous Database runs a capture on the source database and then replays the workload on the target database, until you disable `WORKLOAD_AUTO_REPLAY`.

You can find information about workload captures and replays in the `DBA_CAPTURE_REPLAY_HISTORY` view. See DBA_CAPTURE_REPLAY_HISTORY View for more information.

Autonomous Database automatically applies patches on your database. Oracle provides a service level objective of zero regressions in your production database due to these patches. See Zero-Regression Service Level Objective for more information.

Enable Workload Auto Replay

The `WORKLOAD_AUTO_REPLAY` feature allows you run a workload from your production database and monitor for any divergence on an instance that is patched one week in advance. This feature allows you to test an upcoming patch by running an existing workload that is in production against a patch before the patch reaches production.

To enable `WORKLOAD_AUTO_REPLAY`:

1. Create a refreshable clone of the production database.

   When you create the target refreshable clone, set the patch level to **Early**.

   See Set the Patch Level and Create a Refreshable Clone for an Autonomous Database Instance for more information.

2. Run `DBMS_CLOUD_ADMIN.ENABLE_FEATURE` on the source database.

   For example:

```
BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_FEATURE(
        feature_name => 'WORKLOAD_AUTO_REPLAY',
        params       => JSON_OBJECT(
                          'target_db_ocid' VALUE
'OCID1.autonomousdatabase.REGION..ID1',
                          'capture_duration' VALUE 120,
                          'capture_day' VALUE 'MONDAY',
                          'capture_time' VALUE '15:00'));
END;
/
```

   Where the parameters are:

   - `feature_name`: the value `WORKLOAD_AUTO_REPLAY` enables the workload auto replay feature.
   - `params`: is a JSON Object with the following value pairs:
     - `target_db_ocid`: accepts a `string` value. The value specifies the OCID of the target refreshable clone database on which the captured workload is replayed.

This parameter is mandatory.

- `capture_duration`: accepts a `number` value. The value specifies the duration in minutes for which the workload is captured on the production database. The value must be in the range between 1 and 720 minutes.

  This parameter is mandatory.

- `capture_day`: accepts a `string` value. The value specifies the day of the week the workload capture on the production database should begin.

  This parameter is optional.

- `capture_time`: accepts a value in the `HH24:MM` format. The value specifies the time of the day the workload capture on the production database should begin.

  This parameter is optional.

By default the workload capture starts when you enable `WORKLOAD_AUTO_REPLAY`. When the optional `capture_day` and `capture_time` are specified, the automatic workload capture and replay happen at the specified timestamp.

For example if `capture_day` is Monday and `capture_time` is 15:00, the first capture on the production database starts at 3PM on the next Monday. The same day of week and time are also used to schedule subsequent captures and replays.

See ENABLE_FEATURE Procedure for more information.

A an error value of `ORA-20000: Invalid argument for target_db_ocid` could indicate that the OCID you supplied is not a refreshable clone. In this case, you need to supply an OCID with a value for a refreshable clone.

3. Query the `DBA_CAPTURE_REPLAY_STATUS` view to check the workload replay status.

   See Replay a Workload on an Autonomous Database Instance and DBA_CAPTURE_REPLAY_STATUS View for more information.

This example enables `WORKLOAD_AUTO_REPLAY` on the source Autonomous Database and on the specified target refreshable clone database. With `WORKLOAD_AUTO_REPLAY` enabled, every week Autonomous Database runs a capture on the source database and replays the workload on the target database, until you disable `WORKLOAD_AUTO_REPLAY`.

To find information about captures and replays, subscribe to Information events. Information events provide notification for workload capture and reply events and include a PAR URL where you can download the capture file and replay report. See (Optional) Subscribe to Information Events to be Notified of Capture and Replay Details for more information.

You can find information about workload captures and replays in the `DBA_CAPTURE_REPLAY_HISTORY` view. See DBA_CAPTURE_REPLAY_HISTORY View for more information.

Disable Workload Auto Replay

Run `DBMS_CLOUD_ADMIN.DISABLE_FEATURE` to disable `WORKLOAD_AUTO_REPLAY`.

Run `DBMS_CLOUD_ADMIN.DISABLE_FEATURE` to disable workload auto replay. For example:

```
BEGIN
DBMS_CLOUD_ADMIN.DISABLE_FEATURE(
    feature_name => 'WORKLOAD_AUTO_REPLAY');
END;
/
```

You must be logged in as ADMIN or have DBMS_CLOUD_ADMIN privileges to run DBMS_CLOUD_ADMIN.DISABLE_FEATURE.

See DISABLE_FEATURE Procedure for more information.

Manage and Store Files in a Cloud Code Repository with Autonomous Database

Autonomous Database provides routines to manage and store files in Cloud Code (Git) Repositories. The supported Cloud Code Repositories are: GitHub, AWS CodeCommit, and Azure Repos.

- About Cloud Code Repositories with Autonomous Database
  The DBMS_CLOUD_REPO package provides a single interface for accessing a Cloud Code Repository from Autonomous Database.
- Initialize a Cloud Code Repository
  The DBMS_CLOUD_REPO initialization routines initialize a Cloud Code Repository. After you obtain a Cloud Code Repository handle, you use the handle to access the Cloud Code Repository.
- Create and Manage a Cloud Code Repository
  The DBMS_CLOUD_REPO management routines allow you to manage a Cloud Code Repository by creating, listing, updating, or deleting a repository.
- Create and Manage Branches in a Cloud Code Repository
  The DBMS_CLOUD_REPO management routines allow you to manage Cloud Code Repository branches by creating, listing, merging, or deleting branches in a repository.
- Export Schema Objects to the Cloud Code Repository Branch
  The DBMS_CLOUD_REPO management routine allows you to export metadata of the objects in a schema to the Cloud Code Repository branch. You can filter your list based on the object names or object types.
- Use File Operations with a Cloud Code Repository
  The DBMS_CLOUD_REPO file operations allow you to create, get, list, update, or delete files in a Cloud Code Repository.
- Use SQL Install Operations with a Cloud Code Repository
  The DBMS_CLOUD_REPO SQL Install operations allow you to store and download SQL scripts from a Cloud Code Repository.

About Cloud Code Repositories with Autonomous Database

The DBMS_CLOUD_REPO package provides a single interface for accessing a Cloud Code Repository from Autonomous Database.

The supported Cloud Code Repositories provide the following features:

- Git Version Control System: Git is software for tracking changes in any set of files, usually used for coordinating work among programmers collaboratively developing source code during software development. Its goals include speed, data integrity, and support for distributed, non-linear workflows.
- Git Repository: A Git repository is a virtual storage of your project. It allows you to save versions of your code, which you can access when needed.

The DBMS_CLOUD_REPO APIs use a repository handle (REPO object). The repository handle is an opaque JSON object that represents a Cloud Code Repository of a specific cloud provider. A

REPO object can be passed to different DBMS_CLOUD_REPO APIs. This opaque object ensures that DBMS_CLOUD_REPO procedures and functions are multicloud compatible; you do not have to change your code when you migrate from one Cloud Code Repository provider to another Cloud Code Repository.

Autonomous Database provides the following to help you work with Cloud Code Repositories:

- Repository initialization operations to allow you to initialize a repository.

  See Initialize a Cloud Code Repository for more information.
- Repository management operations that let you create, list, update or delete a repository.

  See Create and Manage a Cloud Code Repository for more information.
- Repository branch management operations that let you create, list, merge or delete branches in a repository.

  See Create and Manage Branches in a Cloud Code Repository for more information.
- Export the metadata DDL of all objects in a schema to a repository.

  See Export Schema Objects to the Cloud Code Repository Branch for more information.
- Repository file management operations to upload, download, update, and delete files.

  See Use File Operations with a Cloud Code Repository for more information.
- SQL install operations that let you export database object metadata DDL to a repository and install SQL statements into the database from a Cloud Code Repository.

  See Use SQL Install Operations with a Cloud Code Repository for more information.

Initialize a Cloud Code Repository

The DBMS_CLOUD_REPO initialization routines initialize a Cloud Code Repository. After you obtain a Cloud Code Repository handle, you use the handle to access the Cloud Code Repository.

To initialize a Cloud Code Repository:

1. Create a credential to access the Cloud Code Repository.

   See CREATE_CREDENTIAL Procedure for information on creating credentials.
2. Depending on the repository, GitHub, Azure Repos, or AWS CodeCommit, call DBMS_CLOUD_REPO.INIT_REPO with the parameters for the particular repository to obtain a repository handle.

   The following examples provide samples for each supported Cloud Code Repository.

   - **GitHub Initialization:**

```
DEFINE repo_name='test_repo';
DEFINE cred_name='GITHUB_CRED';
VAR repo clob
BEGIN
  :repo := DBMS_CLOUD_REPO.INIT_REPO(
               params => JSON_OBJECT('provider' value 'github',
                                     'repo_name' value '&repo_name',
                                     'credential_name' value
'&cred_name',
                                     'owner'           value
'<myuser>')
          );
```

```
END;
/
```

- **Azure Repos Initialization:**

```
DEFINE repo_name='test_repo';
DEFINE cred_name='AZURE_REPO_CRED';
VAR repo clob
BEGIN
  :repo := DBMS_CLOUD_REPO.INIT_REPO(
               params => JSON_OBJECT('provider' value 'azure',
                                     'repo_name' value '&repo_name',
                                     'credential_name' value
'&cred_name',
                                     'organization' value '<myorg>',
                                     'project' value '<myproject>')
          );
END;
/
```

- **AWS CodeCommit Initialization:**

```
DEFINE repo_name='test_repo';
DEFINE cred_name='AWS_REPO_CRED';
VAR repo clob
BEGIN
  :repo := DBMS_CLOUD_REPO.INIT_REPO(
               params => JSON_OBJECT('provider' value 'aws',
                                     'repo_name' value '&repo_name',
                                     'credential_name' value
'&cred_name',
                                     'region' value 'us-east-1')
          );
END;
/
```

See DBMS_CLOUD_REPO Initialization Operations for details on the initialization functions.

Create and Manage a Cloud Code Repository

The `DBMS_CLOUD_REPO` management routines allow you to manage a Cloud Code Repository by creating, listing, updating, or deleting a repository.

First, obtain a Cloud Code Repository handle to provide access a repository. See Initialize a Cloud Code Repository for details.

1. To create a repository:

```
VAR repo clob
BEGIN
  DBMS_CLOUD_REPO.CREATE_REPOSITORY(
        repo => :repo,
        description => 'test repo'
  );
```

```
END;
/
```

2. To update a repository:

```
VAR repo clob
DEFINE repo_name='test_repo';
BEGIN
  DBMS_CLOUD_REPO.UPDATE_REPOSITORY(
        repo => :repo,
        new_name => '&repo_name' || '_new'
  );
END;
/
```

3. To list repositories:

```
col id format a30
col name format a10
col description format a15
select id, name, bytes, private, description from
  DBMS_CLOUD_REPO.LIST_REPOSITORIES(:repo);
```

4. To delete a repository:

```
VAR repo clob
BEGIN
  DBMS_CLOUD_REPO.DELETE_REPOSITORY(
        repo => :repo
  );
END;
/
```

See DBMS_CLOUD_REPO Repository Management Operations for more information.

Create and Manage Branches in a Cloud Code Repository

The `DBMS_CLOUD_REPO` management routines allow you to manage Cloud Code Repository branches by creating, listing, merging, or deleting branches in a repository.

To perform Cloud Code Repository branch management operations you must first:

- Create a credential.

  See CREATE_CREDENTIAL Procedure for details.
- Obtain a handle.

  See Initialize a Cloud Code Repository for details.
- Create a repository.

  See Create and Manage a Cloud Code Repository for details.
- Log in as the ADMIN user or have `EXECUTE` privilege on `DBMS_CLOUD_REPO`.

1. To create a branch in a Cloud Code Repository:

```
BEGIN
  DBMS_CLOUD_REPO.CREATE_BRANCH (
    repo            => l_repo,
    branch_name     => 'test_branch',
    parent_branch   => 'main'
    );
END;
/
```

2. To delete a branch in a Cloud Code Repository:

```
BEGIN
  DBMS_CLOUD_REPO.DELETE_BRANCH (
      repo        => l_repo,
      branch_name => 'test_branch'
  );
END;
/
```

3. To merge a branch into another branch in a Cloud Code Repository:

```
BEGIN
  DBMS_CLOUD_REPO.MERGE_BRANCH (
    repo               => l_repo,
    branch_name        => 'test_branch',
    target_branch_name => 'main'
    );
END;
/
```

4. To list branches in a Cloud Code Repository:

```
SELECT * FROM DBMS_CLOUD_REPO.LIST_BRANCHES (repo => l_repo);
```

5. You can list commits in a branch in Cloud Code Repository based on the following;
   - Repository
   - Branch
   - The file path of the branch
   - SHA/commit_id

> **Note:**
>
> SHA is a checksum 40-character string composed of hexadecimal characters (0–9 and a–f).

To list commits based on the repository:

```
SELECT * FROM DBMS_CLOUD_REPO.LIST_COMMITS (
    repo => :repo);
```

To list commits based on the repository, commit_id and file path of the branch:

```
SELECT * FROM DBMS_CLOUD_REPO.LIST_COMMITS (
    repo        => l_repo,
    commit_id   => '66dd2b23b74cd0afabd11af66c6aa9c550540ba6',
    file_path   => 'sub_dir/test11.sql'
```

To list commits based on the repository, and file path of the branch:

```
SELECT * FROM DBMS_CLOUD_REPO.LIST_COMMITS (
    repo        => :repo,
    branch_name => 'branch1');
```

To list commits based on the repository, branch_name and file path of the branch:

```
SELECT * FROM DBMS_CLOUD_REPO.LIST_COMMITS (
    repo         => :repo,
    branch_name  => 'branch1',
    file_path    => 'sub_dir/test11.sql');
```

To list commits based on the repository and commit_id of the branch:

```
SELECT * FROM DBMS_CLOUD_REPO.LIST_COMMITS (
    repo => :repo,
    commit_id =>'66dd2b23b74cd0afabd11af66c6aa9c550540ba6');
```

See DBMS_CLOUD_REPO Repository Branch Management Operations for more information.

Export Schema Objects to the Cloud Code Repository Branch

The `DBMS_CLOUD_REPO` management routine allows you to export metadata of the objects in a schema to the Cloud Code Repository branch. You can filter your list based on the object names or object types.

To export schema metadata you must first:

- Create a credential.

  See CREATE_CREDENTIAL Procedure for details.
- Obtain a handle.

  See Initialize a Cloud Code Repository for details.
- Create a repository.

  See Create and Manage a Cloud Code Repository for details.
- Log in as the ADMIN user or have `EXECUTE` privilege on `DBMS_CLOUD_REPO`.

Use the `EXPORT_SCHEMA` procedure to export metadata of the objects in your schema to a Cloud Code Repository branch:

```
BEGIN
  DBMS_CLOUD_REPO.EXPORT_SCHEMA(
    repo          => l_repo,
    schema_name   => 'USER1',
    file_path     => 'myschema_ddl.sql'
    filter_list   =>
        to_clob('[
             { "match_type":"equal",
                 "type":"table"
             },
             { "match_type":"not_equal",
                 "type":"view"
             },
             { "match_type":"in",
                 "type":"table",
                 "name": " ''EMPLOYEE_SALARY'',''EMPLOYEE_ADDRESS'' "
             },
             { "match_type":"equal",
                 "type":"sequence",
                 "name": "EMPLOYEE_RECORD_SEQ"
             },
             { "match_type":"like",
                 "type":"table",
                 "name": "%OFFICE%"
             }
        ]'
     );
  );
END;
/
```

This example exports the metadata of the `USER1` schema into the `l_repo` repository. The export includes the metadata of the tables `EMPLOYEE_SALARY` and `EMPLOYEE_ADDRESS`, and any table name that contains `OFFICE`. It also exports the `EMPLOYEE_RECORD_SEQ` sequence and excludes the views in the schema.

See EXPORT_SCHEMA Procedure for more information.

Use File Operations with a Cloud Code Repository

The `DBMS_CLOUD_REPO` file operations allow you to create, get, list, update, or delete files in a Cloud Code Repository.

Obtain a Cloud Code Repository handle before using the file operations. See Initialize a Cloud Code Repository for details.

You also need to create a repository before you work with files. See Create and Manage a Cloud Code Repository for details.

1. To get a file:

```
SELECT DBMS_CLOUD_REPO.GET_FILE(repo => :repo, file_path => 'test1.sql')
```

2. To create a file:

```
VAR repo clob
BEGIN
  DBMS_CLOUD_REPO.PUT_FILE(
        repo => :repo,
        file_path => 'test1.sql',
        contents => UTL_RAW.cast_to_raw('create table t1 (x varchar2(30))'
|| CHR(10) || '/')
  );
END;
/
```

3. To update a file:

```
VAR repo clob
BEGIN
  DBMS_CLOUD_REPO.PUT_FILE(
        repo => :repo,
        file_path => 'test1.sql',
        contents => UTL_RAW.cast_to_raw('create table t2 (x varchar2(30))'
|| CHR(10) || '/')
  );
END;
/
```

4. To list files:

```
SELECT id, name, bytes, url FROM DBMS_CLOUD_REPO.LIST_FILES(repo => :repo);
```

5. To delete a file:

```
VAR repo clob
BEGIN
  DBMS_CLOUD_REPO.DELETE_FILE(
        repo => :repo,
        file_path => 'test1.sql'
  );
END;
/
```

See DBMS_CLOUD_REPO File Operations for more information.

Use SQL Install Operations with a Cloud Code Repository

The `DBMS_CLOUD_REPO` SQL Install operations allow you to store and download SQL scripts from a Cloud Code Repository.

Obtain a Cloud Code Repository handle before using the SQL Install operations. See Initialize a Cloud Code Repository for details.

You also need to create a repository before you work with SQL Install operations. See Create and Manage a Cloud Code Repository for details.

The scripts are intended as schema install scripts and not as generic SQL scripts:

- Scripts cannot contain SQL*Plus client specific commands.
- Scripts cannot contain bind variables or parameterized scripts.
- SQL statements must be terminated with a slash on a new line (/).
- Scripts can contain DDL, DML PLSQL statements, but direct `SELECT` statements are not supported. Using SELECT within a PL/SQL block is supported.

Any SQL statement that can be run using `EXECUTE IMMEDIATE` will work if it does not contain bind variables or defines.

1. To upload DDL metadata to a Cloud Code Repository:

```
VAR repo clob
BEGIN
  DBMS_CLOUD_REPO.EXPORT_OBJECT(
        repo => :repo,
        object_type => 'PACKAGE',
        object_name => 'MYPACK',
        file_path   => 'mypack.sql'
  );
END;
/
```

2. To install SQL statements from a file:

```
VAR repo clob
BEGIN
  DBMS_CLOUD_REPO.INSTALL_FILE(
        repo => :repo,
        file_path     => 'test3.sql',
        stop_on_error => FALSE
  );
END;
/
```

3. To install SQL statements from a buffer:

```
VAR repo clob
BEGIN
  DBMS_CLOUD_REPO.INSTALL_SQL(
        repo => :repo,
        content   => 'create table t1 (x varchar2(30))' || CHR(10) || '/',
        stop_on_error => FALSE
  );
END;
/
```

See DBMS_CLOUD_REPO SQL Install Operations for more information.

Invoke User Defined Functions

You can invoke external functions such as OCI, AWS Lambda, Google Cloud Run functions and Azure remote functions in your Autonomous Database as SQL functions.

- About User Defined Functions
  User defined functions enable you to invoke externally available functions from PL/SQL or SQL code within your database.
- Steps to Invoke OCI Cloud Functions as SQL Functions
  Shows the steps to invoke OCI remote functions as SQL functions in your database.
- Steps to Invoke AWS Lambda as SQL Functions
  Shows the steps to invoke AWS remote functions as SQL functions in your database.
- Steps to Invoke Azure Functions as SQL Functions
  Shows the steps to invoke Azure remote functions as SQL functions in your database using HTTP trigger.
- Steps to Invoke Google Cloud Run Functions as SQL Functions
  Shows the steps to invoke Google Cloud Run functions as SQL functions in your database.
- Invoke External Procedures as SQL Functions
  Shows the steps to invoke external procedures using PL/SQL within your database.

About User Defined Functions

User defined functions enable you to invoke externally available functions from PL/SQL or SQL code within your database.

You can invoke the following external functions using user defined functions:

- **Oracle Cloud Infrastructure Functions**: Oracle Cloud Infrastructure Functions are fully managed, multi-tenant, highly scalable, on-demand, Functions-as-a-Service platform. Oracle Cloud Infrastructure Functions are built on enterprise-grade Oracle Cloud Infrastructure and powered by the Fn Project open-source engine. Use Oracle Cloud Infrastructure Functions also referred to as OCI Functions, to focus on developing code to meet the business requirements.

  See Overview of Functions for more information.
- **AWS Lambda**: AWS Lambda is a serverless, event-driven compute service that lets you run code for virtually any application or backend service without provisioning or managing servers.

  See AWS Lambda for more information.
- **Azure Function**: Azure Function is a serverless compute service that enables users to run event-triggered code without provisioning or managing infrastructure.

  See Azure Functions overview and Create your first function in the Azure portal for more information.
- **External Procedures**: External procedures are functions written in a third-generation language (C, for example) and callable from within PL/SQL or SQL as if they were a PL/SQL routine or function.

  See What Is an External Procedure? for more information.
- **Google Cloud Run Function**: Google Cloud Run function is a serverless compute service that enables user to run event-triggered code without having to provision or manage infrastructure.

See Google Cloud Run functions for more information.

Steps to Invoke OCI Cloud Functions as SQL Functions

Shows the steps to invoke OCI remote functions as SQL functions in your database.

To invoke cloud functions in OCI as SQL functions, we will be creating a catalog of SQL wrapper functions that reference and call their respective cloud function via their API endpoints. Before you create this catalog, it is assumed here that you have created the necessary cloud functions to be referenced by this catalog.

See Creating and Deploying Functions for more information on the creation and deployment of Oracle Cloud Infrastructure Functions and Application (ie a group of OCI Functions).

1. Create credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`. The credential object that you are creating must be of type private key.

```
SET DEFINE OFF
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL (
       credential_name => 'OCI_CRED',
       user_ocid       => 'user_ocid',
       tenancy_ocid    => 'tenancy_ocid',
       private_key     => 'private_key',
       fingerprint     => 'fingerprint'
  );
END;
/

PL/SQL procedure successfully completed.
```

   This creates the `OCI_CRED` credential.

   See CREATE_CREDENTIAL Procedure for more information.

> **Note:**
>
> If you are using a Resource Principal for authentication, the necessary policies required for OCI Function access must be configured. See Details for Functions for more information.

2. Create a catalog.

   A catalog is a collection of wrapper functions that reference and call their respective cloud functions via their API endpoints.

   Example to create a catalog for Oracle Cloud Infrastructure Functions.

```
BEGIN
DBMS_CLOUD_FUNCTION.CREATE_CATALOG (
    credential_name  => 'OCI_CRED',
    catalog_name     => 'OCI_DEMO_CATALOG',
    service_provider => 'OCI',
    cloud_params     => '{"region_id":"phx",
      "compartment_id":"compartment_id"}'
);
```

```
END;
/

PL/SQL procedure successfully completed.
```

This creates the `OCI_DEMO_CATALOG` catalog.

See CREATE_CATALOG Procedure for more information. You can query DBA_CLOUD_FUNCTION_CATALOG View and USER_CLOUD_FUNCTION_CATALOG View views to retrieve the list of all the catalogs in your database.

3. You can list the functions in a catalog.

   Following is the example to list Oracle Cloud Infrastructure Functions:

```
VAR function_list CLOB;

BEGIN
    DBMS_CLOUD_FUNCTION.LIST_FUNCTIONS (
      credential_name  => 'OCI_CRED',
      catalog_name     => 'OCI_DEMO_CATALOG',
      function_list    => :function_list
 );
END;
/

PL/SQL procedure successfully completed.

SELECT JSON_QUERY (:function_list, '$' RETURNING VARCHAR2(32676) pretty)
AS search_results FROM dual;
SEARCH_RESULTS                          --This is a sample output
--------------------------------------------------------------------------------
--------------------
[
  {
    "functionName"   : "create_par",
    "functionId"     : "ocid.funfc.oc1.phx.aaaa_example",
    "invokeEndpoint" : "https://dw.us.func.oci.oraclecloud_example.com"
  },
  {
    "functionName"   : "fintech",
    "functionId"     : "ocid.funfc.oc1.phx.bbbb_example"
SEARCH_RESULTS
--------------------------------------------------------------------------------
--------------------
4ayosyosv4sthmya2lyza",
    "invokeEndpoint" : "https://dw.us.func.oci.oraclecloud.com_example"
  },
  {
    "functionName"   : "jwt_codec",
    "functionId"     : "ocid.funfc.oc1.phx.jwt_code_example",
    "invokeEndpoint" : "https://dw.us.func.oci.oraclecloud_example.com"
  },
SEARCH_RESULTS
--------------------------------------------------------------------------------
--------------------
  {
```

```
    "functionName"   : "oci-objectstorage-create-par-python",
    "functionId"     : "ocid.funfc.oc1.phx.aaaaaaaas_example",
    "invokeEndpoint" : "https://dw.us.func.oci.oraclecloud_example.com"
  },
  {
    "functionName"   : "run_dbt",
    "functionId"     : "ocid.funfc.oc1.phx.aaaaaaaav_example",

SEARCH_RESULTS
--------------------------------------------------------------------------------
----------------------
    "invokeEndpoint" : "https://dw.us.func.oci.oraclecloud_example.com"
  }
]
```

See LIST_FUNCTIONS Procedure for more information.

4. Run the `DBMS_CLOUD_FUNCTION.SYNC_FUNCTIONS` to create wrapper SQL functions. You can use one of the following methods to create the wrapper SQL functions in the catalog, that call their respective cloud functions:

   - **SYNC_FUNCTIONS**: `SYNC_FUNCTIONS` is the quickest and simplest method, which automatically syncs (creates or deletes) wrapper functions in the catalog with the complete list of cloud functions defined in the region, compartment, and tenancy with which the catalog was created. For example:

```
BEGIN
 DBMS_CLOUD_FUNCTION.SYNC_FUNCTIONS (
        catalog_name => 'OCI_DEMO_CATALOG'
 );
END;
/
PL/SQL procedure successfully completed.
```

This creates a PL/SQL wrapper for adding new functions to the catalog and removing wrappers for functions that have been deleted from the catalog.

Run the following query to verify the sync.

```
SELECT object_name FROM sys.all_objects WHERE owner='TEST_USER' AND
object_type='FUNCTION';

OBJECT_NAME
--------------------------------------------------------------------------------
-----
CREATE_PAR
FINTECH
JWT_CODEC
OCI-OBJECTSTORAGE-CREATE-PAR-PYTHON
RUN_DBT
```

> **Note:**
>
> Keep a note of the current user in order to run this command.

See SYNC_FUNCTIONS Procedure for more information.

- You can manually create a `SQL` Function in your catalog that calls its respective cloud function using `DBMS_CLOUD.CREATE_FUNCTION`. For example:

Example to create a function in the `OCI_DEMO_CATALOG` catalog.

```
VAR function_args CLOB;
EXEC :function_args := TO_CLOB('{"command": "VARCHAR2", "value":
"VARCHAR2"}');
BEGIN
    DBMS_CLOUD_FUNCTION.CREATE_FUNCTION (
        credential_name  => 'OCI_CRED',
        catalog_name     => 'OCI_DEMO_CATALOG',
        function_name    => 'fintech_fun',
        function_id      => 'ocid1.fnfunc.oc1.phx.aaabbbcccc_example',
        input_args       => :function_args
 );
END;
/

PL/SQL procedure successfully completed.
```

This creates the `FINTECH_FUN` function in the `OCI_DEMO_CATALOG` catalog.

The `FINTECH_FUN` function in the catalog is a reference to the respective cloud function whose endpoint is referenced by the `FUNCTION_ID` parameter. Invoking the function in the catalog along with its arguments runs the corresponding cloud function and provides the output returned by the cloud function.

Manually creating a function in the catalog also allows you to create custom return types and response handlers. For example:

```
CREATE OR REPLACE TYPE fintech_rt AS OBJECT (
    STATUS  VARCHAR2(1000),
    OUTPUT  CLOB
);
/

Type created.


CREATE OR REPLACE FUNCTION fintech_response_handler(function_response in
CLOB)
RETURN fintech_rt
IS
l_comp   fintech_rt;
l_json_obj  JSON_OBJECT_T;
status  VARCHAR2(1000);
output  CLOB;
BEGIN
```

```
      l_json_obj := JSON_OBJECT_T.parse(function_response);
      status     := l_json_obj.get('STATUS').to_string;
      output     := l_json_obj.get('RESPONSE_BODY').to_string;
      l_comp     := fintech_rt(status,output);
      RETURN l_comp;
END;
/
Function created.

VAR input_param clob;
VAR  l_return_type varchar2(100);
VAR  l_response_handler varchar2(1000);

exec :input_param       := TO_CLOB('{"command": "VARCHAR2", "value":
"VARCHAR2"}');
exec :l_return_type     := 'fintech_rt';
exec :l_response_handler := 'fintech_response_handler';

PL/SQL procedure successfully completed.

PL/SQL procedure successfully completed.

PL/SQL procedure successfully completed.

EXEC DBMS_CLOUD_FUNCTION.CREATE_FUNCTION (CREDENTIAL_NAME => 'OCI_CRED',
FUNCTION_NAME => 'fintech_fun', CATALOG_NAME => 'OCI_DEMO_CATALOG',
FUNCTION_ID => 'ocid1.funfn.oci.phx.aaaaaa_example', INPUT_ARGS
=> :input_param, RETURN_TYPE => :l_return_type ,RESPONSE_HANDLER
=>:l_response_handler);

PL/SQL procedure successfully completed.
```

You can query DBA_CLOUD_FUNCTION View and USER_CLOUD_FUNCTION View views to retrieve the list of all the functions in your database.

See CREATE_FUNCTION Procedure for more information.

5. After the function is created you can `DESCRIBE` and invoke it.

```
DESC fintech_fun
COLUMN STATUS format a30
COLUMN OUTPUT format a30


DECLARE
l_comp fintech_rt;
BEGIN
l_comp := fintech_fun(command=>'tokenize',value => 'PHI_INFORMATION');
DBMS_OUTPUT.put_line ('Status of the function   =  '|| l_comp.status);
DBMS_OUTPUT.put_line ('Response of the function =  '|| l_comp.output);
END;
/

PL/SQL procedure successfully completed.
```

This invokes the `fintech_fun` cloud function by calling the function reference `oocid1.funfn.oci.phx.aaaaaa_example` in the `OCI_DEMO_CATALOG` catalog.

6. You can drop an existing function using `DROP_FUNCTION` procedure. For example:

```
EXEC DBMS_CLOUD_FUNCTION.DROP_FUNCTION (CATALOG_NAME =>
'OCI_DEMO_CATALOG', FUNCTION_NAME => 'fintech_fun');

PL/SQL procedure successfully completed.
```

This drops the `FINTECH_FUN` function from the `OCI_DEMO_CATALOG` catalog.

See DROP_FUNCTION Procedure for more information.

7. You can drop an existing catalog using `DROP_CATALOG` procedure. For example:

```
BEGIN
    DBMS_CLOUD_FUNCTION.DROP_CATALOG (
      catalog_name     => 'OCI_DEMO_CATALOG'
  );
END;
/

PL/SQL procedure successfully completed.
```

This drops the `OCI_DEMO_CATALOG` from your database.

See DROP_CATALOG Procedure for more information.

Steps to Invoke AWS Lambda as SQL Functions

Shows the steps to invoke AWS remote functions as SQL functions in your database.

To invoke AWS Lambda as SQL functions, we will be creating a catalog of SQL wrapper functions that reference and call their respective cloud function via their API endpoints. Before you create this catalog, it is assumed here that you have created the necessary cloud functions to be referenced by this catalog.

> **Note:**
>
> To access AWS lambda functions you need to configure the necessary policies. See Creating an IAM policy to access AWS Lambda resources and Using resource-based policies for Lambda for more information.

1. Create a credential using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`. The credential object that you are creating must be of type secret key.

```
SET DEFINE OFF
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL (
    credential_name => 'AWS_CRED',
    username        => 'access_key_ID',
    password        => 'secret_access_key'
  );
```

```
END;
/
```

This creates the `AWS_CRED` credential.

See CREATE_CREDENTIAL Procedure for more information.

2. Create a catalog.

   A catalog is a collection of wrapper functions that reference and call their respective cloud functions via their API endpoints.

   Example to create a catalog for AWS functions.

```
BEGIN
    DBMS_CLOUD_FUNCTION.CREATE_CATALOG (
        credential_name  => 'AWS_CRED',
        catalog_name     => 'AWS_DEMO_CATALOG',
        service_provider => 'AWS',
        cloud_params     => '{"region_id":"ap-northeast-1"}'
 );
END;
/

PL/SQL procedure successfully completed.
```

   This creates the `AWS_DEMO_CATALOG` catalog.

   See CREATE_CATALOG Procedure for more information. You can query DBA_CLOUD_FUNCTION_CATALOG View and USER_CLOUD_FUNCTION_CATALOG View views to retrieve the list of all the catalogs in your database.

3. You can list the functions in a catalog.

   Following is the example to list AWS Lambda functions:

```
VAR function_list CLOB;

BEGIN
 DBMS_CLOUD_FUNCTION.LIST_FUNCTIONS (
    credential_name  => 'AWS_CRED',
    catalog_name     => 'AWS_DEMO_CATALOG',
    function_list    => :function_list
);
END;
/

PL/SQL procedure successfully completed.

SELECT JSON_QUERY (:function_list, '$' RETURNING VARCHAR2(32676) pretty)
AS search_results FROM dual;

SEARCH_RESULTS                          --This is a sample output
--------------------------------------------------------------------------------
----------------------
[
  {
    "functionName"   : "test3_example",
```

```
    "functionArn"    : "arn:aws:lambda:ap-north-1:378:func:test3_example",
    "invokeEndpoint" : "https://swiy3.lambda-url.ap-
north-1.on.aws_example/"
  },
  {
    "functionName"   : "SumOfNum_example",
    "functionArn"    : "arn:aws:lambda:ap-
north-1:378:func:SumOfNum_example"

SEARCH_RESULTS
------------------------------------------------------------------------------
------------------------
    "invokeEndpoint" : "https://swiy3.lambda-url.ap-
north-1.on.aws_example/"
  },
  {
    "functionName"   : "testlambda_example",
    "functionArn"    : "arn:aws:lambda:ap-
north-1:378:func:testlambda_example",
    "invokeEndpoint" : "https://swiy3.lambda-url.ap-
north-1.on.aws_example/"
  },

SEARCH_RESULTS
------------------------------------------------------------------------------
------------------------
  {
    "functionName"   : "hellp-python_example",
    "functionArn"    : "arn:aws:lambda:ap-north-1:378:func:hellp-
python_example",
    "invokeEndpoint" : "https://swiy3.lambda-url.ap-
north-1.on.aws_example/"
  },
  {
    "functionName"   : "testlam_example",
    "functionArn"    : "arn:aws:lambda:ap-
north-1:378:func:testlam_example",

SEARCH_RESULTS
------------------------------------------------------------------------------
------------------------
    "invokeEndpoint" : "https://swiy3.lambda-url.ap-
north-1.on.aws_example/"
  }
]
```

See LIST_FUNCTIONS Procedure for more information.

4. Run the `DBMS_CLOUD_FUNCTION.SYNC_FUNCTIONS` to create wrapper SQL functions. You can use one of the following methods to create the wrapper SQL functions in the catalog, that call their respective cloud functions:

   - **SYNC_FUNCTIONS**: `SYNC_FUNCTIONS` is the quickest and simplest method, which automatically syncs (creates or deletes) wrapper functions in the catalog with the

complete list of cloud functions defined in the region, compartment, and tenancy with which the catalog was created. For example:

```
BEGIN
 DBMS_CLOUD_FUNCTION.SYNC_FUNCTIONS (
        catalog_name => 'AWS_DEMO_CATALOG'
 );
END;
/
PL/SQL procedure successfully completed.
```

This creates a PL/SQL wrapper for adding new functions to the catalog and removing wrappers for functions that have been deleted from the catalog.

Run the following query to verify the sync.

```
SELECT object_name FROM sys.all_objects WHERE owner='TEST_USER' AND
object_type='FUNCTION';


OBJECT_NAME
--------------------------------------------------------------------------------
-----
TESTLAMBDA
HELLP-PYTHON
TESTLAM
TEST3
SUMOFNUMBERS
```

> **Note:**
>
> Keep a note of the current user in order to run this command.

See SYNC_FUNCTIONS Procedure for more information.

- You can manually create a `SQL` Function in your catalog that calls its respective cloud function using `DBMS_CLOUD.CREATE_FUNCTION`.

Example to create a function in the `AWS_DEMO_CATALOG` catalog.

```
BEGIN
DBMS_CLOUD_FUNCTION.CREATE_FUNCTION (
  credential_name => 'AWS_CRED',
  catalog_name    => 'AWS_DEMO_CATALOG',
  function_name   => 'aws_testlambda',
  function_id     => 'arn:aws:lambda:ap-
northeast-1:378079562280:function:hellp-python'
);
END;
/
PL/SQL procedure successfully completed.
```

This creates the AWS_TESTLAMBDA function in the AWS_DEMO_CATALOG catalog.

The AWS_TESTLAMBDA function in the catalog is a reference to the respective cloud function whose endpoint is referenced by the FUNCTION_ID parameter. Invoking the function in the catalog along with its arguments runs the corresponding cloud function and provides the output returned by the cloud function.

You can query DBA_CLOUD_FUNCTION View and USER_CLOUD_FUNCTION View views to retrieve the list of all the functions in your database.

See CREATE_FUNCTION Procedure for more information.

5. After the function is created you can DESCRIBE and invoke it.

```
DESC AWS_TESTLAMBDA
COLUMN STATUS format a30
COLUMN OUTPUT format a30


SELECT AWS_TESTLAMBDA(NULL) FROM dual;

AWS_TESTLAMBDA(NULL)
----------------------------------------------------
{"STATUS":"200","RESPONSE_BODY":"Hello Python!!"}
```

This invokes the AWS_TESTLAMBDA function by calling the function reference arn:aws:lambda:ap-northeast-1:378079562280:function:hellp-python in the AWS_DEMO_CATALOG catalog.

6. You can drop an existing function using DROP_FUNCTION procedure. For example:

```
EXEC DBMS_CLOUD_FUNCTION.DROP_FUNCTION (CATALOG_NAME =>
'AWS_DEMO_CATALOG', FUNCTION_NAME => 'AWS_TESTLAMBDA');

PL/SQL procedure successfully completed.
```

This drops the AWS_TESTLAMBDA function from the AWS_DEMO_CATALOG catalog.

See DROP_FUNCTION Procedure for more information.

7. You can drop an existing catalog using DROP_CATALOG procedure. For example:

```
BEGIN
    DBMS_CLOUD_FUNCTION.DROP_CATALOG (
      catalog_name     => 'AWS_DEMO_CATALOG'
  );
END;
/

PL/SQL procedure successfully completed.
```

This drops the AWS_DEMO_CATALOG from your database.

See DROP_CATALOG Procedure for more information.

Steps to Invoke Azure Functions as SQL Functions

Shows the steps to invoke Azure remote functions as SQL functions in your database using HTTP trigger.

To invoke Azure Function as SQL functions, create a catalog of SQL wrapper functions that reference and call their respective cloud function through their API endpoints. Before you create this catalog, it is assumed here that you have created the necessary Azure functions to be referenced by this catalog.

1. To access Azure functions you need to use Azure Service Principal with Autonomous Database. You must grant the **Website Contributor** role to the Azure Service Principal for the Azure function app under its **Access control** (IAM).

   See the following for more information:

   - Use Azure Service Principal to Access Azure Resources
   - Azure built-in roles
   - Assign Azure roles using the Azure portal

2. Create a catalog.

   A catalog is a collection of wrapper functions that reference and call their respective cloud functions via their API endpoints.

   Example to create a catalog for Azure functions.

```
BEGIN
    DBMS_CLOUD_FUNCTION.CREATE_CATALOG (
        credential_name  => 'AZURE$PA',
        catalog_name     => 'AZURE_DEMO_CATALOG',
        service_provider => 'AZURE',
        cloud_params     =>
'{"subscription_id":"XXXXXXXXXXXXXXXXXXXXXXXX_example"}'
 );
END;
/
```

   The `SUBSCRIPTION_ID` value in the `CLOUD_PARAMS` is the subscription_id of the Azure function app.

   This creates the `AZURE_DEMO_CATALOG` catalog and catalogs all the functions of the provided `SUBSCRIPTION_ID`.

   See CREATE_CATALOG Procedure for more information. You can query DBA_CLOUD_FUNCTION_CATALOG View and USER_CLOUD_FUNCTION_CATALOG View views to retrieve the list of all the catalogs in your database.

3. You can list the functions in a catalog.

   Following is the example to list Azure functions:

```
VAR function_list CLOB;

BEGIN
 DBMS_CLOUD_FUNCTION.LIST_FUNCTIONS (
    credential_name  => 'AZURE$PA',
    catalog_name     => 'AZURE_DEMO_CATALOG',
```

```
    function_list    => :function_list
);
END;
/

SELECT JSON_QUERY (:function_list, '$' RETURNING VARCHAR2(32676) pretty)
AS search_results FROM dual;
```

See LIST_FUNCTIONS Procedure for more information.

4. Run the `DBMS_CLOUD_FUNCTION.SYNC_FUNCTIONS` to create wrapper SQL functions. You can use one of the following methods to create the wrapper SQL functions in the catalog, that call their respective cloud functions:

   - **SYNC_FUNCTIONS**: `SYNC_FUNCTIONS` is the quickest and simplest method, which automatically syncs (creates or deletes) wrapper functions in the catalog with the complete list of Azure functions. For example:

```
BEGIN
 DBMS_CLOUD_FUNCTION.SYNC_FUNCTIONS (
        catalog_name => 'AZURE_DEMO_CATALOG'
 );
END;
/
```

This creates a PL/SQL wrapper for adding new functions to the catalog and removing wrappers for functions that have been deleted from the catalog.

> **Note:**
>
> When an Azure function is added, deleted, or disabled on the Azure function app, you must run the `SYNC_FUNCTIONS` to update the corresponding PL/SQL wrappers in the catalog.

Run the following query to verify the sync.

```
SELECT object_name FROM sys.all_objects WHERE owner='TEST_USER' AND
object_type='FUNCTION';
```

> **Note:**
>
> Keep a note of the current user in order to run this command.

See SYNC_FUNCTIONS Procedure for more information.

   - You can manually create a `SQL` Function in your catalog that calls its respective Azure function using `AZURE_DEMO_CATALOG`.

Example to create a function in the AZURE_DEMO_CATALOG catalog.

```
BEGIN
DBMS_CLOUD_FUNCTION.CREATE_FUNCTION (
  credential_name => 'AZURE$PA',
  catalog_name    => 'AZURE_DEMO_CATALOG',
  function_name   => 'azure_testfunc',
  function_id     => 'function_id_path',
  input_args      => :function_args
);
END;
/
```

> **Note:**
>
> The maximum length of the function name is limited to 100 characters.

This creates the AZURE_TESTFUNC function in the AZURE_DEMO_CATALOG catalog.

The AZURE_TESTFUNC function in the catalog is a reference to the respective Azure function whose subscription is referenced by the FUNCTION_ID parameter. Invoking the function in the catalog along with its arguments runs the corresponding Azure function and provides the output returned by the function.

You can query DBA_CLOUD_FUNCTION View and USER_CLOUD_FUNCTION View views to retrieve the list of all the functions in your database.

See CREATE_FUNCTION Procedure for more information.

See Azure Functions HTTP trigger for more information.

5. After the function is created you can DESCRIBE and invoke it.

```
DESC AZURE_TESTFUNC


SELECT AZURE_TESTFUNC(NULL) FROM dual;
```

This invokes the AZURE_TESTFUNC function by calling the function reference /subscriptions/44496e556-8ssp-4262-b389-0f15f685c879/resources/ADBStest/providers/Microsoft.Web/sites/AZUREADBS/functions/HttpTrigger_example in the AZURE_DEMO_CATALOG catalog.

6. You can drop an existing function using DROP_FUNCTION procedure. For example:

```
EXEC DBMS_CLOUD_FUNCTION.DROP_FUNCTION (CATALOG_NAME =>
'AZURE_DEMO_CATALOG', FUNCTION_NAME => 'AZURE_TESTFUNC');
```

This drops the AZURE_TESTFUNC function from the AZURE_DEMO_CATALOG catalog.

See DROP_FUNCTION Procedure for more information.

7. You can drop an existing catalog using `DROP_CATALOG` procedure. For example:

```
BEGIN
    DBMS_CLOUD_FUNCTION.DROP_CATALOG (
      catalog_name     => 'AZURE_DEMO_CATALOG'
  );
END;
/
```

This drops the `AZURE_DEMO_CATALOG` from your database.

See DROP_CATALOG Procedure for more information.

Steps to Invoke Google Cloud Run Functions as SQL Functions

Shows the steps to invoke Google Cloud Run functions as SQL functions in your database.

To invoke Google Cloud Run functions as SQL functions, you create a library of SQL wrapper functions that reference and call their respective cloud functions. Before you create this catalog, it is assumed here that you have created the necessary cloud functions to be referenced by this catalog. See Create a Cloud Run function by using the Google Cloud console for more information.

1. To access Google Cloud Run functions from an Autonomous Database instance use your Google service account. You must grant the `cloudfunctions.functions.list` permission, and `Cloud Functions Invoker` and `Cloud Run Invoker` roles to the Google service account for the Google Cloud Platform (GCP) app.

   See the following for more information:

   - About Using a Google Service Account to Access Google Cloud Resources
   - Authorize access with IAM (1st gen)
   - Authorize access with IAM

2. Create a catalog. To create a catalog for Google Cloud Run functions with a user other than `ADMIN` you need to grant read privileges on the `CLOUD_INTEGRATIONS` view to that user.

   A catalog is a collection of wrapper functions that reference and call their respective cloud functions using their API endpoints.

```
BEGIN
    DBMS_CLOUD_FUNCTION.CREATE_CATALOG (
        credential_name  => 'GCP$PA',
        catalog_name     => 'GCP_DEMO_CATALOG',
        service_provider => 'GCP',
        cloud_params     => '{"project_id":"example_XXXXXX"}'
 );
END;
/
```

The `project_id` value in the `cloud_params` is the project ID of the Google Cloud Platform (GCP).

This creates the `GCP_DEMO_CATALOG` catalog and catalogs all the functions of the provided `project_id`.

See CREATE_CATALOG Procedure for more information. You can query DBA_CLOUD_FUNCTION_CATALOG View and USER_CLOUD_FUNCTION_CATALOG View views to retrieve the list of all the catalogs in your database.

3. List the functions in a catalog.

   Following is the example to list Google Cloud Run functions:

```
VAR function_list CLOB;

BEGIN
    DBMS_CLOUD_FUNCTION.LIST_FUNCTIONS (
        credential_name  => 'GCP$PA',
        catalog_name     => 'GCP_DEMO_CATALOG',
        function_list    => :function_list
 );
END;
/

SELECT JSON_QUERY (:function_list, '$' RETURNING VARCHAR2(32676) pretty)
AS search_results FROM dual;
```

   See LIST_FUNCTIONS Procedure for more information.

4. Run the `DBMS_CLOUD_FUNCTION.SYNC_FUNCTIONS` to create wrapper SQL functions. You can use one of the following methods to create the wrapper SQL functions in the catalog, that call their respective cloud functions:

   - **SYNC_FUNCTIONS**: `SYNC_FUNCTIONS` method to create wrapper SQL function is the quickest and simplest method, which automatically syncs (creates or deletes) wrapper functions in the catalog with the complete list of cloud functions defined in the region, compartment, and tenancy with which the catalog was created. For example:

```
BEGIN
    DBMS_CLOUD_FUNCTION.SYNC_FUNCTIONS (
        catalog_name => 'GCP_DEMO_CATALOG'
 );
END;
/
```

   This creates a PL/SQL wrapper for adding new functions to the catalog and removing wrappers for functions that have been deleted from the catalog.

   Run the following query to verify the sync.

```
SELECT object_name FROM sys.all_objects WHERE owner='TEST_USER' AND
object_type='FUNCTION';
```

> **Note:**
>
> Keep a note of the current user in order to run this command.

See SYNC_FUNCTIONS Procedure for more information.

- You can manually create a SQL Function in your catalog that calls its respective cloud function using DBMS_CLOUD.CREATE_FUNCTION.

Example to create a function in the GCP_DEMO_CATALOG catalog.

```
EXEC :function_args  := TO_CLOB('{"name": "VARCHAR2"}');


BEGIN
    DBMS_CLOUD_FUNCTION.CREATE_FUNCTION (
        credential_name => 'GCP$PA',
        catalog_name    => 'GCP_DEMO_CATALOG',
        function_name   => 'gcp_testfunc',
        function_id     => 'function_id_path',
        input_args      => :function_args
 );
END;
/
```

This creates the GCP_TESTFUNC function in the GCP_DEMO_CATALOG catalog.

The GCP_TESTFUNC function in the catalog is a reference to the respective Google Cloud Run function whose project id is referenced by the FUNCTION_ID parameter. Invoking the function in the catalog along with its arguments runs the corresponding Google Cloud Run function and provides the output returned by the function.

You can query DBA_CLOUD_FUNCTION View and USER_CLOUD_FUNCTION View views to retrieve the list of all the functions in your database.

See CREATE_FUNCTION Procedure for more information.

5. After the function is created you can DESCRIBE and invoke it.

```
DESC GCP_TESTFUNC
COLUMN STATUS format a30
COLUMN OUTPUT format a30


SELECT GCP_TESTFUNC(NULL) FROM dual;
```

This invokes the GCP_TESTFUNC function by calling the function reference specified in the function_id parameter.

6. You can drop an existing function using DROP_FUNCTION procedure. For example:

```
EXEC DBMS_CLOUD_FUNCTION.DROP_FUNCTION (CATALOG_NAME =>
'GCP_DEMO_CATALOG', FUNCTION_NAME => 'GCP_TESTFUNC');
```

This drops the GCP_TESTFUNC function from the GCP_DEMO_CATALOG catalog.

See DROP_FUNCTION Procedure for more information.

7. You can drop an existing catalog using DROP_CATALOG procedure. For example:

```
BEGIN
    DBMS_CLOUD_FUNCTION.DROP_CATALOG (
```

```
        catalog_name     => 'GCP_DEMO_CATALOG'
 );
END;
/
```

This drops the `GCP_DEMO_CATALOG` from your database.

See DROP_CATALOG Procedure for more information.

Invoke External Procedures as SQL Functions

Shows the steps to invoke external procedures using PL/SQL within your database.

- External Procedures Overview
  External procedures are functions written in a third-generation language and callable from within PL/SQL or SQL as if they were a PL/SQL routine or function.
- About Using External Procedures in Autonomous Database
  You can invoke and use external procedures in your Autonomous Database with user defined functions.
- Define the C Procedure
  Define the C procedure using one of these prototypes.
- Create a Shared Library (.so) File
  Create a shared object (`.so` file) library. The shared object library contains the C procedure (external procedure) which was defined in the previous step.
- Get the OCI Marketplace EXTPROC Stack Application
  Shows the steps to get the OCI Marketplace `EXTPROC` Stack Application.
- Launch EXTPROC Stack Application
  Launch the `EXTPROC` Stack Application from the EXTPROC Application Details page.
- Create Stack for EXTPROC Agent Application
  Shows the steps to create Stack for `EXTPROC` instance.
- Upload Wallet to Create Secure Connection to the EXTPROC Agent Instance
  A self-signed wallet is created as part of the `EXTPROC` agent application creation. This wallet allows you to access the `Extrpoc` agent instance.
- Steps to Invoke an External Procedure as a SQL Function
  Shows the steps to invoke an external Procedure as a SQL function.

External Procedures Overview

External procedures are functions written in a third-generation language and callable from within PL/SQL or SQL as if they were a PL/SQL routine or function.

External procedures promote reusability, efficiency, and modularity. Existing dynamic link libraries (DLLs) written in other languages can be called from PL/SQL programs. The DLLs are loaded only when needed and they can be enhanced without affecting the calling programs.

Using external procedures also enhances performance, because third-generation languages perform certain tasks more efficiently than PL/SQL, which is better suited for SQL transaction processing.

External procedures are useful when:

- Solving scientific and engineering problems

- Analyzing data
- Controlling real-time devices and processes

See What Is an External Procedure? for more information.

About Using External Procedures in Autonomous Database

You can invoke and use external procedures in your Autonomous Database with user defined functions.

You do not install external procedures on an Autonomous Database instance. To use an external procedure, the procedure is hosted remotely on a VM running in an Oracle Cloud Infrastructure Virtual Cloud Network (VCN).

External procedures are only supported when your Autonomous Database is on a private endpoint. The `EXTPROC` agent instance is hosted on a private subnet and the Autonomous Database access the `EXTPROC` agent through a Reverse Connection Endpoint (RCE).

> **Note:**
>
> Autonomous Database only supports C language external procedures.

External procedures are deployed by using:

- An Oracle provided container image with `EXTPROC` agent installed and configured as a part of the Oracle Cloud Infrastructure (OCI) Marketplace stack.

  The `EXTPROC` agent instance is hosted remotely on a VM running in an Oracle Cloud Infrastructure Virtual Cloud Network (VCN). The secure communication between your Autonomous Database and the `EXTPROC` agent instance is ensured by setting Network Security Group (NSG) rules such that the traffic is allowed from your Autonomous Database instance running on a private endpoint to the `EXTPROC` agent instance.

  The `EXTPROC` agent image is pre-configured to host and execute external procedures on port 16000.
- PL/SQL procedures to create a library and to register and invoke external functions and procedures.

  See DBMS_CLOUD_FUNCTION Package for more information.

Follow these steps to invoke an external procedure on Autonomous Database:

- Define the C procedure. See Define the C Procedure.
- Create a shared object (`.so` file) library. See Create a Shared Library (.so) File.
- Launch the Autonomous Database `EXTPROC` stack application. See Get the OCI Marketplace EXTPROC Stack Application.
- Provision and configure Oracle Autonomous Database `EXTPROC` agent. See Create Stack for EXTPROC Agent Application for more information.
- Configure your Autonomous Database to connect to the `EXTPROC` agent instance. See Upload Wallet to Create Secure Connection to the EXTPROC Agent Instance for more information.
- Create a remote library using `DBMS_CLOUD_FUNCTION.CREATE_CATALOG`. See Steps to Invoke an External Procedure as a SQL Function for more information.

- Use the user defined function you created in the previous step. See Steps to Invoke an External Procedure as a SQL Function for more information.

Define the C Procedure

Define the C procedure using one of these prototypes.

- Kernighan & Ritchie style prototypes. For example:

```
void UpdateSalary(x)
 float x;
...
```

- ISO/ANSI prototypes other than numeric data types that are less than full width (such as `float`, `short`, `char`). For example:

```
void UpdateSalary(double x)
...
```

- Other data types that do not change size under default argument promotions.

  This example changes size under default argument promotions:

```
void UpdateSalary(float x)
...
```

Create a Shared Library (.so) File

Create a shared object (`.so` file) library. The shared object library contains the C procedure (external procedure) which was defined in the previous step.

You generate a shared object library using the following command:

```
gcc -I/u01/app/oracle/extproc_libs/ -shared -fPIC -o extproc.so UpdateSalary.c
```

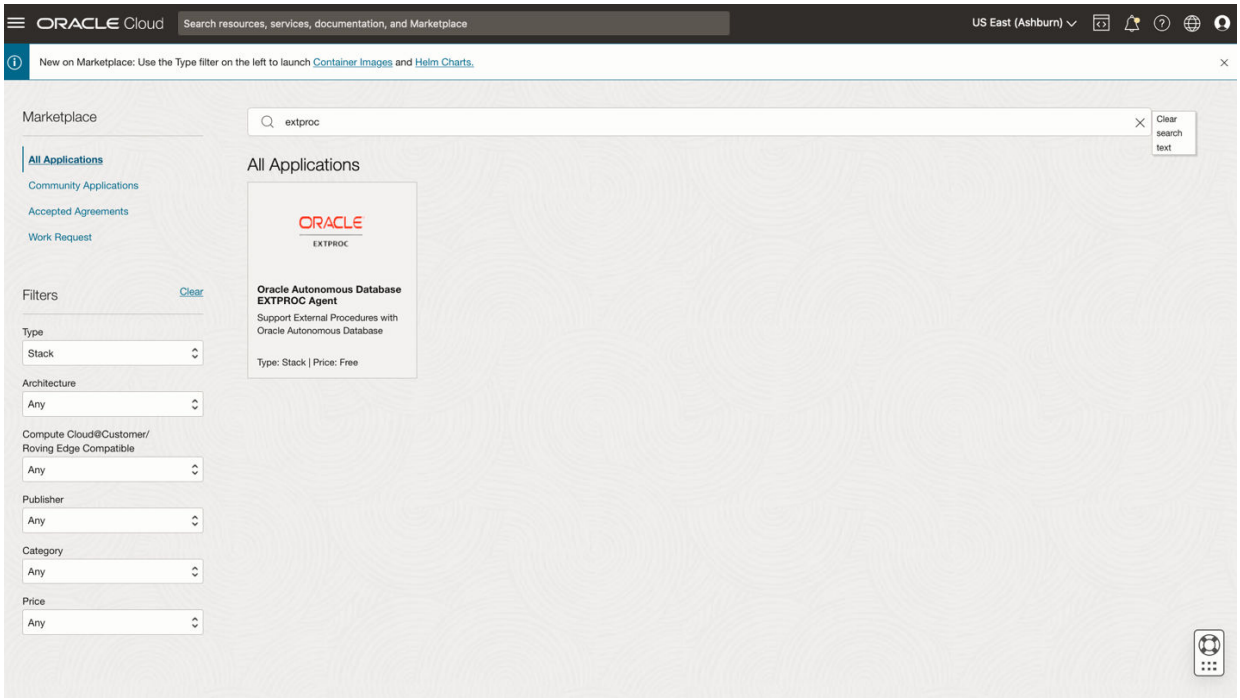
This creates the shared object (`.so`), `extproc.so` library. The `UpdateSalary` procedure, defined in the previous step, is contained in the `extproc.so` library. The shared object (`.so`) libraries are dynamically loaded at run time.

Get the OCI Marketplace EXTPROC Stack Application

Shows the steps to get the OCI Marketplace `EXTPROC` Stack Application.

Perform the following steps:

1. Sign in to the OCI Console at http://cloud.oracle.com. See Sign in to the Oracle Cloud Infrastructure Console for more information.
2. From the Oracle Cloud Infrastructure left navigation menu click **Marketplace** and then, under **Marketplace** click **All Applications**. This takes you to the Marketplace All Applications dashboard.
3. Enter `EXTPROC` in the search field and click **search**.
4. Click the `EXTPROC` widget of **Type: Stack**.



This takes you to the **Oracle Autonomous Database EXTPROC Agent** details page.

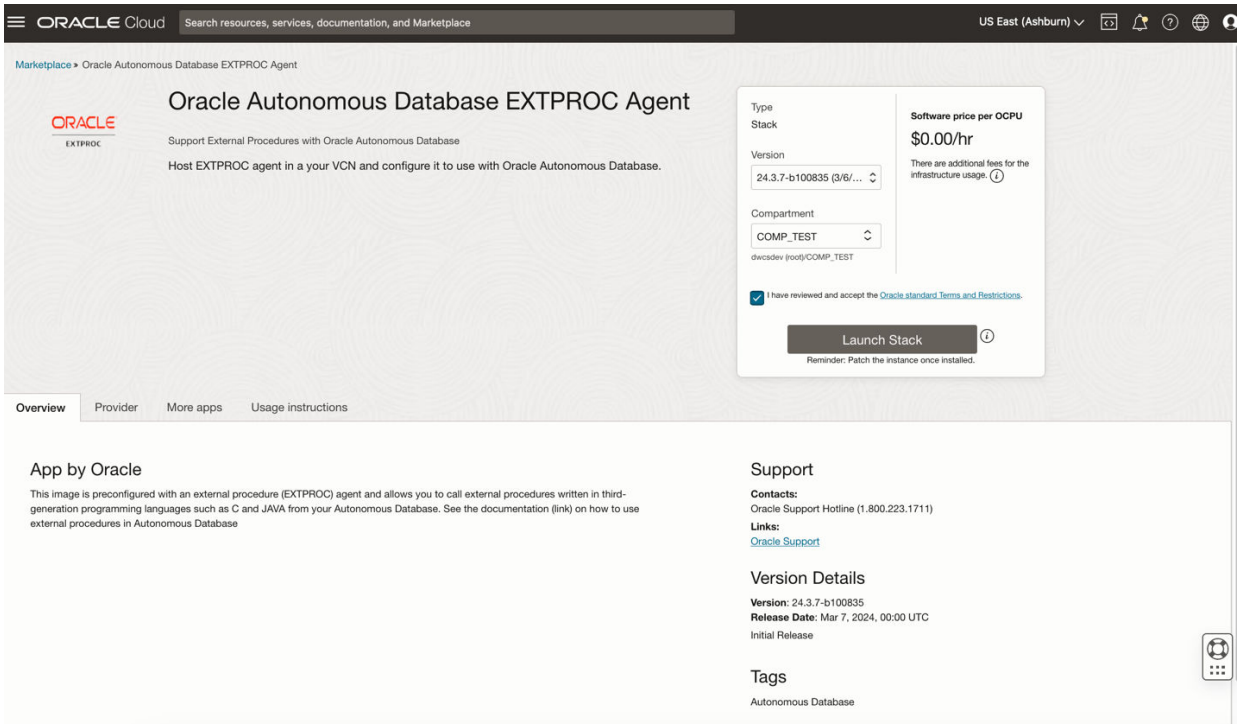
Launch EXTPROC Stack Application

Launch the `EXTPROC` Stack Application from the EXTPROC Application Details page.

1. On the Oracle Autonomous Database `EXTPROC` Agent page, under **Type Stack**, perform the following:
   - From the **Version** drop-down list, select the package version of the stack. By default, the menu displays the latest version.
   - From the **Compartment** drop-down list, select the name of the compartment where you want to launch the instance.

     > **Note:**
     >
     > If you don't have permission to launch the instance in the selected compartment, the instance is launched in the root compartment.

   - Select the **I have reviewed and accept the Oracle Standard Terms and Restrictions** checkbox.
2. Click **Launch Stack**.



This takes you to the **Create stack** page that allows you to create stack for the `EXTPROC` agent.

Create Stack for EXTPROC Agent Application

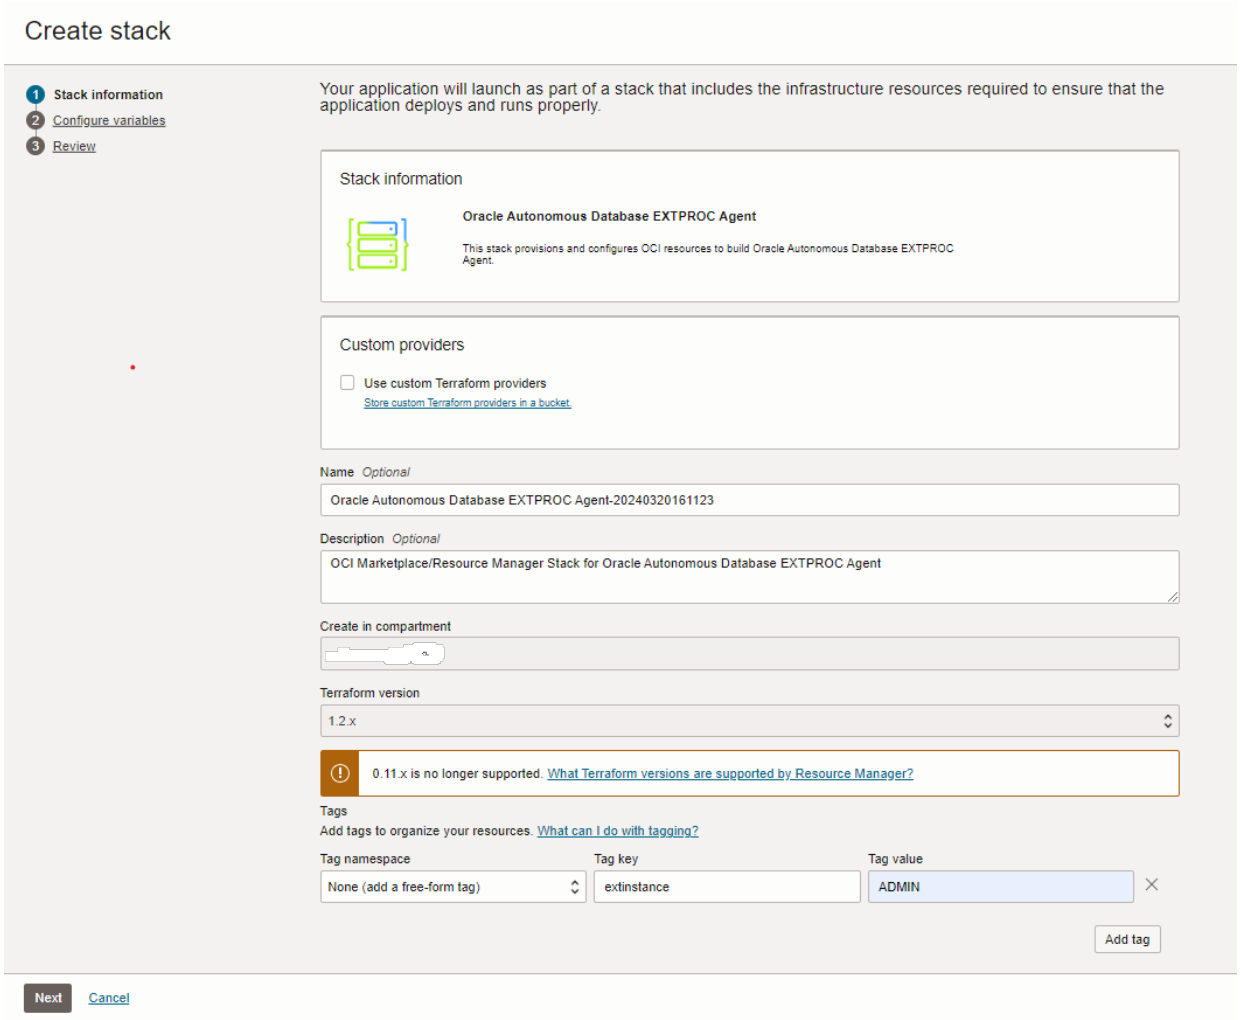
Shows the steps to create Stack for `EXTPROC` instance.

In the **Create stack** wizard, perform the following steps:

1. On the **Stack information** page, review and edit the following information as necessary:
   - **Stack information**
   - **Custom providers**
   - **Name (Optional)**: You can edit the default stack name. Avoid entering confidential information.
   - **Description (Optional)**: You can edit the default stack description. Avoid entering confidential information.
   - **Create in compartment**
   - **Terraform version**
   - **Tags**: Provide the following to assign tags to the stack.
     - **Tag namespace**: To add a defined tag, select an existing namespace. To add a free-from tag, leave the value blank.
     - **Tag key**: To add a defined tag, select an existing tag key. To add a free-form tag, type the key name that you want.
     - **Tag value**: Type the tag value that you want.

     Add tag: Click to add another tag.

     See Resource Tags for more information on tagging.



2. Click **Next**.

   This takes you to the **Configure variables** page which enables you to configure variables for the infrastructure resources that the stack creates when you run the apply job for this execution plan.

3. On the Configure variables page enter the information in the areas: **Configure the EXTPROC Agent**, **Network Configuration**, and **Compute configuration**.

   a. Provide information in the **Configure the EXTPROC Agent** area.

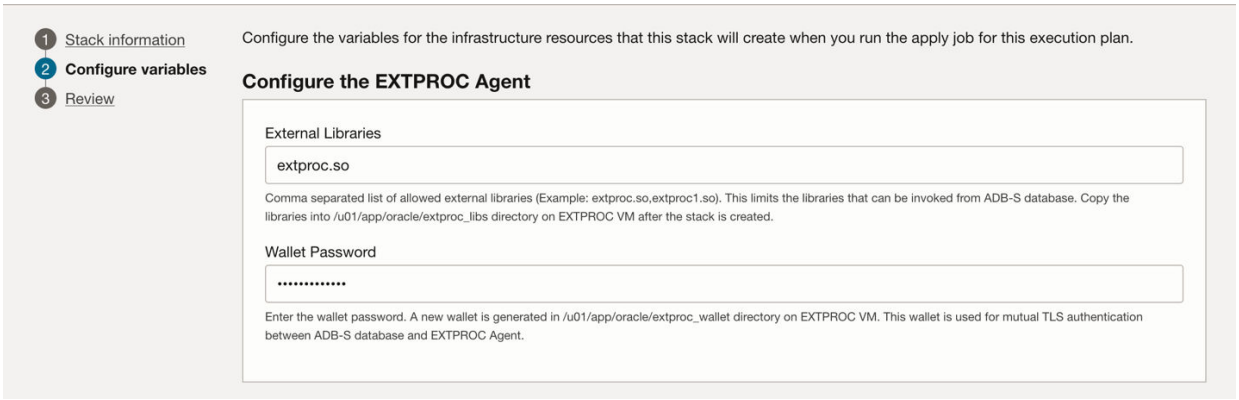
      - **External Libraries**: Provide a list of libraries, separated by comma (,), which you want to allow to be invoked from your Autonomous Database. For example, `extproc.so`, `extproc1.so`.

        After you create the stack, you must copy the libraries to the `/u01/app/oracle/extproc_libs` directory on the `EXTPROC` agent VM.

      - **Wallet Password:** Provide the wallet password.

        The wallet and a self-signed certificate is generated for mutual TLS authentication between the Autonomous Database and the `EXTPROC` agent VM. The wallet is created in the `/u01/app/oracle/extproc_wallet` directory.

> **Note:**
>
> After the wallet is created, the wallet password cannot be changed.

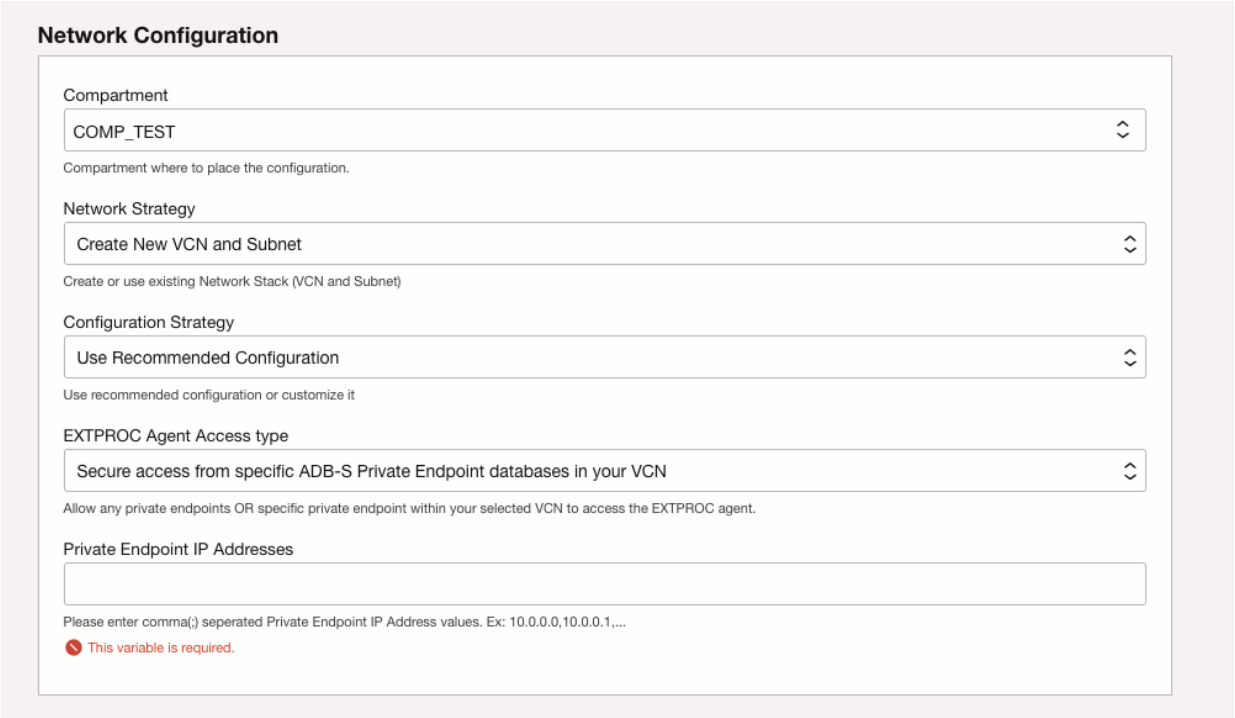


b. Provide information in the **Network Configuration** area.

- **Compartment**: From the drop-down list, choose the compartment where you want to place the configuration.
- **Network Strategy**: Choose one of the options from the drop-down list, **Create New VCN and Subnet** or **Use Existing VCN and Subnet**.
  - **Create New VCN and Subnet**: Choose this option if a private endpoint is not configured for your Autonomous Database. This creates a new VCN with public and private subnet that are preconfigured with security rules.

    If you select this option the page also shows the Configuration Strategy drop-down list:

    Choose **Use Recommended Configuration** from the **Configuration Strategy** drop-down list.



- **Use Existing VCN and Subnet**: Select this option to create the EXTPROC agent using an existing VCN. This creates the EXTPROC agent instance in the provided subnet.

  When you select this option, provide the following information for the existing VCN and Subnet:

  * Under **Virtual Cloud Network**:

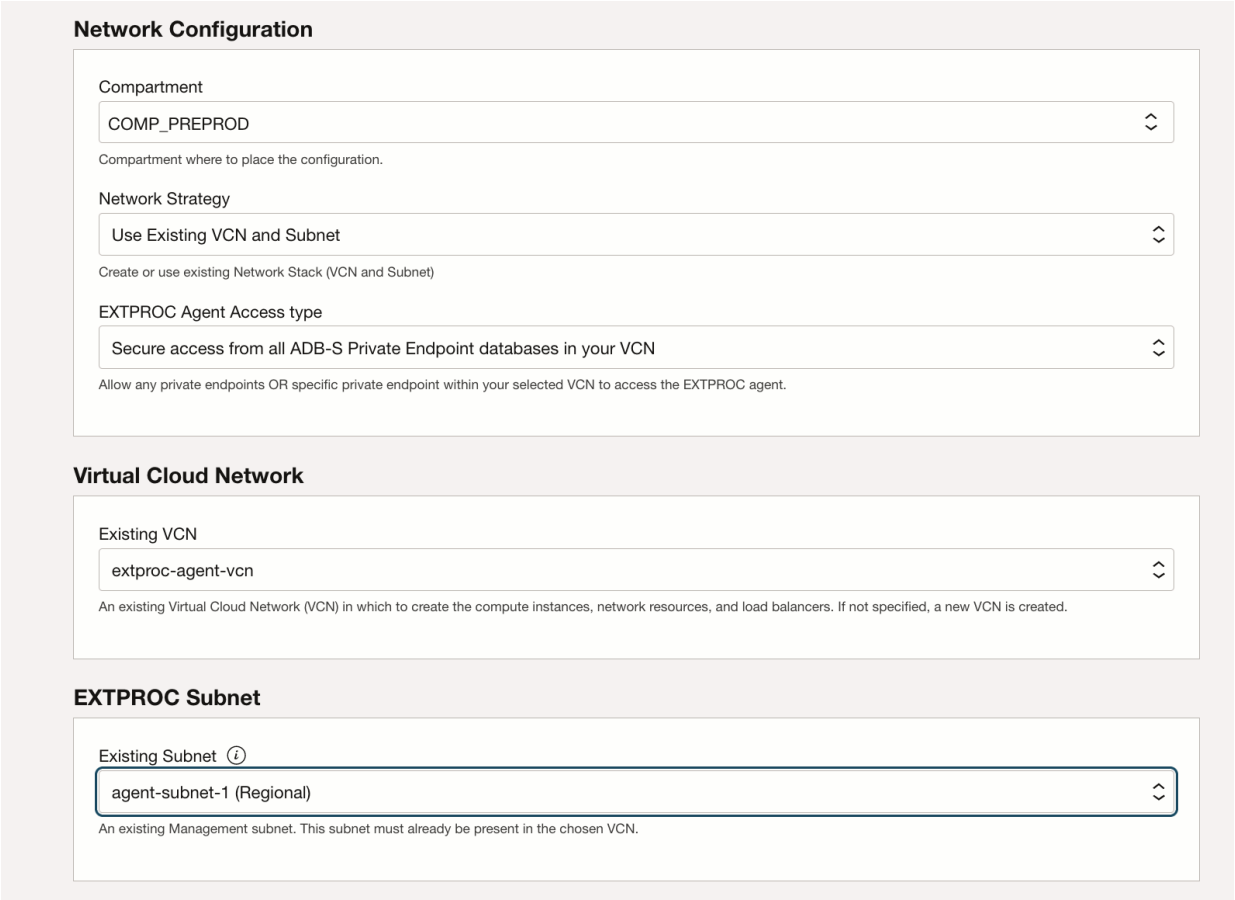
    From the **Existing VCN** drop-down list choose an existing VCN. If the specified VCN does not exist, a new VCN is created.

  * Under **EXTPROC Subnet**:

    From the **Existing Subnet** drop-down list choose an existing subnet.

    When you choose to use an existing VCN and subnet, add an ingress rule for the `EXTPROC` agent instance's port `16000`. You also add an egress rule on public subnet.

    See Configuring Network Access with Private Endpoints for more information.



- **EXTPROC Agent Access type**: Choose one of the following options from the drop-down list.
  - **Secure access from specific ADB-S Private Endpoint databases in your VCN**: Choose this option to allow only specified private endpoint IPs inside your Virtual Cloud Network (VCN) to connect to your `EXTPROC` agent.

    When this option is chosen, you provide a list of allowed private endpoint IP Addresses in the next step.
  - **Secure access from all ADB-S Private Endpoint databases in your VCN**: Choose this option to allow any private endpoint inside your Virtual Cloud Network (VCN) to connect to your `EXTPROC` agent.
- **Private Endpoint IP Addresses**

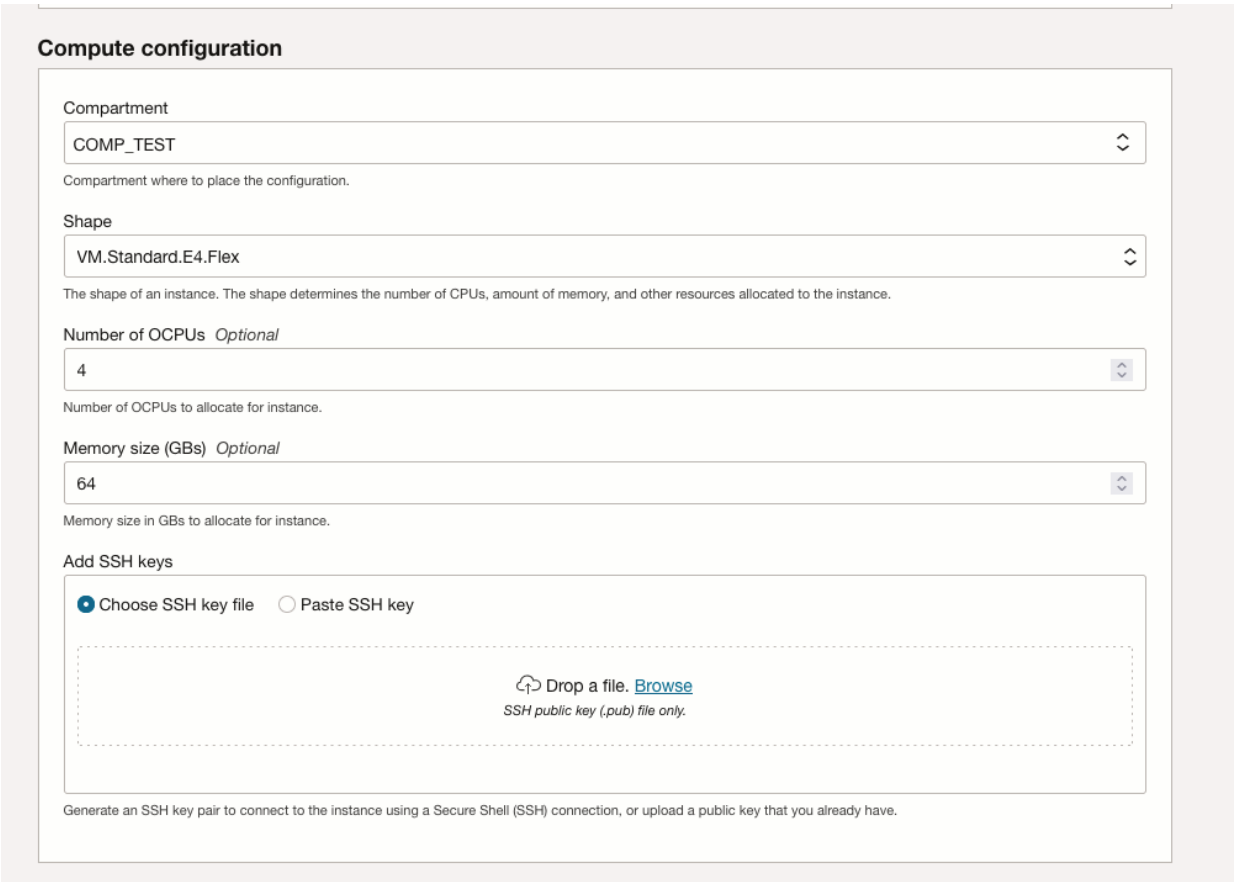
  Provide a list of private endpoint IP addresses separated by comma (,) for the Private Endpoint IP Addresses variable. For example, `10.0.0.0`, `10.0.0.1`.

> **Note:**
>
> This field only shows when you select **Secure access from specific ADB-S Private Endpoint databases in your VCN** for the **EXTPROC Agent Access type**.

c. Provide the **Compute configuration** information.

- **Compartment**: Select the compartment where you want to create the stack.

- **Shape**: Select a shape based on the workload requirements of the `EXTPROC` agent instance. The shape determines the resources allocated to the `EXTPROC` agent instance.
- **Number of OCPUs**: Choose the number of OCPUs that you want to allocate to the `EXTPROC` agent instance.
- **Memory size (GBs)**: Choose the amount of memory in Gigabytes (GB) that you want to allocate to the `EXTPROC` agent instance.
- **Add SSH keys**: Upload an SSH public key or paste the public key. Select one of the following options:
  - **Choose SSH key file**: Upload the public key portion of your key pair. Either browse to the key file that you want to upload, or drag and drop the file into the box.
  - **Paste SSH key**: Paste the public key portion of your key pair in the box.

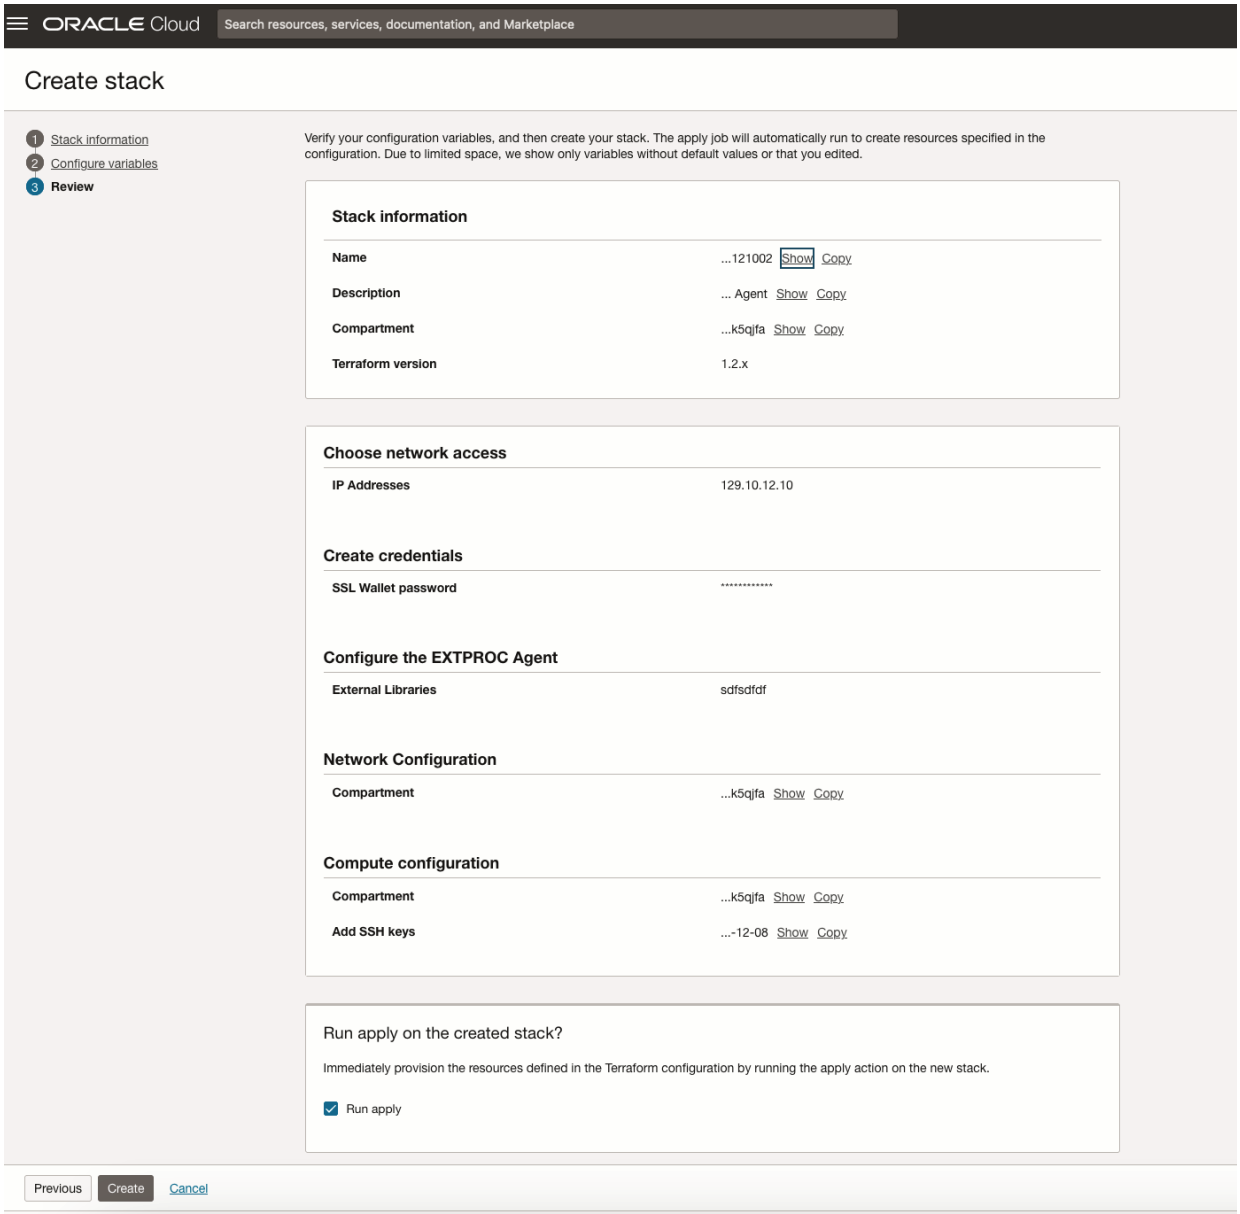


4. Click **Next**.

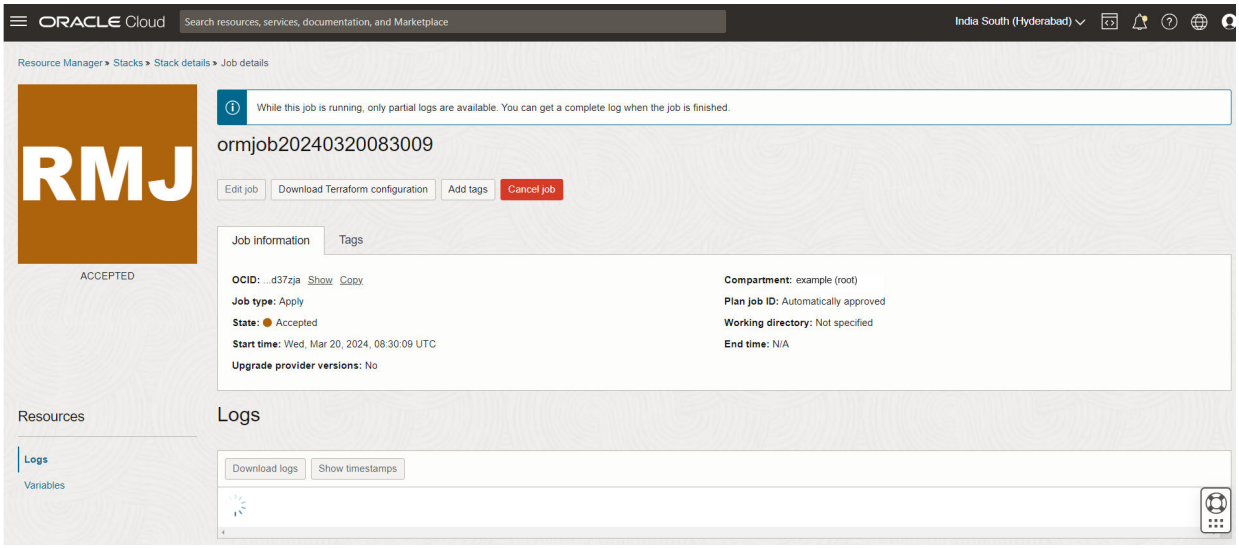
   This takes you to the **Review** page.

5. On the **Review** page perform the following steps:
   a. Verify the configuration variables.
   b. Select the **Run apply** checkbox under **Run apply on the created stack?**
   c. Click **Create**.

> **Note:**
>
> This area does not show variables that have default values or variables that you have not changed.



Resource Manager runs the apply job to create stack resources accordingly. This takes you to the **Job details** page and the job state is **Accepted**. When the apply job starts the status is updated to **In Progress**.



> **Note:**
>
> The information you need to connect to the instance created as part of the stack is available in the **Application Information** tab.

Upload Wallet to Create Secure Connection to the EXTPROC Agent Instance

A self-signed wallet is created as part of the `EXTPROC` agent application creation. This wallet allows you to access the `Extrpoc` agent instance.

To execute remote procedures at the `EXTPROC` agent instance, the Autonomous Database and the `EXTPROC` agent connect using Mutual Transport Layer Security (mTLS). When using Mutual Transport Layer Security (mTLS), clients connect through a TCPS (Secure TCP) database connection using standard TLS 1.2 with a trusted client certificate authority (CA) certificate. See About Connecting to an Autonomous Database Instance for more information.

> **Note:**
>
> You can also obtain and use a public certificate issued by a Certificate Authority (CA).

As a prerequisite, you must export the wallet to Object Storage from the `/u01/app/oracle/extproc_wallet` directory on the VM where `EXTPROC` runs.

Follow these steps to upload the wallet to your Autonomous Database:

1. Import the wallet, `cwallet.sso`, containing the certificates for the `EXTPROC` agent instance from Object Storage in your Autonomous Database. Note the following for the wallet file:
   - The wallet file, along with the Database user ID and password provide access to the `EXTPROC` agent instance. Store wallet files in a secure location and share them only with authorized users.
   - Do not rename the wallet file. The wallet file in Object Storage must be named `cwallet.sso`.

2. Create credentials to access your Object Storage where you store the wallet file `cwallet.sso`. See CREATE_CREDENTIAL Procedure for information about the username and password parameters for different object storage services.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See About Using Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

3. Create a directory on Autonomous Database for the wallet file `cwallet.sso`.

   ```
   CREATE DIRECTORY wallet_dir AS 'directory_location';
   ```

   See Create Directory in Autonomous Database for more information creating directories.

4. Use `DBMS_CLOUD.GET_OBJECT` to upload the wallet. For example:

   ```
   BEGIN
     DBMS_CLOUD.GET_OBJECT (
       credential_name     => 'DEF_CRED_NAME',
       object_uri          => 'https://objectstorage.us-
   phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/cwallet.sso',
       directory_name      => 'WALLET_DIR'
   );
   END;
   /
   ```

   In this example, `namespace-string` is the Oracle Cloud Infrastructure Object Storage namespace and `bucketname` is the bucket name. See Object Storage Namespaces for more information.

   The wallet is copied to the directory created in the previous step, *WALLET_DIR*. The wallet that allows you to connect to the `EXTPROC` agent instance is now available on your Autonomous Database instance.

Steps to Invoke an External Procedure as a SQL Function

Shows the steps to invoke an external Procedure as a SQL function.

After you launch the OCI Marketplace `EXTPROC` stack application and configure it to run external procedures, you create a library of SQL wrapper functions that reference and call their respective external procedures.

As a prerequisite, the whitelisted libraries must be copied into the `/u01/app/oracle/extproc_libs` directory on the `EXTPROC` VM.

Follow these steps to create a library in your Autonomous Database and register C routines as an external procedure in the library:

1. Create a library.

   An external procedure is a C language routine stored in a library. To invoke external procedures with Autonomous Database, you create a library.

   Run `DBMS_CLOUD_FUNCTION.CREATE_CATALOG` to create a library. For example:

   ```
   BEGIN
       DBMS_CLOUD_FUNCTION.CREATE_CATALOG (
           library_name               => 'demolib',
   ```

```
        library_listener_url       => 'remote_extproc_hostname:16000',
        library_wallet_dir_name    => 'wallet_dir',
        library_ssl_server_cert_dn => 'CN=VM Hostname',
        library_remote_path        => '/u01/app/oracle/extproc_libs/
library name'
);
END;
/
```

This creates the `demolib` library in your Autonomous Database and registers the dynamic link library in your database. The `EXTPROC` agent instance is pre-configured to host external procedures on port 16000.

See CREATE_CATALOG Procedure for more information.

Query DBA_CLOUD_FUNCTION_CATALOG View and USER_CLOUD_FUNCTION_CATALOG View views to retrieve the list of all the catalogs and libraries in your database.

Query the USER_CLOUD_FUNCTION_ERRORS View view to list any errors generated during the connection validation to the remote library location.

2. Create the function.

   For example:

```
CREATE OR REPLACE FUNCTION ftest(
      x VARCHAR2,
      y VARCHAR2)
RETURN VARCHAR2 AS LANGUAGE C
      LIBRARY test
      NAME "demolib"
      PARAMETERS(
          x STRING,
          y STRING)
      AGENT IN (x);
/
```

3. You can drop an existing library using `DROP_CATALOG` procedure. For example:

```
BEGIN
    DBMS_CLOUD_FUNCTION.DROP_CATALOG (
      LIBRARY_NAME  => 'demolib'
  );
END;
/
```

   This drops the `DEMOLIB` library.

   See DROP_CATALOG Procedure for more information.

Make External Calls Using a Customer-Managed Wallet

When your Autonomous Database instance is on a private endpoint you can use a customer-managed wallet with procedures in `UTL_HTTP`, `DBMS_LDAP`, `UTL_SMTP`, or `UTL_TCP`. You can also use a customer-managed wallet when the scheduler sends SMTP email notifications for various events related to scheduler jobs.

- About Using a Customer-Managed Wallet with External Calls
  When your Autonomous Database instance is on a private endpoint you can use a customer-managed wallet to handle external calls, or with the scheduler when the scheduler sends email for various events related to scheduler jobs.
- Prerequisites to Use a Customer-Managed Wallet with External Calls
  Shows the prerequisite steps to use a customer-managed wallet with external calls or with scheduler SMTP email notifications.
- Use a Customer-Managed Wallet for External Calls with UTL_HTTP
  When your Autonomous Database instance is on a private endpoint you can use a customer-managed wallet to handle external calls.
- Use a Customer-Managed Wallet with Scheduler Email Notifications
  Describes the steps to use the scheduler SMTP email server with a customer-managed wallet.
- Notes for Using a Customer-Managed Wallet with External Calls
  Provides notes for using a customer-managed wallet with external calls.

About Using a Customer-Managed Wallet with External Calls

When your Autonomous Database instance is on a private endpoint you can use a customer-managed wallet to handle external calls, or with the scheduler when the scheduler sends email for various events related to scheduler jobs.

On Autonomous Database you can make external calls for any of the of the following purposes:

- To use web services with `UTL_HTTP`.
- To access data from LDAP servers using `DBMS_LDAP`.
- To send email with `UTL_SMTP`.
- To communicate with external TCP/IP-based servers using TCP/IP with `UTL_TCP`.
- For Oracle Scheduler job email notifications.

By default when you use procedures in these packages Autonomous Database maintains an internal wallet and always uses secure connections (the Oracle managed wallet contains more than 90 of the most common trusted root and intermediate SSL certificates). When your Autonomous Database resides on a private endpoint you have the option to use the default Oracle managed wallet with the trusted and intermediate SLL certifications, or you can supply a customer-managed wallet.

When your Autonomous Database resides on a private endpoint you can specify a customer-managed wallet for `UTL_HTTP`, `UTL_SMTP`, `DBMS_LDAP`, and `DBMS_NETWORK_ACL_ADMIN` using the following PL/SQL procedures:

```
UTL_HTTP.SET_WALLET (
   path                 IN VARCHAR2,
   password             IN VARCHAR2 DEFAULT NULL);
UTL_HTTP.REQUEST (
   wallet_path          IN VARCHAR2 DEFAULT NULL,
   wallet_password      IN VARCHAR2 DEFAULT NULL)
UTL_HTTP.REQUEST_PIECES (
   wallet_p ath         IN VARCHAR2 DEFAULT NULL,
   wallet_password      IN VARCHAR2 DEFAULT NULL,
UTL_HTTP.CREATE_REQUEST_CONTEXT (
 wallet_path            IN VARCHAR2 DEFAULT NULL,
```

```
 wallet_password        IN VARCHAR2 DEFAULT NULL)
UTL_TCP.OPEN_CONNECTION
   wallet_path          IN  VARCHAR2 DEFAULT NULL,
   wallet_password      IN  VARCHAR2 DEFAULT NULL);
UTL_SMTP.OPEN_CONNECTION
   wallet_path          IN  VARCHAR2 DEFAULT NULL,
   wallet_password      IN  VARCHAR2 DEFAULT NULL)
DBMS_LDAP.OPEN_SSL(
  sslwrl                IN VARCHAR2,
  sslwalletpasswd       IN VARCHAR2)
DBMS_NETWORK_ACL_ADMIN.APPEND_WALLET_ACE
  wallet_path IN VARCHAR2
DBMS_NETWORK_ACL_ADMIN.APPEND_WALLET_ACL(
  wallet_path IN VARCHAR2
```

In these calls, depending on the procedure, use either the `path` or `wallet_path` parameter to specify a customer-managed wallet. You set the wallet directory with the `DIR:` prefix and include a wallet path. For example:

```
UTL_HTTP.set_wallet('DIR:WALLET_DIR', 'password');
```

The current user invoking the API must have `READ` privileges on the directory object or be granted the `CREATE ANY DIRECTORY` system privilege.

The `DIR:` prefix is the preferred form to use to specify a customer-managed wallet. In addition, the `file:` prefix is supported. For example:

```
UTL_HTTP.set_wallet('file:WALLET_DIR/wallet.sso', 'password');
```

See Use a Customer-Managed Wallet for External Calls with UTL_HTTP for more information:

About Using a Customer-Managed Wallet with Scheduler Email Server

The scheduler email server is available to send email notifications for various events related to the scheduler, such as notification of started, failed, or completed jobs. By default, the scheduler SMTP email server uses the wallet defined in the `SSL_WALLET` property for SSL/TLS communications. Optionally, you can use a customer-managed wallet with the scheduler SMTP email server.

The following global attributes support using a customer-managed wallet:

- `EMAIL_SERVER_WALLET_DIRECTORY`: is set to a directory object that specifies the path where the SSL wallet resides.
- `EMAIL_SERVER_WALLET_CREDENTIAL`: is set to a credential object with a *username*/*password* pair, where the *username* is any value and the *password* is the SSL wallet password.

You set values for these attributes using `DBMS_SCHEDULER.SET_SCHEDULER_ATTRIBUTE`.

See the following for more information:

- Use a Customer-Managed Wallet with Scheduler Email Notifications
- SET_SCHEDULER_ATTRIBUTE Procedure

Prerequisites to Use a Customer-Managed Wallet with External Calls

Shows the prerequisite steps to use a customer-managed wallet with external calls or with scheduler SMTP email notifications.

Perform the prerequisite steps:

1. Verify that the Autonomous Database instance is configured with a private endpoint.

   See Configure Network Access with Private Endpoints for more information.

2. Set the `private_target` parameter to value `TRUE` when you grant the required ACLs with `dbms_network_acl_admin.append_wallet_ace`, or set `ROUTE_OUTBOUND_CONNECTIONS` to the value `PRIVATE_ENDPOINT`.

   See Use a Customer-Managed Wallet for External Calls with UTL_HTTP and Enhanced Security for Outbound Connections with Private Endpoints and for more information.

3. Obtain or create a customer-managed wallet.

   For example, to create a wallet with `orapki`:

   ```
   -- Create an SSL Wallet and load the Root CERTs using orapki utility
   $ORACLE_HOME/bin/orapki wallet create -wallet /u01/web/wallet -pwd ********
   $ORACLE_HOME/bin/orapki wallet add -wallet /u01/web/wallet -trusted_cert -
   cert MyWebServer.cer -pwd ********
   -- Store the credentials in the SSL Wallet using mkstore utility
   $ORACLE_HOME/bin/mkstore -wrl /u01/web/wallet -createCredential secret-
   from-the-wallet 'example@oracle.com' ********
   Enter wallet password: ********
   ```

   See Managing Public Key Infrastructure (PKI) Elements for more information.

4. Store the wallet to a bucket on Cloud Object Storage.

   After you obtain or create the customer-managed wallet containing self-signed certificates, store the wallet to a location on your Cloud Object Storage.

Use a Customer-Managed Wallet for External Calls with UTL_HTTP

When your Autonomous Database instance is on a private endpoint you can use a customer-managed wallet to handle external calls.

These steps describe using a customer-managed wallet with `UTL_HTTP`. The steps are the same for the other supported packages, including: `DMBS_LDAP`, `UTL_SMTP`, and `UTL_TCP`.

Perform the prerequisite steps to use a customer-managed wallet. See Prerequisites to Use a Customer-Managed Wallet with External Calls for more information.

To configure an Autonomous Database to use a customer-managed wallet:

1. Create a credential using `DBMS_CLOUD.CREATE_CREDENTIAL` to access Cloud Object Storage.

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'DEF_CRED_NAME',
       username => 'user1@example.com',
       password => 'password'
   ```

```
  );
END;
/
```

The values you provide for `username` and `password` depend on the Cloud Object Storage service you are using.

This creates the credential you use to access the Cloud Object Storage where the customer-managed wallet resides.

Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

2. Use an existing directory or create a new a directory for the wallet file.

   For example:

   ```
   CREATE DIRECTORY wallet_dir AS 'directory_path_of_your_choice';
   ```

   See Create Directory in Autonomous Database for information on creating directories.

3. Use `DBMS_CLOUD.GET_OBJECT` to upload the customer-managed wallet to the directory you created in the previous step, *WALLET_DIR*.

   For example:

   ```
   BEGIN
       DBMS_CLOUD.GET_OBJECT(
           credential_name => 'DEF_CRED_NAME',
           object_uri => 'https://objectstorage.us-
   phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/cwallet.sso',
           directory_name => 'WALLET_DIR');
   END;
   /
   ```

   In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

4. Grant the required ACLs to allow you to read the credentials from the specified directory.

   ```
   BEGIN
     dbms_network_acl_admin.append_wallet_ace(
         wallet_path => 'dir:WALLET_DIR',
         ace         => xs$ace_type(
              privilege_list => xs$name_list('use_client_certificates',
   'use_passwords'),
              principal_name => 'USER_NAME',
              principal_type => xs_acl.ptype_db)
       );
   END;
   /
   ```

5. Set the wallet path for use with `UTL_HTTP` procedures.

```
BEGIN
    UTL_HTTP.set_wallet('DIR:WALLET_DIR', 'password');
END;
/
```

   To specify the wallet directory with `UTL_HTTP.set_wallet` you can use either the `dir:` prefix or the `file:` prefix.

   The `dir:` prefix is only available on Autonomous Database. See UTL_HTTP for information on the `file:` prefix.

6. Now you can run `UTL_HTTP` procedures to access an endpoint using the customer-managed wallet with a self-signed certificate.

   For example:

```
SELECT UTL_HTTP.REQUEST('https://example.com') from dual;
```

Use a Customer-Managed Wallet with Scheduler Email Notifications

Describes the steps to use the scheduler SMTP email server with a customer-managed wallet.

Perform the prerequisite steps to use a customer-managed wallet. See Prerequisites to Use a Customer-Managed Wallet with External Calls for more information.

To use a customer-managed wallet with the scheduler email server:

1. Create a credential using `DBMS_CLOUD.CREATE_CREDENTIAL` to access your Cloud Object Storage.

   For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'user1@example.com',
    password => 'password'
  );
END;
/
```

   The values you provide for `username` and `password` depend on the Cloud Object Storage service you are using.

   This creates the credential you use to access Cloud Object Storage where the customer-managed wallet resides.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

2. Use an existing directory or create a new a directory for the wallet file.

For example:

```
CREATE DIRECTORY wallet_dir AS 'directory_path_of_your_choice';
```

See Create Directory in Autonomous Database for information on creating directories.

3. Use `DBMS_CLOUD.GET_OBJECT` to upload the wallet to the directory you created in the previous step, *WALLET_DIR*.

   For example:

```
BEGIN
    DBMS_CLOUD.GET_OBJECT(
        credential_name => 'DEF_CRED_NAME',
        object_uri => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/cwallet.sso',
        directory_name => 'WALLET_DIR');
END;
/
```

   In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

4. Run the commands to set up the scheduler to send SMTP email for scheduler job notifications.

   For example:

```
BEGIN
    DBMS_SCHEDULER.set_scheduler_attribute('EMAIL_SERVER','smtp.email.us-
ashburn-1.oci.oraclecloud.com:587');
    DBMS_CLOUD.create_credential('EMAIL_CRED', 'ocid1.user.oc1..username',
'password');

DBMS_SCHEDULER.set_scheduler_attribute('EMAIL_SERVER_CREDENTIAL','ADMIN.EMA
IL_CRED');

DBMS_SCHEDULER.set_scheduler_attribute('EMAIL_SERVER_ENCRYPTION','STARTTLS'
);
END;
/
```

   These commands set the scheduler email SMTP server, create the credential object containing the SMTP credentials and set the scheduler attribute for the SMTP credentials, and specify to use TLS for email sent for scheduler job notification.

   See SET_SCHEDULER_ATTRIBUTE Procedure for more information.

5. Create a credential to store the password for the customer-managed wallet.

```
BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
         credential_name  => 'WALLET_CRED',
         username         =>  'any_user',
         password         => 'password');
```

```
END;
/
```

This creates the credential you use in the next step to provide the password for the customer-managed wallet.

6. Set the scheduler wallet directory and the wallet credential.

```
BEGIN

DBMS_SCHEDULER.set_scheduler_attribute('EMAIL_SERVER_WALLET_DIRECTORY','WAL
LET_DIR');

DBMS_SCHEDULER.set_scheduler_attribute('EMAIL_SERVER_WALLET_CREDENTIAL',
'ADMIN.WALLET_CRED');
END;
/
```

7. Query the `DBA_SCHEDULER_GLOBAL_ATTRIBUTE` view to verify the values you set in the previous steps.

```
SELECT attribute_name, value
       FROM DBA_SCHEDULER_GLOBAL_ATTRIBUTE
       WHERE attribute_name LIKE 'EMAIL_SERVER%' ORDER BY 1, 2;


ATTRIBUTE_NAME                 VALUE
------------------------------
-----------------------------------------------
EMAIL_SERVER                   smtp.email.us-
ashburn-1.oci.oraclecloud.com:587
EMAIL_SERVER_CREDENTIAL        "ADMIN"."EMAIL_CRED"
EMAIL_SERVER_ENCRYPTION        STARTTLS
EMAIL_SERVER_WALLET_CREDENTIAL "ADMIN"."WALLET_CRED"
EMAIL_SERVER_WALLET_DIRECTORY  "WALLET_DIR"
```

Notes for Using a Customer-Managed Wallet with External Calls

Provides notes for using a customer-managed wallet with external calls.

- `DBMS_CLOUD` requests do not honor the custom wallet you set with `UTL_HTTP.set_wallet`. This includes `DBMS_CLOUD.SEND_REQUEST` and all object storage access for `DBMS_CLOUD` external tables that you define with `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`, `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE`, or `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE`. When you perform a query on an external table that you created with a `DBMS_CLOUD` procedure, the queries do not honor the custom wallet you set with `UTL_HTTP.set_wallet`.
- `APEX_WEB_SERVICE` requests do not honor the custom wallet you set with `UTL_HTTP.set_wallet`.

  To use a customer-managed wallet with APEX, specify the `p_wallet_path` parameter in `APEX_WEB_SERVICE` API calls or set the Wallet Path instance setting in APEX Administration Services.

  See Access Oracle APEX Administration Services for more information.

- Both auto-login wallets and password protected wallets are supported. When you use an auto-login wallet, specify `NULL` for the `wallet_password` parameter.
- The current user invoking the `UTL_HTTP.set_wallet` API must have `READ` privileges on the directory object or be granted the `CREATE ANY DIRECTORY` system privilege.
- The `UTL_HTTP.SET_AUTHENTICATION_FROM_WALLET` API is allowed. See SET_AUTHENTICATION_FROM_WALLET Procedure for more information.
- The `file:` prefix is supported with `UTL_HTTP.set_wallet` as long as the specified file path is `PATH_PREFIX` compliant.

  You can determine the `PATH_PREFIX` compliance for a path provided as input with the `DBMS_PDB_IS_VALID_PATH` procedure (granted to `PUBLIC` including the `ADMIN` user).

  For example:

```
with function check_path_prefix_compliance(file_path varchar2)
  return varchar2 as
BEGIN
  if dbms_pdb_is_valid_path(file_path) then
    return 'YES';
  else
    return 'NO';
  end if;
END;
SELECT
  check_path_prefix_compliance('/u03/dbfs/1276CDexample/data/dpdump') as
PATH_PREFIX_COMPLIANT,
  check_path_prefix_compliance('/u01/app/oracle/diag') as
PATH_PREFIX_COMPLIANT
FROM
  dual;
/
```

- To ensure backward compatibility when you use `UTL_HTTP.set_wallet`, in cases where the wallet path is ignored, input values such as `file:`, `NULL`, and so on, are accepted. These values are ignored and specify to use the default SSL wallet path with `UTL_HTTP.set_wallet`.
- The `DBMS_NETWORK_ACL_ADMIN` wallet ACL APIs, such as `APPEND_WALLET_ACL` are supported. These procedures allow you to grant/revoke wallet ACL privileges. See DBMS_NETWORK_ACL_ADMIN for more information.
- To support using the password credentials in an SSL wallet for authentication, the current user invoking `UTL_HTTP` APIs must have the "`use-passwords`" ACL privilege on the wallet path.
- Notes for setting `EMAIL_SERVER_WALLET_DIRECTORY` and `EMAIL_SERVER_WALLET_CREDENTIAL` with `DBMS_SCHEDULER.SET_SCHEDULER_ATTRIBUTE`:
  - To set the attribute values with `DBMS_SCHEDULER.SET_SCHEDULER_ATTRIBUTE` you must be an administrative user or a user with `MANAGE SCHEDULER` privileges (the ADMIN user has these privileges).
  - In addition to `MANAGE SCHEDULER` privilege, the user invoking `DBMS_SCHEDULER.SET_SCHEDULER_ATTRIBUTE` must have of `READ` privilege on the directory object you set with `EMAIL_SERVER_WALLET_DIRECTORY`, and `EXECUTE` privilege on the credential object you set with `EMAIL_SERVER_WALLET_CREDENTIAL`.

Use Cloud Tables to Store Logging and Diagnostic Information

You can create Cloud Tables where table data resides on Oracle managed Cloud Storage and the table data does not consume database storage.

- About Cloud Tables
  You can create Cloud Tables as a complementary alternative to in-database tables. All Cloud Table data is stored in Oracle managed Object Storage. Oracle managed Object Storage is external storage, outside of the database, that Autonomous Database creates and manages.
- Create Cloud Tables
  Shows the steps to create a Cloud Table on Autonomous Database.
- Cloud Table Notes
  Provides notes for Cloud Tables:

About Cloud Tables

You can create Cloud Tables as a complementary alternative to in-database tables. All Cloud Table data is stored in Oracle managed Object Storage. Oracle managed Object Storage is external storage, outside of the database, that Autonomous Database creates and manages.

You can use Cloud Tables to store infrequently used application logging data, diagnostic information, or to store other data. In some existing applications that do not run on Autonomous Database you might store this kind of information in files on a local file system (for example using `UTL_FILE` APIs). Such logging mechanisms and the associated files can be very helpful when you need to diagnose and resolve application errors. However, storing information in database tables can use large amounts of database storage for data that is infrequently used. Using Cloud Tables the persistent data is saved in Oracle managed Object Storage, without consuming database storage.

SELECT and DML Restrictions for Cloud Tables

Cloud Tables function like ordinary database tables with some restrictions. You can use SELECT and DML, data manipulation statements, with the following exceptions:

- `MERGE` statements are not supported.
- `LOB` columns are limited to 10MB of data.
- DML concurrency control is different, and therefore:
  - `LOCK TABLE` may not prevent concurrent DML as it does for a database table.
  - `INSERT` does not acquire a lock on the table, and therefore INSERT is never blocked by concurrent DML operations.
  - `UPDATE` and `DELETE` operations both acquire an exclusive lock on a Cloud Table. Therefore, `UPDATE` or `DELETE` transactions block concurrent `UPDATE` or `DELETE` operations on a Cloud Table.
- Only NOT NULL constraints are enforced.
- DML is allowed in a Read-Write Autonomous Database as it is for any other table; Cloud Tables also allow DML operations in a Refreshable Clone.

Cloud Tables do not support the following:

- Indexes

- Invisible columns
- Virtual columns
- DML triggers
- More than 996 columns
- Boolean data type columns

Lifecycle Management Operations and Cloud Tables

Cloud Table data is stored in Oracle managed Object Storage. This means certain operations on Autonomous Database handle Cloud Tables differently than in-database tables, as follows:

- Cloud Table data is excluded from an Autonomous Database instance's backup and recovery (the data is not backed up and you cannot restore Cloud Table data).
- Cloud Table Data is protected through Oracle managed Object Storage.
- The lifecycle management operations that impact the state of an Autonomous Database instance do not have an impact on the data stored in Cloud Tables.

Cloud Table naming in Object Storage is defined uniquely for each Autonomous Database instance, based on its OCID. This means that any operation that changes or introduces a new OCID for an existing database has an impact on Cloud Tables. The following illustrates the impact of lifecycle operations on Cloud Table data.

Lifecycle Operation	Cloud Table Data Availability
Same region database clone	Cloud Table is cloned without Cloud Table data
Cross-region database clone	Cloud Table is cloned without Cloud Table data
Same region (local) Autonomous Data Guard Standby	Cloud Table and Cloud Table data are accessible
Cross-region Autonomous Data Guard Standby	Cloud Table is available on the standby, without the Cloud Table data
Same region (local) Backup-Based Disaster Recovery peer	Cloud Table and Cloud Table data are accessible
Cross-region Backup-Based Disaster Recovery peer	Cloud Table is available on the standby, without Cloud Table data
Lifecycle management operations impacting the SCN/timestamp of an Autonomous Database instance, including: • Long term backup • Restore database (point in time restore) • Clone from backup	Cloud Table will continue to be updated and the old state of Cloud Table data is not preserved or restored. This means only the current Cloud Table data is available.
Lifecycle Management operations, including: • Manage resource allocation • Move • Shrink • Rename • Mode: Read-only/read-write • Change workload type: for example from Data Warehouse to Transaction Processing	No impact on Cloud Tables or on Cloud Table data

Create Cloud Tables

Shows the steps to create a Cloud Table on Autonomous Database.

To create a Cloud Table:

1. Run the `CREATE_CLOUD_TABLE` procedure.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CLOUD_TABLE(
      table_name  => 'CLOUD_TABLE_TEST',
      column_list => 'I INTEGER, STR1 VARCHAR2(32)' );
   END;
   /
   ```

   See CREATE_CLOUD_TABLE Procedure for more information.

2. Insert data into the Cloud Table.

   ```
   INSERT INTO cloud_table_test VALUES (1, 'xyz');
   ```

3. Select data from a Cloud Table.

   ```
   SELECT * FROM cloud_table_test;


   I          STR1
   ---------- --------------------------------
   1          xyz
   ```

Use `DROP TABLE` when you want to drop a Cloud Table.

For example:

```
DROP TABLE CLOUD_TABLE_TEST;
```

Cloud Tables do not support the recycle bin.

See Cloud Table Notes for additional information.

Cloud Table Notes

Provides notes for Cloud Tables:

- You can grant `SELECT`, `INSERT`, and `UPDATE` privileges for a Cloud Table. No other privileges can be granted to a Cloud Table.

  See Configuring Privilege and Role Authorization for more information.

- Cloud Table constraints are limited to constraints in `RELY DISABLE NOVALIDATE` mode, which means the constraint is not enforced. The only exception to this is for `NOT NULL` constraints.

  Cloud Tables support all `NOT NULL` constraint modes including the default mode (`ENABLE VALIDATE`). `PRIMARY KEY`, `UNIQUE`, `FOREIGN KEY`, and `NOT NULL` constraints are supported; `CHECK` constraints are not supported.

  You can declare constraints inline as part of `COLUMN_LIST`.

For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CLOUD_TABLE(
        table_name  => 'CLOUD_TAB_WITH_CONSTRAINTS',
        column_list => 'PK INTEGER,
            DATE_ID INT REFERENCES DATE_DIM(DATE_ID) RELY DISABLE
NOVALIDATE,
            VAL NUMBER NOT NULL,
            CONSTRAINT CLOUD_TAB_PK PRIMARY KEY(PK) RELY DISABLE
NOVALIDATE');
END;
/
```

See Cloud Table Notes for additional Cloud Table limitations.

- The `DBMS_CLOUD` package is an invoker's rights package. The `DBMS_CLOUD.CREATE_CLOUD_TABLE` procedure only allows you to create a table in the invoker's schema.

  See Invoker's Rights and Definer's Rights Clause for more information.

- The `column_list` parameter in a `DBMS_CLOUD.CREATE_CLOUD_TABLE` procedure call can include the `DEFAULT` clause, which functions like the `DEFAULT` clause in `CREATE TABLE`. See CREATE TABLE for more information.

  For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CLOUD_TABLE(
   table_name  => 'CLOUD_TABLE_TEST_DEFAULT',
   column_list => 'I INTEGER, STR2 VARCHAR2(32) DEFAULT ''ABC''');
END;
/
```

  Then you can insert with default values. For example:

```
INSERT INTO cloud_table_test_default (i) VALUES (1);
1 row created.
INSERT INTO cloud_table_test_default VALUES (2, default);
1 row created.
INSERT INTO cloud_table_test_default VALUES (3, null);
1 row created.
INSERT INTO cloud_table_test_default VALUES (4, 'xyz');
1 row created.
COMMIT;
Commit complete.
SELECT * FROM cloud_table_test_default ORDER BY i;

I STR2
- ----
1 ABC
2 ABC
3 null
4 xyz
```

Analyze

Provides information on the built-in tools that allow you to explore your data with reports, analytics, and visualization.

- Build Reports and Dashboards Using Oracle Analytics
  You can use Oracle Analytics Cloud and Oracle Analytics Desktop with Autonomous Database.
- Creating Dashboards, Reports, and Notebooks with Autonomous Database
  Autonomous Database includes support for building dashboards, reports, and notebooks for data analysis using Oracle Analytics Desktop and Oracle Machine Learning Notebooks.
- Use Oracle Spatial with Autonomous Database
  Oracle Spatial Database is included in Autonomous Database, allowing developers and analysts to get started easily with location intelligence analytics and mapping services.
- The Data Analysis Tool
  The Data Analysis tool enables you to create Analytic Views with multidimensional metadata.
- The Catalog Tool
  The Catalog Tool allows you browse, search and discover data in, and connected to, the Oracle Autonomous Database.
- The Data Insights Page
  The Data Insights page displays information about patterns and anomalies in the data of entities in your Oracle Autonomous Database.
- Use Oracle OLAP with Autonomous Database
  Oracle OLAP is available on Autonomous Database. Oracle OLAP supports uses such as planning, forecasting, budgeting, and analysis.

Build Reports and Dashboards Using Oracle Analytics

You can use Oracle Analytics Cloud and Oracle Analytics Desktop with Autonomous Database.

Working with Oracle Analytics Cloud

You can use Oracle Analytics Cloud with Autonomous Database. Use Oracle Analytics Cloud to select interactive visualizations and automatically create advanced calculations to reveal the insights in your data.

For more information, see Use Oracle Analytics Cloud with Autonomous Database.

Working with Oracle Analytics Desktop

You can use Oracle Analytics Desktop with Autonomous Database. Just connect to Autonomous Database, select the elements that you're interested in, and let Oracle Analytics Desktop find the best way to visualize it. Choose from a variety of visualizations to look at data in a specific way.

For more information see Use Oracle Analytics Desktop with Autonomous Database.

Creating Dashboards, Reports, and Notebooks with Autonomous Database

Autonomous Database includes support for building dashboards, reports, and notebooks for data analysis using Oracle Analytics Desktop and Oracle Machine Learning Notebooks.

Oracle Analytics Desktop is designed for business users and report developers. Oracle Analytics Desktop enables you to gain insight into your data and create reports and build interactive visualizations. It also gives you a preview of the self-service visualization capabilities included in Oracle Analytics Cloud, which extends the data exploration and visualization experience by offering secure sharing and collaboration across the enterprise, additional data sources, and a full mobile experience including proactive self-learning analytics delivered to your device.

Oracle Machine Learning Notebooks is a built-in browser-based interactive data analysis environment that is part of Autonomous Database. Oracle Machine Learning Notebooks is designed for data scientists who want use the extensive library of in-database machine learning features.

- Create Dashboards and Reports to Analyze and Visualize Your Data
  Gain insight into your data with Oracle Analytics Desktop and Oracle Analytics Cloud. These tools let you explore your Autonomous Database data through advanced analytics and interactive visualizations.
- Create Notebooks, Workspaces, and Projects with Oracle Machine Learning Notebooks
  Oracle Machine Learning Notebooks provides a collaborative interface for creating notebooks and for performing real-time analytics.

Create Dashboards and Reports to Analyze and Visualize Your Data

Gain insight into your data with Oracle Analytics Desktop and Oracle Analytics Cloud. These tools let you explore your Autonomous Database data through advanced analytics and interactive visualizations.

- Use Oracle Analytics Desktop with Autonomous Database
  Gain insight into your data with Oracle Analytics Desktop. Oracle Analytics Desktop lets you explore your Autonomous Database data through interactive visualizations.
- Use Oracle Analytics Cloud with Autonomous Database
  You can use Oracle Analytics Cloud with Autonomous Database. Oracle Analytics Cloud provides a complete set of tools for deriving and sharing data insights.

Use Oracle Analytics Desktop with Autonomous Database

Gain insight into your data with Oracle Analytics Desktop. Oracle Analytics Desktop lets you explore your Autonomous Database data through interactive visualizations.

Oracle Analytics Desktop provides powerful personal data exploration and visualization in a simple per-user desktop download. Oracle Analytics Desktop is the perfect tool for quick exploration of sample data from multiple sources or for rapid analysis and investigation of your own local data sets.

Oracle Analytics Desktop makes it easy to visualize your Autonomous Database data so you can focus on exploring interesting data patterns. Just connect to Autonomous Database, select the elements that you're interested in, and let Oracle Analytics Desktop find the best way to visualize it. Choose from a variety of visualizations to look at data in a specific way.

Oracle Analytics Desktop also gives you a preview of the self-service visualization capabilities included in Oracle Analytics Cloud, Oracle's industrial-strength cloud analytics platform. Oracle Analytics Cloud extends the data exploration and visualization experience by offering secure sharing and collaboration across the enterprise, additional data sources, greater scale, and a full mobile experience including proactive self-learning analytics delivered to your device. Try Oracle Analytics Desktop for personal analytics and to sample a taste of Oracle's broader analytics portfolio.

Oracle Analytics Desktop's benefits include:

- A personal, single-user desktop application.
- Offline availability.
- Completely private analysis.
- Full control of data source connections.
- Lightweight single-file download.
- No remote server infrastructure.
- No administration tasks.

See the *User's Guide for Oracle Analytics Desktop* for information on connecting Oracle Analytics Desktop to Autonomous Database.

To get started with Oracle Analytics Desktop use the following link to visit the software download page which includes more information about system requirements and provides instructions for installing Oracle Analytics Desktop on different platforms:

Oracle Analytics Desktop Download Details

Use Oracle Analytics Cloud with Autonomous Database

You can use Oracle Analytics Cloud with Autonomous Database. Oracle Analytics Cloud provides a complete set of tools for deriving and sharing data insights.

- Data preparation: Analysts can ingest, profile, and cleanse data using a variety of algorithms.
- Data flow: Analysts can prepare, transform and aggregate data, and then run machine-learning models at scale.
- Data discovery: Subject matter experts can easily collaborate with other business users, blending intelligent analysis at scale, machine learning, and statistical modeling.
- Data visualization: Analysts can visualize any data, on any device, on-premise and in the cloud.
- Data collaboration: Large organizations and small teams can share data more simply, as you don't need to manage or consolidate multiple versions of spreadsheets.
- Data-driven: Application developers can utilize interfaces that enable them to extend, customize, and embed rich analytic experiences in the application flow.

See *Visualizing Data and Building Reports in Oracle Analytics Cloud* for details on connecting Autonomous Database with Oracle Analytics Cloud.

Create Notebooks, Workspaces, and Projects with Oracle Machine Learning Notebooks

Oracle Machine Learning Notebooks provides a collaborative interface for creating notebooks and for performing real-time analytics.

- Using Oracle Machine Learning User Interface on Autonomous Database
  Oracle Machine Learning Notebooks provide a browser-based interactive data analysis environment where you can develop, document, share, and automate analytical methodologies.

- Work with Oracle Machine Learning User Interface for Data Access, Analysis, and Discovery
  You can use Oracle Machine Learning Notebooks to access data and for data discovery, analytics, and notebooks.

Using Oracle Machine Learning User Interface on Autonomous Database

Oracle Machine Learning Notebooks provide a browser-based interactive data analysis environment where you can develop, document, share, and automate analytical methodologies.

Oracle Machine Learning Notebooks includes:

- **Oracle Machine Learning User Administration Application**
  - Web based administrative UI for managing (list, create, update, delete) Oracle Machine Learning users
  - Notebook users map to database users.
  - Access to the User Management feature is limited to the database administrator account (ADMIN). See Create and Update User Accounts for Oracle Machine Learning Components on Autonomous Database for more information.
- **Oracle Machine Learning Notebooks Application**
  - Web based application for data scientists
  - Allows for creation of workspaces, projects, and notebooks

  See Work with Oracle Machine Learning User Interface for Data Access, Analysis, and Discovery for information on accessing Oracle Machine Learning user workspaces, projects, and notebooks from your Autonomous Database.

Work with Oracle Machine Learning User Interface for Data Access, Analysis, and Discovery

You can use Oracle Machine Learning Notebooks to access data and for data discovery, analytics, and notebooks.

To access Oracle Machine Learning Notebooks you can use the Oracle Cloud Infrastructure Console or Database Actions.

To access Oracle Machine Learning Notebooks from the Oracle Cloud Infrastructure Console:

1. From the Display Name column, select an Autonomous Database.
2. On the Autonomous Database Details page click the **Tools** tab.
3. In the **Tools** column, select the **Oracle ML User Administration** link.
4. Enter your **Username** and **Password**.

   The Administrator creates OML users. See Create and Update User Accounts for Oracle Machine Learning Components on Autonomous Database for details.
5. Click **Sign in**.

To access Oracle Machine Learning Notebooks from Database Actions:

1. On the Oracle Cloud Infrastructure Console, from the Display Name column select an Autonomous Database.
2. On the Autonomous Database Details page select **Database Actions** and click **View all database actions**.

3. On the Database Actions Launchpad, under **Development**, select **Oracle Machine Learning**.

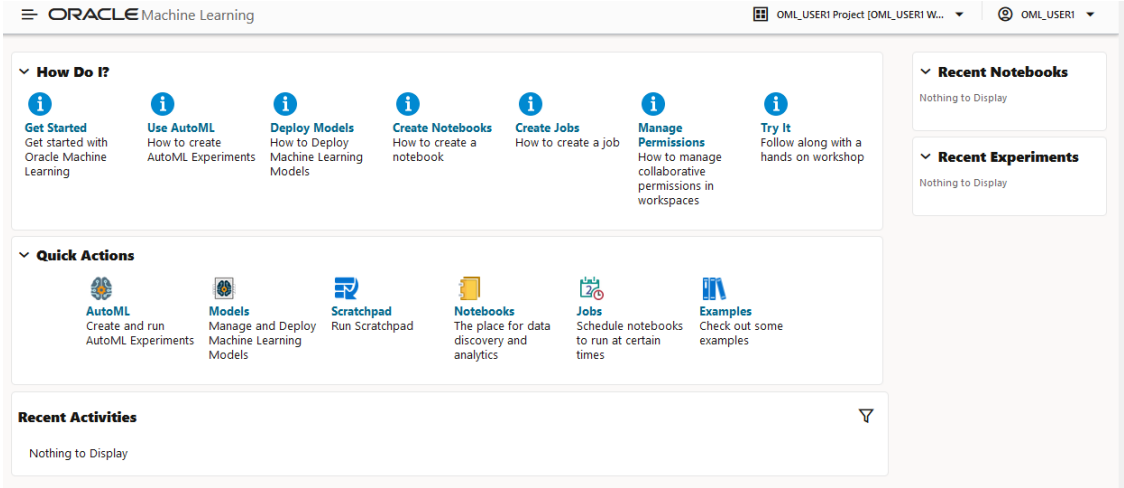
   If you do not see the **Oracle Machine Learning** card then your database user is missing the required `OML_DEVELOPER` role. See Required Roles to Access Tools from Database Actions for more information.

4. Enter your username and password.

   The Administrator creates OML users. See Create and Update User Accounts for Oracle Machine Learning Components on Autonomous Database for details.

5. Click **Sign In**.

This shows Oracle Machine Learning user application.



Oracle Machine Learning Notebooks allows you to access your data in your database and build notebooks with the following:

- Data Ingestion and Selection
- Data Viewing and Discovery
- Data Graphing, Visualization, and Collaboration
- Data Analysis

You can also create and run SQL statements and create and run SQL scripts that access your data in your database.

Use Oracle Spatial with Autonomous Database

Oracle Spatial Database is included in Autonomous Database, allowing developers and analysts to get started easily with location intelligence analytics and mapping services.

- About Oracle Spatial with Autonomous Database
- Oracle Spatial Limitations with Autonomous Database
  Autonomous Database includes Oracle Spatial, with some limitations.

About Oracle Spatial with Autonomous Database

Oracle Spatial allows developers and analysts to use spatial analysis in every application from basic spatial search and analysis to advanced geospatial applications and Geographic Information Systems (GIS).

Organizations can manage different types of geospatial data, perform hundreds of spatial analytic operations, and use interactive map visualization tools with the spatial features in Oracle Autonomous Database and Oracle Database.

The spatial features provide a schema and functions that facilitate the storage, retrieval, update, and query of collections of spatial features in Autonomous Database. Spatial consists of the following:

- A schema (MDSYS) that prescribes the storage, syntax, and semantics of supported geometric data types
- A spatial indexing mechanism
- Operators, functions, and procedures for performing area-of-interest queries, spatial join queries, and other spatial analysis operations
- Utility functions and procedures for validating, loading, extracting, and working with spatial data
- GeoRaster, a feature that lets you store, index, query, analyze, and deliver GeoRaster data, that is, raster image and gridded data and its associated metadata
- Features for geocoding address data and for reverse geocoding longitude/latitude data to a street address:
  - SDO_GCDR.ELOC_GEOCODE
  - SDO_GCDR.ELOC_GEOCODE_AS_GEOM

> **Note:**
>
> `SDO_GCDR.ELOC_GEOCODE` and `SDO_GCDR.ELOC_GEOCODE_AS_GEOM` functions are only available on Autonomous Database.

See the following for more information:

- Spatial Concepts
- Oracle Database Spatial
- Geocoding Address Data

Oracle Spatial Limitations with Autonomous Database

Autonomous Database includes Oracle Spatial, with some limitations.

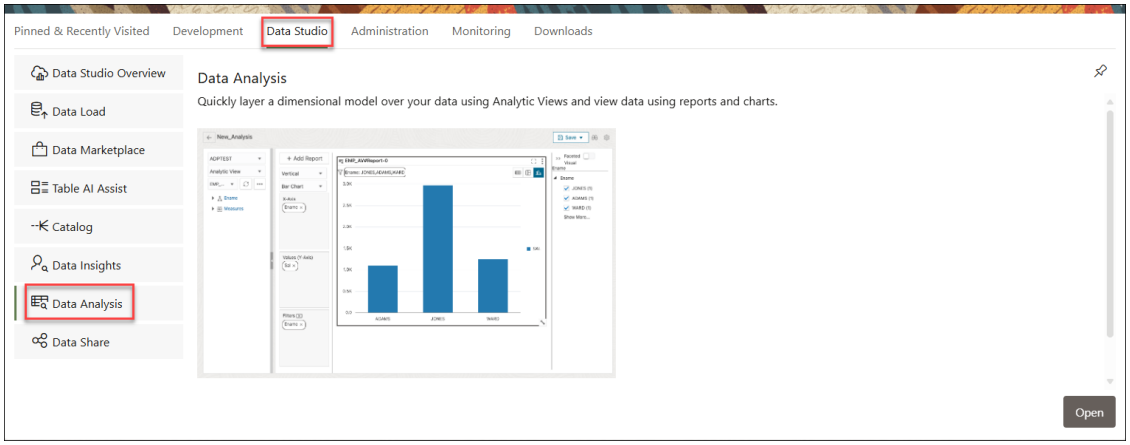
You can use the Oracle spatial features with Autonomous Database, with the following limitations:

- The following Oracle Spatial components require that you download applications from Oracle Cloud Marketplace:
  - Oracle Spatial Studio
  - Oracle Spatial Map Visualization
  - Network Data Model
  - Geospatial Consortium Web Services
- The following is not available in Autonomous Database:
  - Router

The Data Analysis Tool

The Data Analysis tool enables you to create Analytic Views with multidimensional metadata.

To reach the **Data Analysis** page, select the **Data Analysis** menu in the Data Studio tab of the Launchpad.

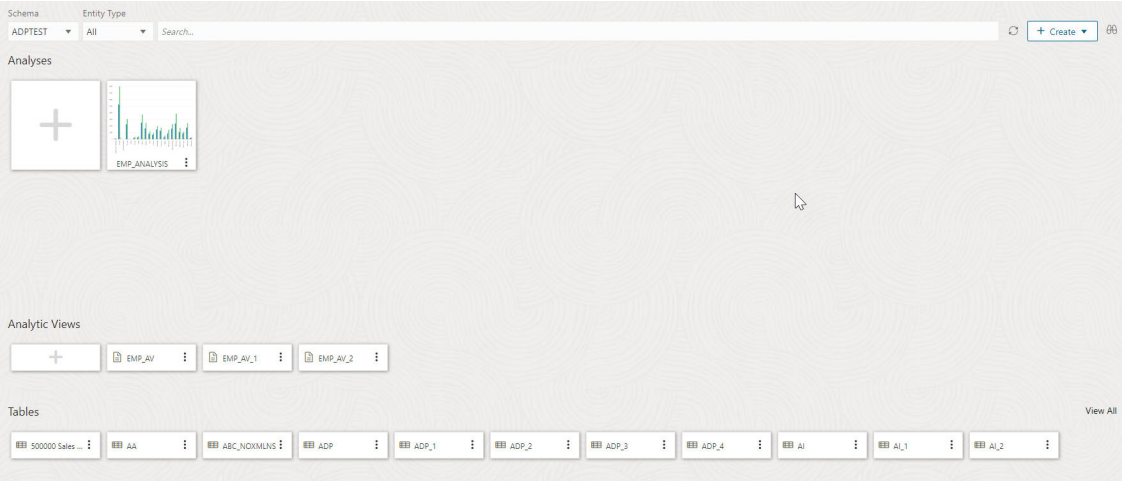


You create Analytic Views on top of a fact table with several dimensions and hierarchies. Analytic views refer to tables in the database and allow users to create hierarchies for dimensions. You can also create Analyses and reports using information from the Analytic Views. The Data Analysis homepage enables you to search for Analyses, view and perform tasks such as edit, delete, view or rename Analyses. You can also analyze, find errors, export, edit, compile and delete Analytic Views. You can analyze tables and generate SQL reports from them.

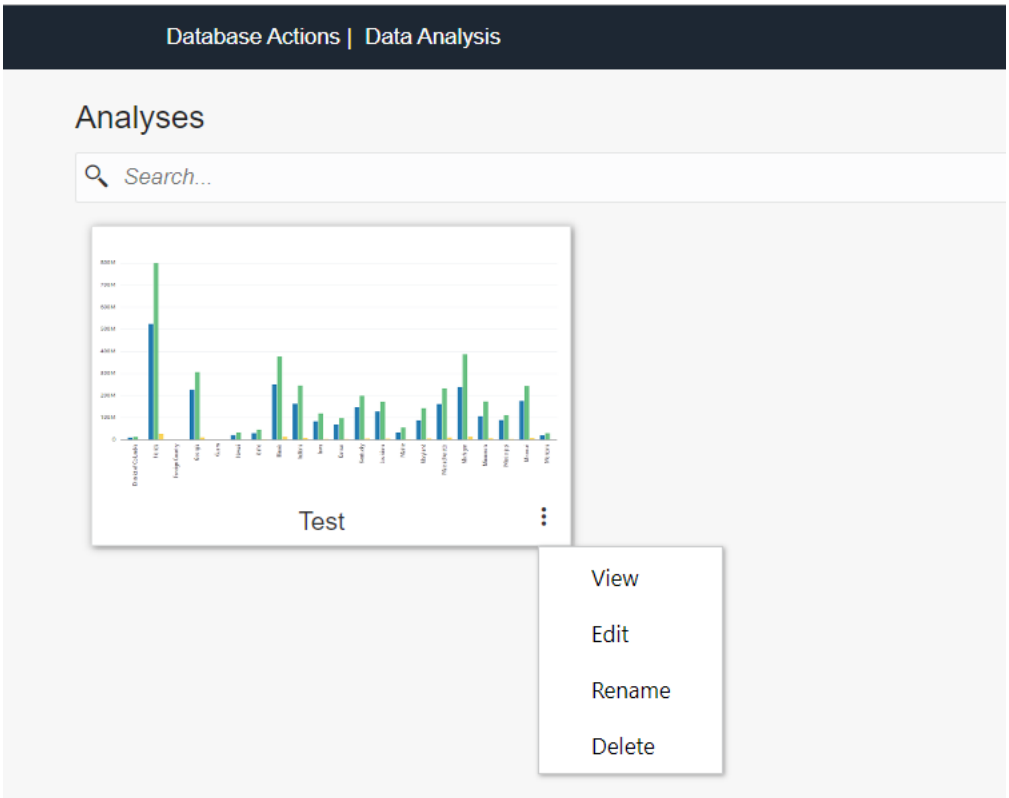
Select the **Data Analysis** card from the Data Studio suite to access this tool. You can also access it by clicking the Selector icon and selecting Data Analysis from the Data Tools menu in the navigation pane.

> **Note:**
>
> If you do not see the Data Analysis card then your database user is missing the required DWROLE role.

The Data Analysis home page consists of three parts: Analyses, Analytic Views and Tables.



Analyses



The top section of the homepage comprises of a list of Analyses. Use the search field to search for Analyses you create. The top section of the homepage comprises of a list of Analyses. Use the search field to search for Analyses you create.

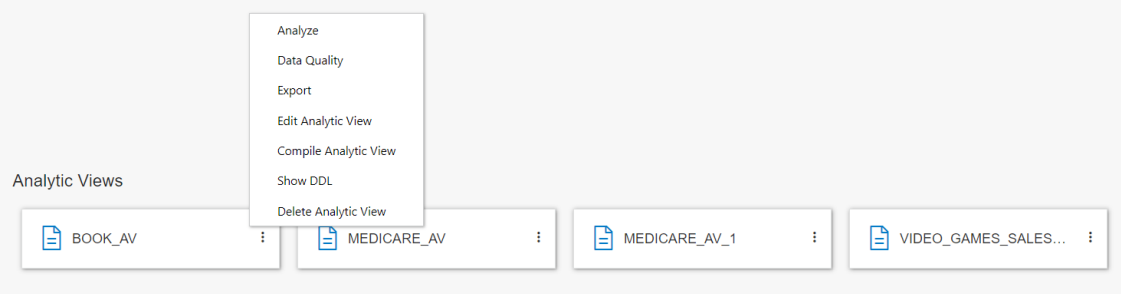
Analyses are analysis of multiple Analytic Views. The Analyses card displays the name of the analysis. Click **Actions** (three vertical dots) to open the context menu.

The actions available are:

- **View:** Opens the Analysis View page in a new window where you can view the analysis.
- **Edit:** Opens the selected Analysis page where you can edit the reports present in the analysis.

- **Rename:** Allows you to rename the Analysis. Click **save** to modify the new name.
- **Delete:** Opens the delete Analysis dialog where you can delete the analysis.

Analytic Views



The bottom section of the homepage displays list of existing Analytic Views. Each Analytic View card displays the name of the Analytic View. The Actions icon enables you to manage the Analytic View. Click **Actions** (three vertical dots) to open the context menu. The actions available are:

- **Analyze:** Opens the Analytic View browser and the Analysis View page in a new window where you can view the Analysis.
- **Data Quality:** Opens the Data Quality page where the tool validates the selected Analytic View for errors and lists them out.
- **Export:** Allows you to export the Analytic View to Tableau and PowerBI.
- **Edit Analytic View:** Opens the Edit Analytic View dialog box where you can edit the properties of the selected Analytic View.
- **Compile Analytic View:** This option compiles the Analytic View and returns compilation errors if there are any.
- **Show DDL:** Displays the DDL statements for the Analytic View.
- **Delete Analytic View:** Deletes the selected Analytic View.

The **+Create** button enables you to create **Analysis** and create **Analytic View** from the home page.

You can select both hierarchies and measures from Analytic Views. Hierarchies are DB objects that allow users to define relationships between various levels or generations of dimension members. As the name implies, hierarchies organize data using hierarchical relationships. With this tool you can analyze and visualize data in different Points of View (POV). You can export the metadata and visualize it with sophisticated tools like Oracle Analytics Cloud (OAC) and Tableau.

Advantages of Data Analysis tool

With Data Analysis tool you can:

- Visualize, analyze and inspect your data clearly and efficiently with pivot tables
- Calculate total number of errors present in the Analytic View you create and provide solutions to minimize the errors
- Automatically display meaningful insights to help you make better decisions
- Analyze your data across dimensions with support for hierarchical aggregation and drill-down
- Share your Analytic Views with the tool of your choice over various options of raw data consumption to draw meaningful insights and make them accessible to any user

By identifying relationships among tables and columns, Analytic Views enable your system to optimize queries. They also open new avenues for analyzing data. These avenues include data insights, improved hierarchy navigation, and the addition of hierarchy-aware calculations.

This tool runs complex and hierarchical SQL queries along with SQL extensions in the background, which simplifies real-time calculations. It makes complex data more accessible and easier to understand.

The Data Analysis Page

The following section describes searching and obtaining information about Analytic Views, creating Analytic Views, inspecting your data, discovering insights and visualizing data using tools like Oracle Analytics Cloud (OAC), Tableau, and Microsoft Power BI.

> **Note:**
>
> - OAC has in-built tools to search and utilize Analytic Views.
> - We have no direct support for Microsoft Power BI, yet its users can map their tool to the AV transparency views to avail some of the benefits of Analytic Views.

Read these topics for detailed descriptions on Analyses and Analytic Views:

- Searching and obtaining information about Analytic Views
  When you first open the Data Analysis page, it displays the list of schemas and Analytic Views. With Select Schema, you can select a preferred Schema from a list of schemas available in the drop-down.
- Creating Analytic Views
  You can create Analytic Views and view information about them. You can also edit and perform other actions on them.
- Working with Analyses
  Analyses are a collection of multiple reports on a single page, which provides quick access to multiple data analyses collected from different Analytic Views.
- Viewing Analyses
  Analyses provide you an insight into the performance of your data.
- Workflow to build Analyses
  Here is the workflow to build an analyses.
- Creating Analyses
  Use the Data Analysis tool to create and edit your Analyses. The analysis provides you customized view of Analytic View data. An analysis consists of one or more reports that displays the results of analysis.
- Creating Reports
  A single report you generate analyzes an AV based on the Levels and measures you select.
- Run Natural Language Query in the Data Analysis Tool
  You can query the Oracle Autonomous Database by using Natural Language Query rather than having to write SQL Query.
- Generate Natural Language explanation from SQL Query using AI Explain
  You can translate SQL queries to natural language that is understood by you.

- Vector Search in the Data Analysis Tool
  You can use Oracle AI Vector Search capabilities to search for relevant text from your source data on a specific column.
- Format Numeric Columns in the Data Analysis Tool
  The custom formatting of numeric columns enables you to select your format type such as **Number**, **Currency** and **Percentage** using Mapping numeric column wizard in the Data Analysis tool.
- Using Calculation Templates
  The Data Analysis tool provides templates for all of the calculations typically in demand for business intelligence applications.
- Oracle Autonomous Database for Excel
  The Oracle Autonomous Database for Excel integrates Microsoft Excel spreadsheets with the Autonomous Database to retrieve and analyze data from Analytic Views in the database. You can also directly run SQL queries to view their results in the worksheet.
- Oracle Autonomous Database for Google Sheets
  The Oracle Autonomous Database add-on  enables you to query tables using SQL or Analytic Views using a wizard directly from Google Sheets for analysis.

Searching and obtaining information about Analytic Views

When you first open the Data Analysis page, it displays the list of schemas and Analytic Views. With Select Schema, you can select a preferred Schema from a list of schemas available in the drop-down.

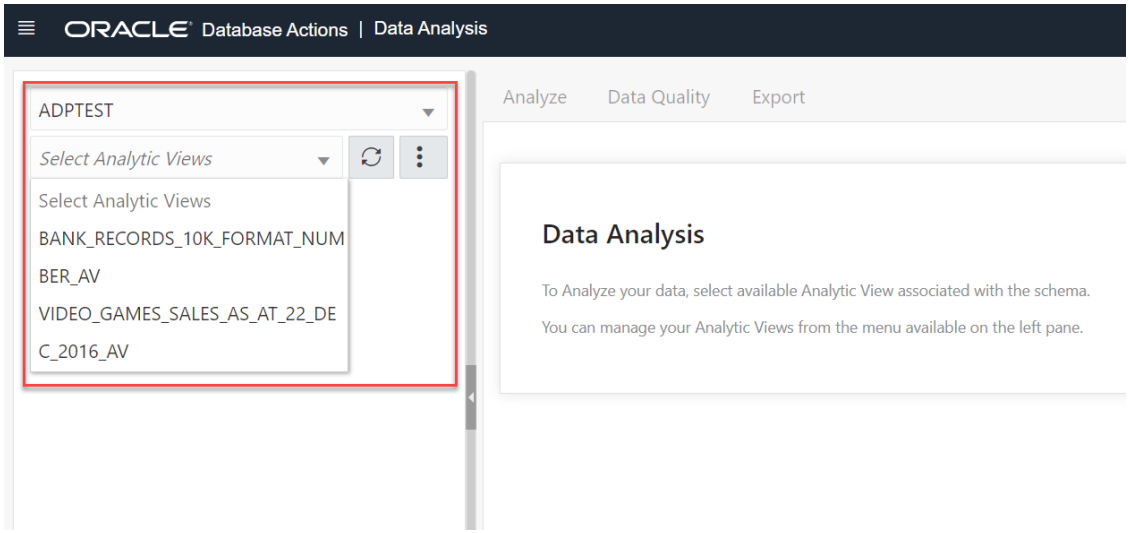
The **Select Analytic Views** drop-down enables you to select an available Analytic View associated with the schema. When you create an Analytic View, it appears in the drop-down option with your schema. The Refresh AV icon refreshes the contents of the selected Analytic View.

The **Action** icon next to the Refresh AV button enables you to manage Analytic Views. You can Create Analytic View, Edit Analytic View, Compile Analytic View, Show the Data Definition Language (DDL) that generates the Analytic View or Delete Analytic View from the menu.

Obtain information about Analytic Views

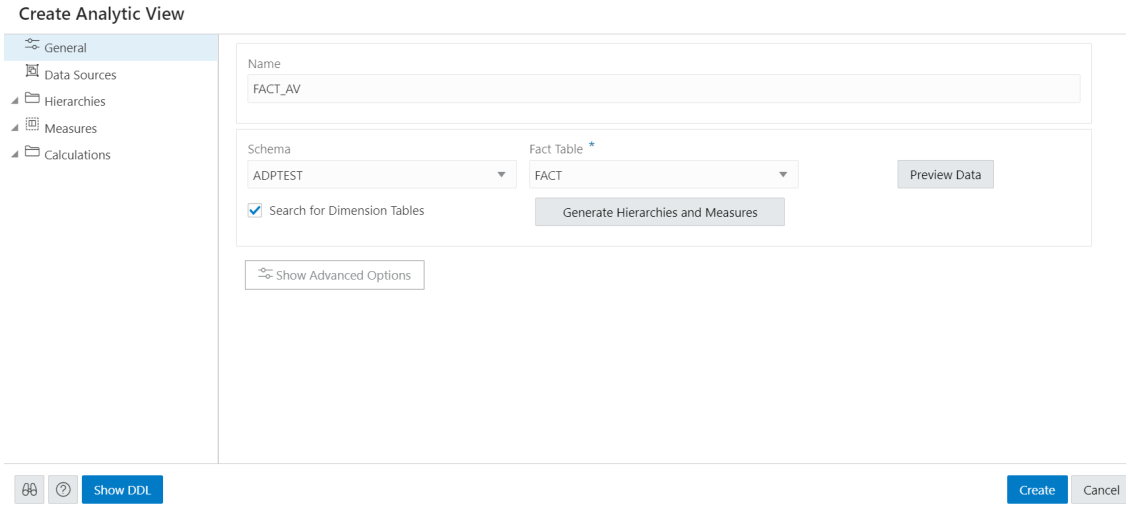
By default, Analytic Views are filtered by the current user's schema, as indicated by the schema list below the menu-bar. You can remove the selected schema filter by selecting another user's schema. To search for Analytic Views in other schemas, select one of the schemas from the drop-down.

If there is no Analytic View associated with the schema selected, the tool prompts you to create an Analytic View.



Creating Analytic Views

You can create Analytic Views and view information about them. You can also edit and perform other actions on them.



When you create an Analytic View, you identify a fact table that contains the data to inspect. The **Generate Hierarchies and Measures** button looks at the contents of that table, identifies any hierarchies in the fact table, and searches for other tables that may contain related hierarchies.

While creating an Analytic View, you can enable or disable the following advanced options:

- Autonomous Aggregate Cache, which uses the dimensional metadata of the Analytic View to manage a cache and that improves query response time.
- Analytic View Transparency Views, which presents Analytic Views as regular database views and enables you to use your analytic tools of choice while gaining the benefits of Analytic Views.
- Analytic View Base Table Query Transformation, which enables you to use your existing reports and tools without requiring changes to them.

Create Analytic View

To create Analytic View, click **Create** from the Data Analysis home page and select **Create Analytic View** to begin the process.

Click **Cancel** to cancel the creation of the Analytic View at any time.

Specify Attributes of the Analytic View

On the **General** tab of the **Create Analytic View** pane, specify the following:

- The name for the Analytic View
- The fact table for the view
- Advanced options

You can also preview the data of the fact table and see statistics about that data.

In the **Name** field, specify a name of your choice.

The **Schema** field has the current user's schema. You can only create an Analytic View in that schema.

In the **Fact Table** field, expand the drop-down list and click **More Sources**. The **Select Sources** dialog box has a list of the available tables and views. Select a table or view from the list.

To filter the list, begin typing characters in the Filter field. As you type, the list changes to show the tables or views that contain the characters. Clear the field to show the complete list again. After you select a table or view, click OK.

To enable or disable the advanced options, on the **Create Analytic View** pane, click the **Show Advanced Options** icon at the bottom left. Select or deselect options as desired.

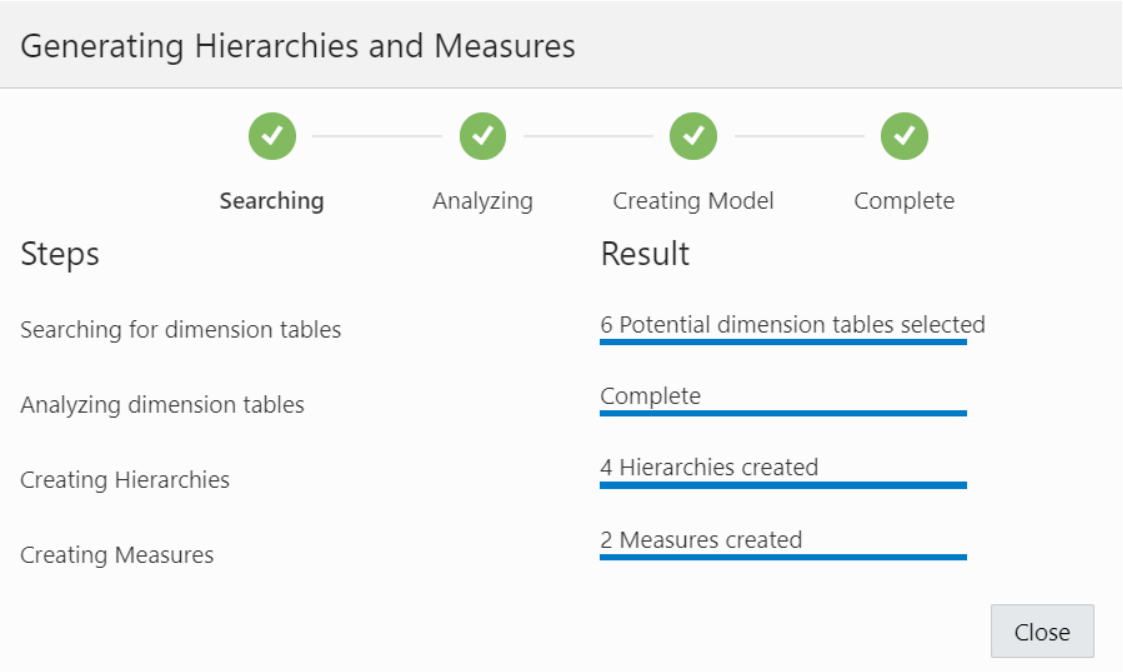
To view the data in the fact table and statistics about the data, click the **Preview Data** button. In the Preview and Statistics pane, the Preview tab shows the columns of the table and the data in the columns.

The Statistics tab shows the size of the table and the number of rows and columns. The statistics may take a few moments to appear, during which time the message, "No statistics available..." may appear. The statistics include the names of the columns, their data types, the number of distinct values and null values, the maximum and minimum values, and other information. The bar graph displays the top unique column values and the number of their occurrences for the selected column. Point to a bar in the graph to see the number of occurrences of the unique value.

Click **Close** to close the Preview and Statistics pane and return to the Create Analytic View pane.

Click on **Generate Hierarchies and Measures** icon.

The Generating Hierarchies and Measures dialog box displays the progress of searching for dimension tables, analyzing the dimension tables and identifying and creating the data sources, joins, hierarchies, and measures to use. When the process completes, click **Close**.



The **Search for Dimension Tables** check box when selected, enables you to search for dimension tables while generating hierarchies and measures.

After the hierarchies and measures are generated, they are displayed under their respective tabs. Review the hierarchies and measures created for you.

Specify the **Name**, **Fact Table** and select **Advanced Options** in the **General** tab of **Create Analytic View** pane. Click **Create** to generate an Analytic View.

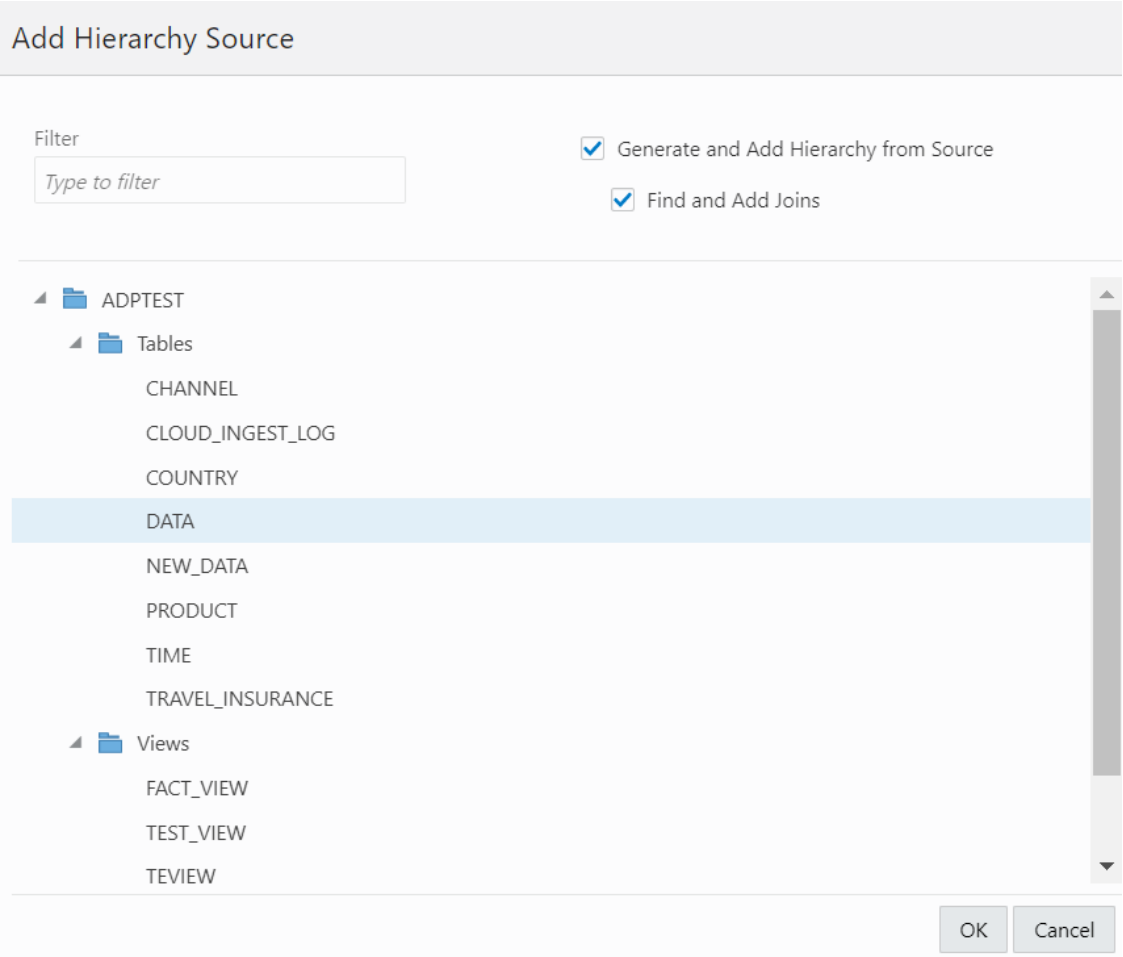
View Data Sources

The Data Sources tab displays the sources of the data and the relationships among them. It has a graphical display of the fact table and the related dimension tables. For example, a fact table of health insurance data might have columns for geography identifiers, income codes, and gender codes. The Data Sources tab would display items for the fact table and for the geography, income, and gender dimension tables.

You can add hierarchies from data sources even after generating hierarchies from the existing fact table. You can add one or more hierarchies to your new or existing analytic view. Multiple hierarchies can be defined and used in an analytical view, however only one will be used by default.

Right-click the Data Sources tab and select **Add Hierarchy Sources** or select **Add Hierarchy Sources**.

Selecting **Add Hierarchy Sources** launches an **Add Hierarchy Source** dialog box.



You can view all the fact tables and views associated with the analytic view.

In the filter field, you can either manually look for the source or start typing to search for the fact table or views from the list of available fact tables and views. After typing the full name of the source, the tool automatically matches the fact table or view.

Select **Generate and Add hierarchy from Source** to generate analysis and hierarchies associated with the source data you select.

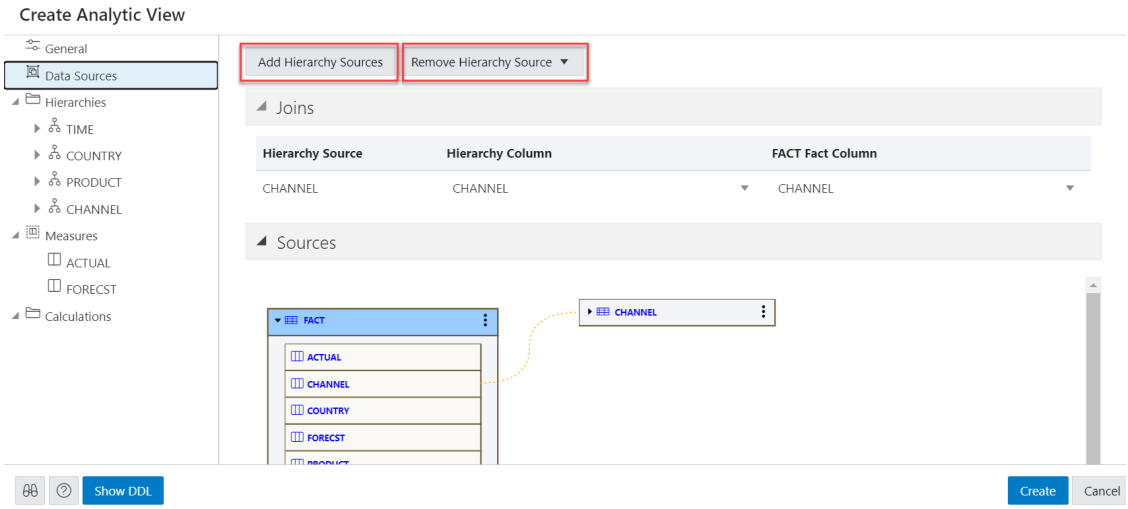
Select **Find and Add Joins** to link all the data sources with the fact table. You can add multiple join entries for a single hierarchy.

Click **OK** to select the source.

The Generating Hierarchies and Measures dialog box displays the progress of analyzing the dimension tables and creating the hierarchies. When the process completes, click **Close**.

> **Note:**
>
> When you add a hierarchy from the data source, you see the new hierarchy in the list of hierarchies in the Hierarchies tab. You can navigate between the Data Sources tab, the Hierarchies tab, the Measures tab, the Calculations tab. You can add a hierarchy from a source that is not connected by navigating back to the Data Sources tab.

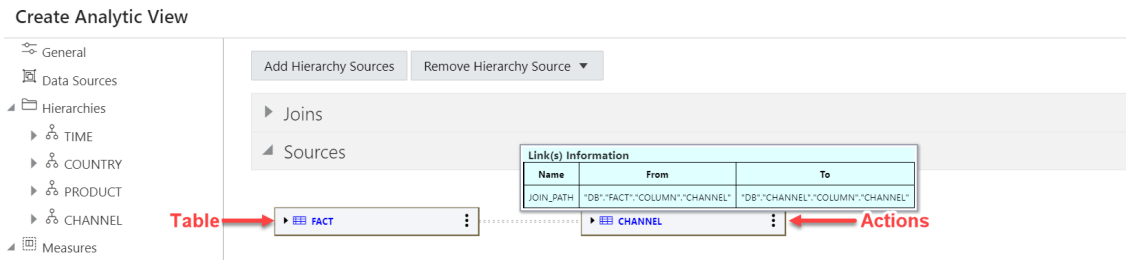
Select **Remove Hierarchy Source** to remove the hierarchies you create from the data sources. You cannot remove hierarchies generated from the fact table you select from this option.



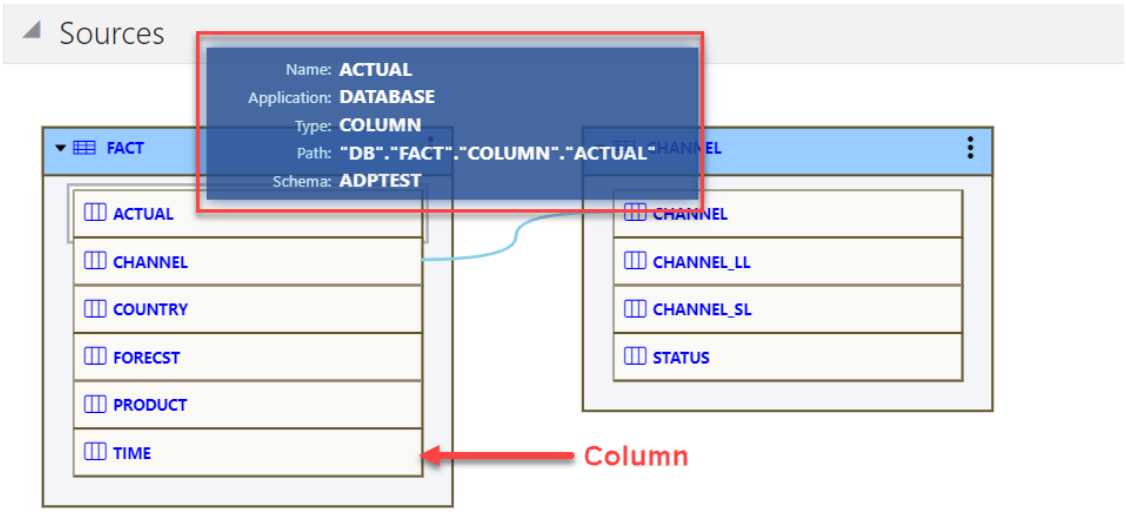
Expand **Joins** to view the **Hierarchy Source**, **Hierarchy Column** and the **Fact column** mapped with the Analytic View. The **Joins** is visible only when the hierarchy table differs from the fact table. You can add multiple join entries for a single hierarchy.

Expand **Sources** to view the fact table associated with the Analytic View. The data model expands to include the data from the source that you added.

Pointing to an item displays the name, application, type, path and the schema of the table. Click the **Actions** (three vertical dots) icon at the right of the item to display a menu to expand or collapse the view of the table.



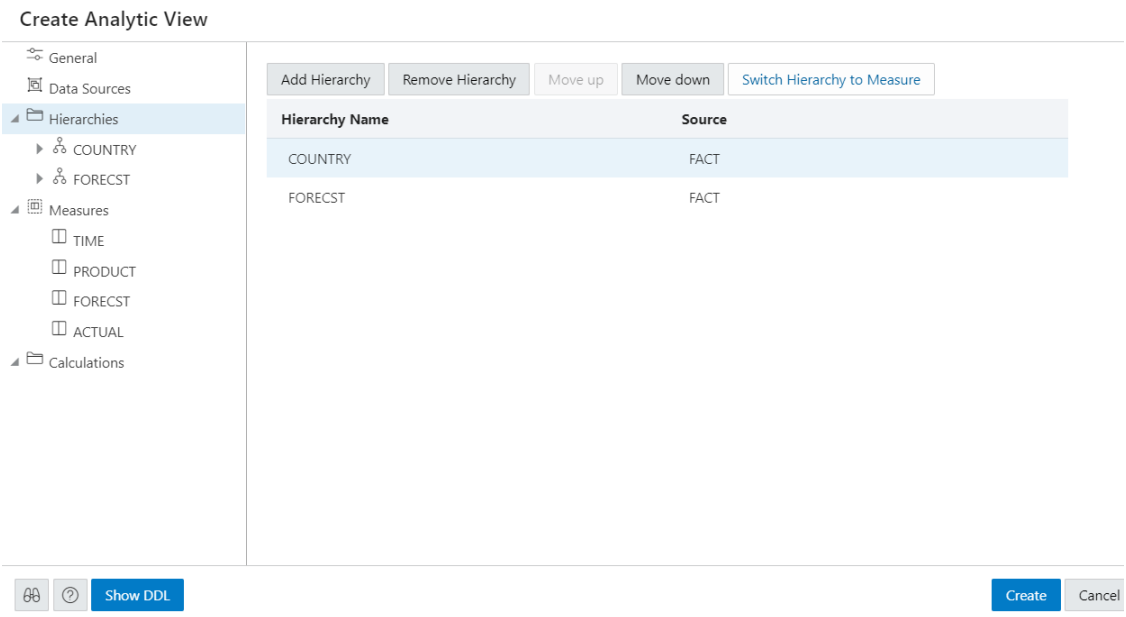
An expanded item displays the columns of the table. Pointing to a column displays the name, application, type, path, and schema of the column.



The lines that connect the dimension tables to the fact table indicate the join paths between them. Pointing to a line displays information about the join paths of the links between the tables. If the line connects a table that is collapsed, then the line is dotted. If the line connects two expanded tables, then the line is solid and connects the column in the dimension table to the column in the fact table.

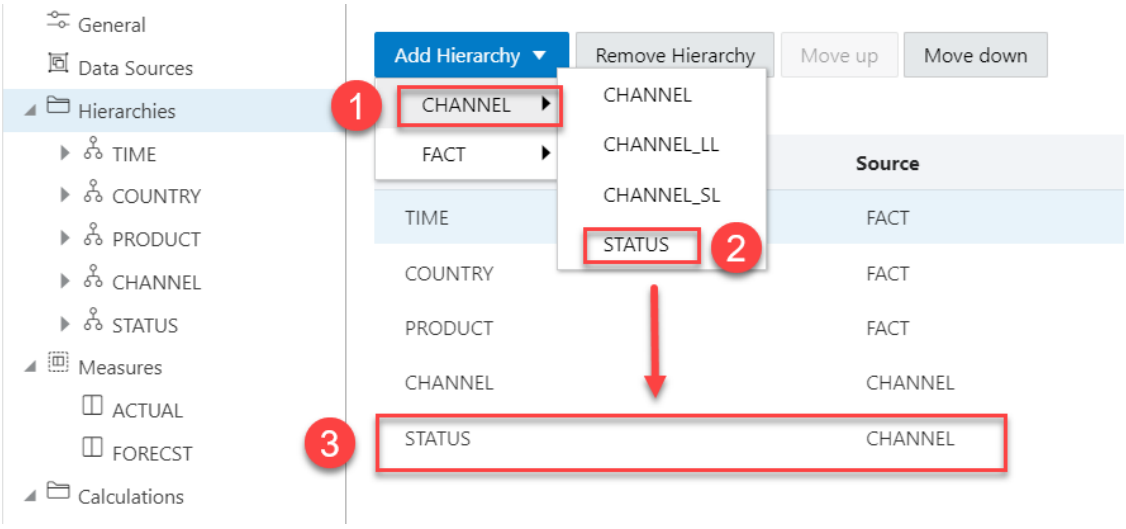
View and Manage Hierarchies

The Hierarchies tab displays the hierarchies generated by the Analytic View creation tool. The display includes the name of the hierarchy and the source table.



An analytic view must include at least one hierarchy.

To add a Hierarchy, click **Add Hierarchy**. This results in a display as a list of column in that table. Select a column that operates as the detailed level of the hierarchy and be the join-key to the fact table.



To remove the hierarchy, select the hierarchy you want to remove from the list and click **Remove Hierarchy**

Select **Move Up** or **Move Down** to position the order of the Hierarchy in the resulting view.

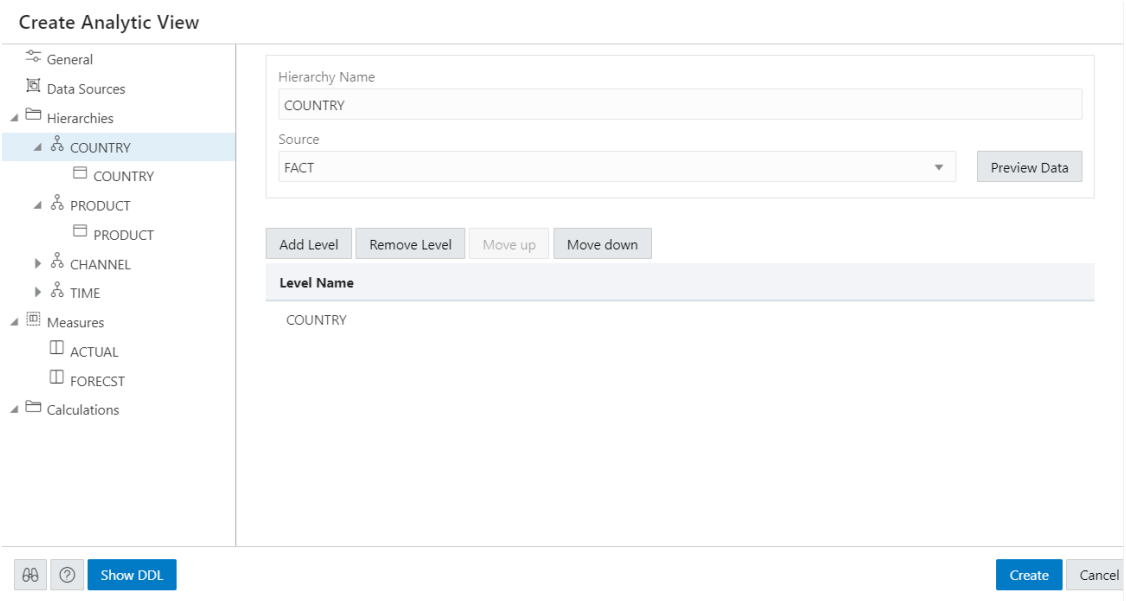
Click **Switch Hierarchy to Measure** to change the hierarchy you select to a measure in the Measures list.

You can also **Add Hierarchy** and **Add Hierarchy From Table** by right-clicking the Hierarchy tab.

If you click on a hierarchy name, a dialog box displays the Hierarchy Name and Source.

To change the source, select a different source from the drop-down list.

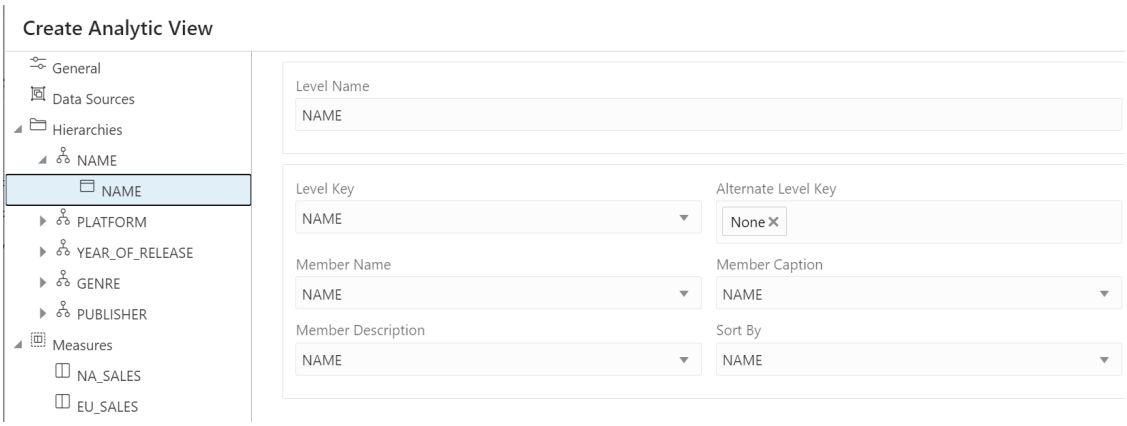
Select **Add Level** to add a level to the hierarchy. Click **Remove Level** to remove the selected level from the hierarchy.



To view the data in the fact table and statistics about the data, click the **Preview Data** button. In the Preview and Statistics pane, the Preview tab shows the columns of the table and the data in the columns. The Statistics tab shows the size of the table and the number of rows and columns.

If you click on a particular level in the Hierarchy tab, a dialog box displays it's respective Level Name, Level Key, Alternate Level Key, Member Name, Member Caption, Member Description, source, and Sort By drop-down. To change any of the field values, enter the value in the appropriate field.

> **Note:**
>
> You can enter multiple level keys Member Name, Member Caption, Member Description and Sort By.



Member Captions and Member Descriptions generally represent detailed labels for objects. These are typically end-user-friendly names. For example, you can caption a hierarchy representing geography areas named GEOGRAPHY_HIERARCHY as "Geography" and specify its description as "Geographic areas such as cities, states, and countries."

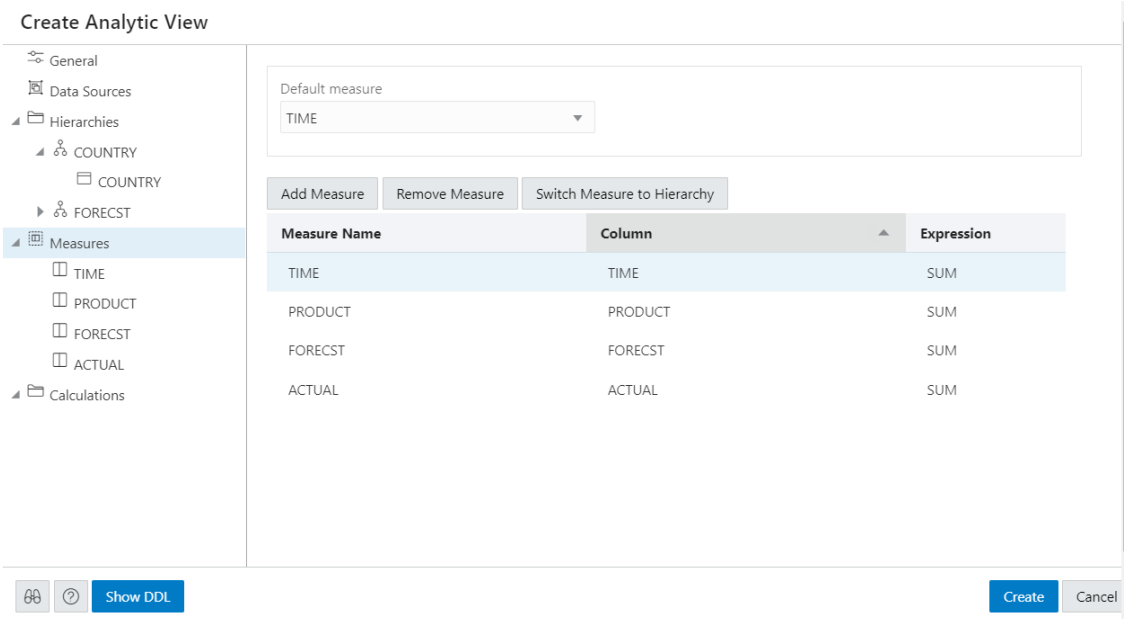
To see the measures for the Analytic View, click **Measures** tab. To immediately create the Analytic View, click **Create**. To cancel the creation, click **Cancel**.

View and Manage Measures

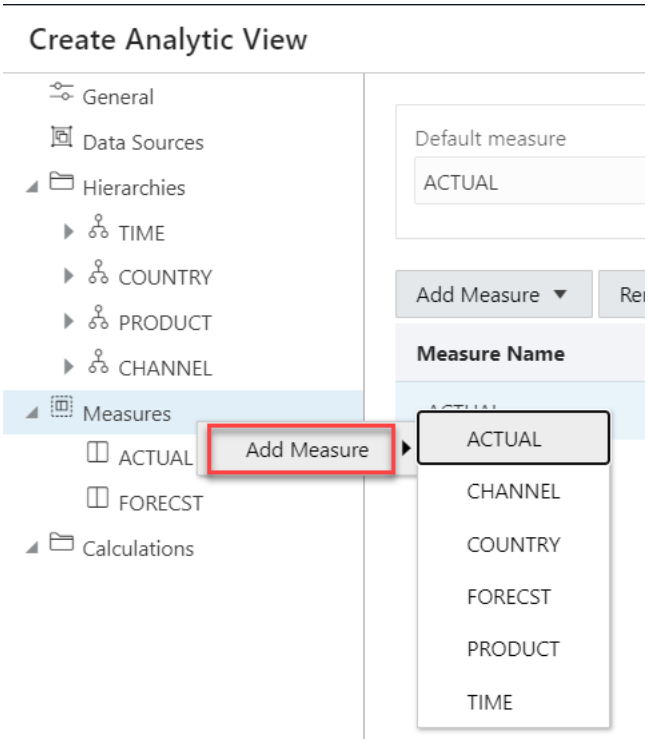
The Measures tab displays the measures suggested for the Analytic View. It displays the Measure Name, Column, and operator Expression for each measure.

The measures specify fact data and the calculations or other operations to perform on the data.

To add measures, click **Add Measure**. You can view a new measure at the bottom of the measures list. To remove the measure, select the measure you want to remove from the list and click **Remove Measure**.



To alternatively add a measure from the data source, right- click the Measures tab. This pops up a list of columns that can be used as measures. Select one measure from the list.



You can exclude a column from the measures on right-clicking the Measures tab and selecting Remove Measure.

Click **Switch Measure to Hierarchy** to change the measure you select to hierarchy in the Hierarchies list.

You must specify a measure as the default measure for the analytic view; otherwise, the first measure in the definition is the default. Select **Default Measure** from the drop-down.

To add a measure, right-click the Measures tab and select **Add Measure**. To remove a measure, select the particular measure you want to remove, right-click on it and select **Remove Measure**.

You can select a different column for a measure from the Column drop-down list. You can select a different operator from the Expression drop-down list.

In creating an analytic view, you must specify one or more hierarchies and a fact table that has at least one measure column and a column to join to each of the dimension tables outside of the fact table.

Create new calculated measures

You can add measure calculations to a query of an analytic view.

The measures and hierarchies associated with the analytic views enable us to create new calculated measures.

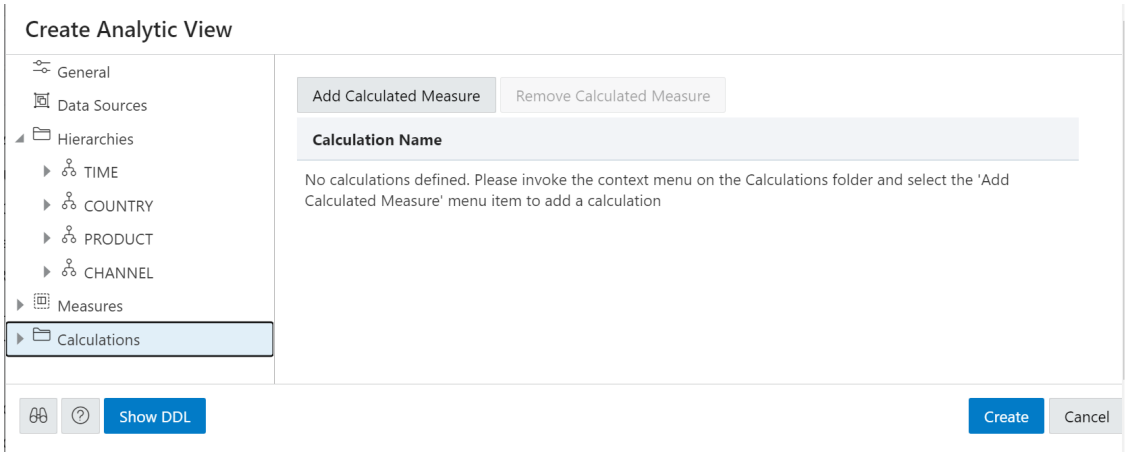
Calculated measures return values from data stored in one or more measures. You compute these measures at run time.

> **Note:**
>
> You can create the measures without increasing the size of the database since the calculated measures do not store the data. However, they may slow performance. You need to decide which measures to calculate on demand.

The Analytic Views provides easy-to-use templates for creating calculated measures.

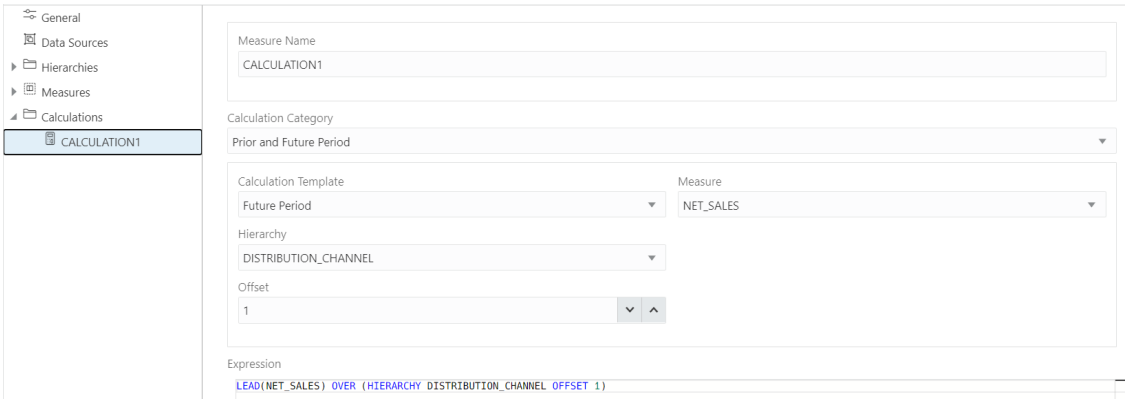
Once you create a calculated measure, it appears in the list of measures of the Analytic View .You can create a calculated measure at any time which is available for querying in SQL.

The Data Analysis tool provides easy-to-use templates for creating calculated measures.



Click **Add Calculated Measure** to add calculations to the measures. You can view the new calculation with system generated name in the **Calculations** tab.

Click the newly created calculated measure.



In the **Measure Name** field, enter the name of the calculated measure.

You can select preferred category of calculation from a list of options such as Prior and Future Period, Cumulative Aggregates, Period To Date, Parallel Period, Moving Aggregates, Share, Qualified Data Reference, and Ranking using the **Calculation Category** drop-down.

Your choice of category of calculation dynamically changes the **Calculation Template**.

For more details on how to use Calculation templates, see Using Calculation Templates.

Select the **Measure** and **Hierarchy** on which you want to base the calculated measures.

Select **Offset** value by clicking the up or the down arrow. The number specifies the number of members to move either forward or backward from the current member. The ordering of members within a level is dependent on the definition of the attribute dimension used by the hierarchy. The default value is 0 which represents POSITION FROM BEGINNING.

The Expression field lists the expressions which the calculated measure uses.

On the creation of the Analytic view, the calculated measure appears in the navigation tree in the Calculated Measures folder.

Click **Create**. A confirmation dialog box appears that asks for your confirmation. Select **Yes** to proceed with the creation of Analytic View.

After creating the Analytic View, you will view a success message informing you of its creation.

On editing the Analytic View you create, you can view the calculated measure in the navigation tree in the Calculations folder.

Click the **Tour** icon for a guided tour of the worksheet highlighting salient features and providing information if you are new to the interface.

Click the **help** icon to open the contextual or online help for the page you are viewing.

Click **Show DDL** to generate Data Definition Language statements for the analytic view.

Edit Analytic View

You might want to edit an Analytic View to make changes to the data sources, the hierarchies, or the measures.

To edit an Analytic View, click the **Action** icon on the Analytic View item, then click **Edit Analytic View**. On the Edit Analytic View screen, select a tab and make changes as desired.

When you have completed the changes, click **Update**.

Working with Analyses

Analyses are a collection of multiple reports on a single page, which provides quick access to multiple data analyses collected from different Analytic Views.

Analyses enable you to monitor performance, create reports and set estimates and targets for future work. It provides you a visual representation of performance with charts and graphs.

You can access the Analyses page by clicking the Analyses tile on the Data Analysis home page.

Viewing Analyses

Analyses provide you an insight into the performance of your data.

You can use the Analysis and the Analyze pane to search or browse Analytic Views, view their analysis, or reports you have access to. Clicking on the Analyses takes you to a page where you can view the Analyze pane. Here you can view default hierarchy level and measures selected. You can drag and drop any levels and measures from the Analytic View browser to rows/columns and Values in the drop area respectively. This defines your analysis criteria. Once the values are dropped, the Data Analysis tool generates a query internally. The tool displays the results of the analysis in the form of reports in the Analyses that matches your analysis criteria. You can add multiple reports to the Analysis. You can also examine and analyze the reports and save them as a new analysis. You can just save the Analysis and not a single report. Once you save all the reports, it will be part of that single Analysis. Reports are unnamed.

Workflow to build Analyses

Here is the workflow to build an analyses.

Following are few common tasks to start building Analyses:

- **Create a useful analysis:** Before creating your first analysis, you can construct a useful analysis over a single Analytic View. This way you can generate analyses on which you can create reports that you display on the Analyses.
- **Create Analysis:** Create an analysis to display data from analysis.
- **Create Reports:** An Analysis can have multiple reports that are independent of one another. This can be used to compare and analyze data generated from different Analytic Views.
- **Save Analysis:** Create customized Analyses that enable you to view reports and their analyses in the current state and save it for future reference.

Format Numeric Columns in the Data Analysis Tool

The custom formatting of numeric columns enables you to select your format type such as **Number**, **Currency** and **Percentage** using Mapping numeric column wizard in the Data Analysis tool.

> **Note:**
>
> Mapping wizard can be viewed only in tabular mode and Chart mode of viewing output.

To map the numeric column:

1. From the Table browser column in the left, drag and drop the numeric column to the measure drop area.
2. Click the respective column on the drop area. The Formatting dialog box appears that enables you to perform changes to the numeric column.
3. The options available depends on the Format Type of the respective column you select. The various format types are **Number**, **Currency** and **Percentage**.
4. The following options are available from the Formatting dialog box when you select **Number** as format type:
   - **Precision**:
     It represents the total number of digits. You can select **Default**, **0.0**, **0.00** or **0.000**.

     For example: To view the number 1.2, map the precision to **0.0**.
   - **Thousand Split**: With this formatting you can insert a comma after every third significant figure.. For example, select **1,234** from the drop-down to view the number as 3,454 instead of 3454.
   - Select **Abbreviate** to view the abbreviated form of the number. For example, you will view 1000 as 1k.
   - **Aggregation**: This column performs aggregation operations such as **Sum**, **Min**, **Max**, **Avg**, **Count** and **Distinct Count** on the selected input column.
5. If your column is of **Currency** Type, you can view the following options:
   - **Select Currency**: You can select various currencies from the drop-down list.

     Select **Highlight Negative Numbers** to highlight negative Numbers.

     You can also view **Precision**, **Thousand Split**, **Abbreviate** and **Aggregation** options.
6. If your column is of **Percentage** Type, you can view **Precision**, **Thousand Split**, **Abbreviate** and **Aggregation** options.

Creating Analyses

Use the Data Analysis tool to create and edit your Analyses. The analysis provides you customized view of Analytic View data. An analysis consists of one or more reports that displays the results of analysis.

You can create a Basepay analysis and add content to track your team's pay. You can view the analysis in a pivot view or tabular view or in the form of charts. You can create an analysis that displays these three views.

In this example, you create a new analysis called New_Analysis.

1. On the Data Analysis home page, click the **Create Analysis** button.
2. You can make changes to existing Analyses by adding different values from hierarchies and measures.
3. The AV name you view on the report represents the AV used to create the report. With a different Analytic View you can create a different report.

> **Note:**
>
> You must have at least one report to build an analysis.

4. To edit the existing analysis, In the Analytic View browser, select the objects to analyze in the navigation pane and drag and drop them to the drop area in Columns, Rows or Values and Filters of the Analyze tab.
5. The report updates based on the artifacts (levels and measures) you select.

The new analysis which contains the updated report will now be visible in the Data Analysis home page for further editing.

- Saving Analyses
  You can save the personalized settings you made for the Analysis and use them on any other Analyses for future reference.

Saving Analyses

You can save the personalized settings you made for the Analysis and use them on any other Analyses for future reference.

You won't have to make these decisions manually every time you open the Analyses page if you save these preferences.

1. Open the analysis for editing from the Data Analysis home page. Select **Edit** from the Analysis tile you wish to edit.
2. Click on the **Save As** icon. Enter a descriptive name for your Analysis.
3. Click **Save**.

Creating Reports

A single report you generate analyzes an AV based on the Levels and measures you select.

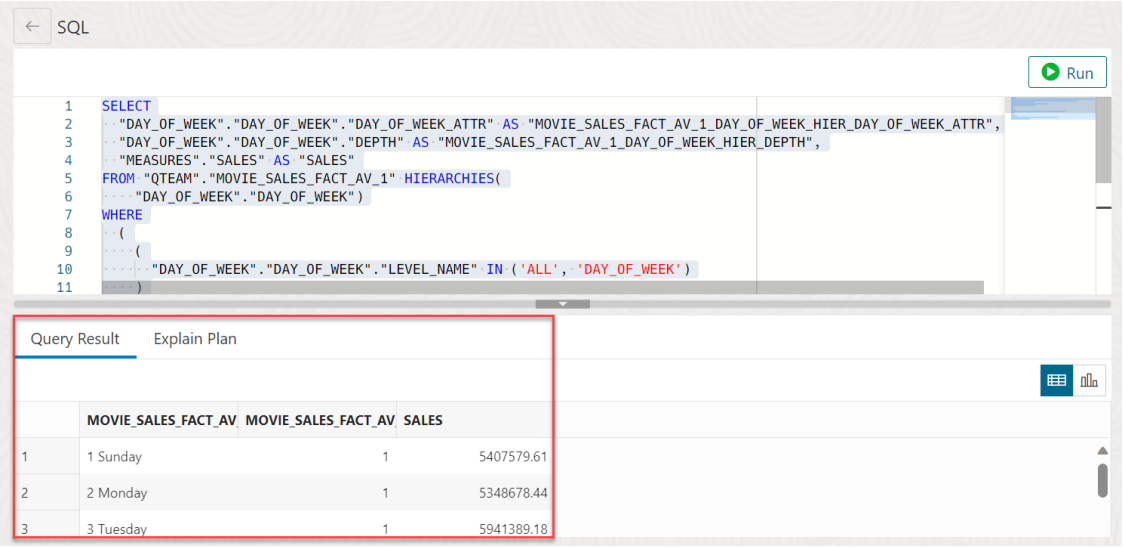
You can add multiple reports to the newly created analysis. One report is independent of another report. To add a report in the Analysis.

1. Open the Analysis for editing from the Data Analysis home page. Select **Edit** from the Analysis tile you wish to edit.
2. Click on the **+ Report** icon to add one or more reports to the Analysis. You can use a report to add configured Analyses to the Analyses page.
3. Click on the report to select it. The resize arrow in the report resizes the report window.
4. Click the cross icon on the selected report to delete the report from the Analysis.
5. The header displays the name of the Analytic View you select.
6. You can expand or collapse the report with their respective arrows.

When you add a report to the Analysis, the report provides the following actions:

- Edit SQL
- Performance
- Rename Report: Click Rename Report to rename it. Click **Save Report** to save the current report.
- Delete Report: Click **Delete Report** to delete the report.

Edit SQL
You can view the SQL output when you click **Edit SQL**. The lower right pane in SQL displays the output of the operation executed in the SQL editor. The following figure shows the output pane on the SQL page.



The output pane has the following tabs:

- **Query Result**: Displays the results of the most recent Run Statement operation in a display table.
- **Explain Plan**: Displays the plan for your query using the Explain Plan command. The default view is the diagram view. For more information, see the description of Explain Plan Diagram in subsequent sections.

Performance menu
The Performance menu displays the PL/SQL procedures in the worksheet area which describes the reports associated with the Analytic Views.

The top part of the performance output consists of the worksheet editor for running SQL statements and an output pane to view the results in different forms. You can view the results in a Diagram View, Chart View, Clear Output from the SQL editor, Show info about the SQL statements and Open the performance menu in a new tab.

The following figure shows the output pane of the performance menu:



The output pane has the following tabs:

- **Diagram View**: Displays the plan of your query in the diagram view.
- **Chart View**: Displays the plan of the query in a chart view.
- **Clear Output**: Clears the PL/SLQ statements from the worksheet.
- **Show info**: Displays the SQL statement for which the output is displayed.
- **Open in new tab**: Opens the explain plan in a new window. An Explain Plan displays the plan for your query.

Explain Plan Diagram
The Explain Plain diagram view is a graphical representation of the contents of the insert row statements in the SQL Query. The plan depicts the hierarchical nature of the steps in the execution plan.

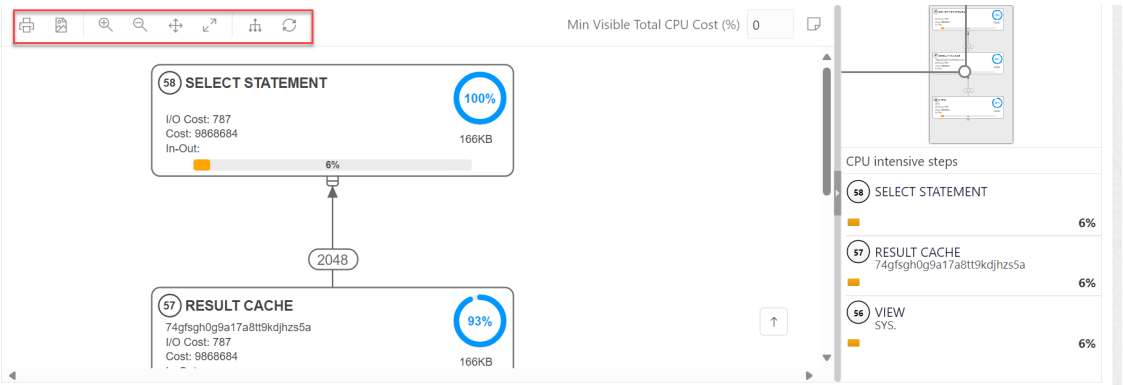
By default, three levels of steps are visible in the diagram. You can use the **+/-** signs at the

bottom of each step (available when the step has children) to expand or collapse. Use Expand All in the toolbar to view all steps in the diagram.
The diagram also provides the following details:

- Cardinality (number on the arrow to the parent step), which is the number of rows processed
- Operation and options applied in that step
- Execution order, which is the sequential number in the order of execution
- Access predicates CPU cost in percentage (orange bar)
- Total CPU cost for the step in percentage (blue circle)
- Estimated I/O Cost, Bytes processed, and Cost metrics

You can see a brief description pop-up when you hover over any of these statistics in a step.

The icons in the toolbar are:



- **Advanced View**: This is the default view of the query when you click Performance. Displays data from SQL Query in mixed tabular/tree view. There is a Diagram View icon that you can use to switch back to the diagram view.
- **Chart View**: Displays data from the SQL query in the form of charts.
- **Print Diagram**: Prints the diagram.
- **Save to SVG**: Saves the diagram to file

- **Zoom In, Zoom Out**: If a step is selected in the diagram, clicking the Zoom In icon ensures that it remains at the center of the screen.
- **Fit Screen**: Fits the entire diagram in the visible area.
- **Actual Size**: Sets the zoom factor to 1.
- **Expand All**: Displays all steps in the diagram.
- **Reset Diagram**: Resets the diagram to the initial status, that is, only three levels of steps are displayed.
- **Show Info**: Shows the SELECT statement used by the Explain Plan functionality.
- **Open in New Tab**: Opens the diagram view in a new tab for better viewing and navigation. The diagram is limited to the initial SELECT statement.
- **Min Visible Total CPU Cost(%)**: Defines the threshold to filter steps with total CPU cost less than the provided value. Enter a value between 0 and 100. There is no filtering for 0.
- **Plan Notes**: Displays the Explain Plan notes.

Properties in Explain Plan Diagram

Double-click or press **Enter** on a selected step to open the Properties slider, which provides more information about that step. See PLAN_TABLE in Oracle Database Reference for a description of each property. The Properties slider shows:

- Displays information for that step extracted from PLAN_TABLE in a tabular format. Nulls are excluded. You can select **JSON** to view the properties in JSON format.
- Information from OTHER_XML column of PLAN_TABLE. The information is displayed in JSON format.

- Working with Reports
  Reports help you in analyzing Analytic Views and Queries.
- Creating Reports on a Query
  This section describes the steps to create reports on an SQL query.
- Creating Reports on an Analytic View
  This section describes the steps to create reports on an Analytic View:

Working with Reports

Reports help you in analyzing Analytic Views and Queries.

The reports are based on the levels and measures you select for the Analytic View and the columns you select for a query.

Click **Analyze** in the Analytic View and click the Table you want to analyze to view the Analyses page.

The Analyses page comprises the following components:

1. **Analytic View Browser**: Select **Analytic View** from the drop-down if you choose to create reports on Analytic Views to view an Analytic View browser. The Analytic View browser displays the Hierarchies, Levels, and Measures associated with the selected Analytic View.

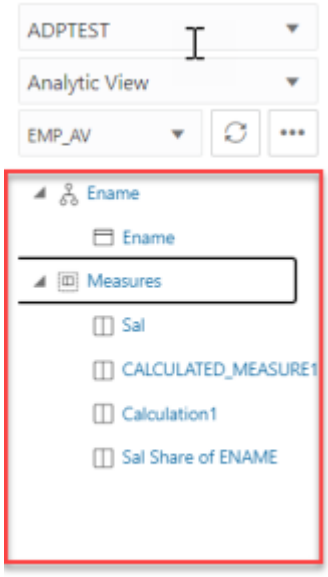
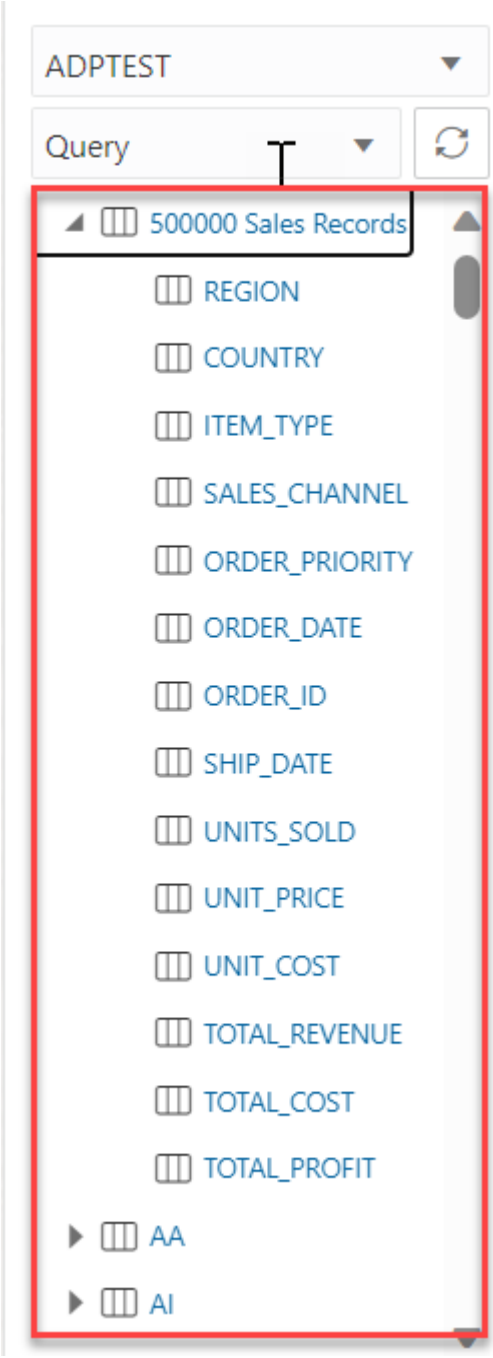
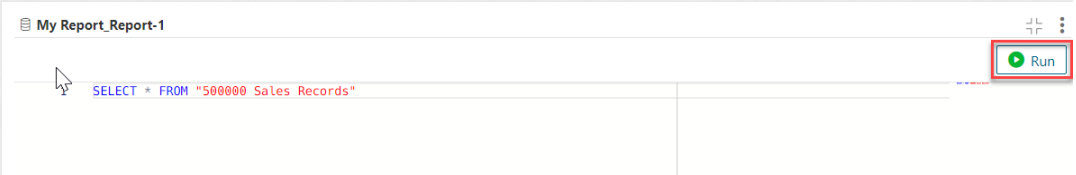


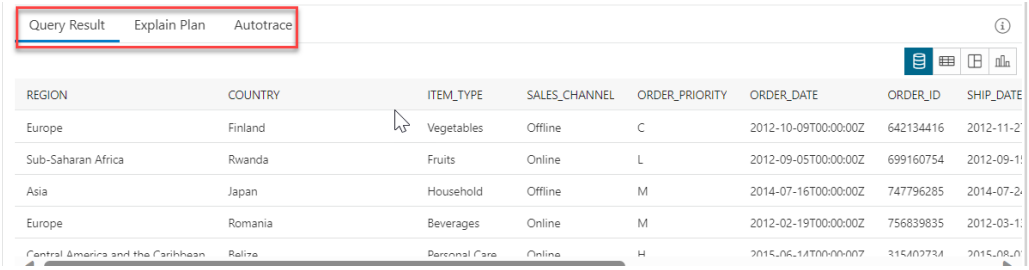
Table Browser: Select **Query** from the drop-down if you choose to create a report on SQL query to view a Table browser. If you select SQL Query, the Table browser displays the available tables and their corresponding columns. You can drill down tables to view their corresponding columns.



2. **SQL Worksheet editor with the Run icon:** You can view this component only when you generate a report on a SQL Query and not an Analytic View.
The SQL editor area enables you to run SQL statements and PL/SQL scripts from the tables you want to query displayed on the Table browser. By default, the SQL editor displays the Select * statement to display all the columns from the first table. Click **Run** to run the statements in the editor.



3. **Output pane:** The output pane, when you view the results of a SQL Query, consists of the following tabs:
   - **Query Result:** Displays the results of the most recent Run Statement operation in a display table.
   - **Explain Plan:** Displays the plan for your query using the Explain Plan command. For more information, refer to the **Explain Plan Diagram** in Creating Reports section.
   - **Autotrace:** Displays the session statistics and execution plan from `v$sql_plan` when running an SQL statement using the Autotrace feature.



4. **Modes of visualization in the Query Result tab**: You can select any of the four modes to visualize the results of the SQL query report you generate.



The four modes of visualization, when you view the reports generated on a SQL query, are:

- **Base Query**: This type of view is by default. Query written in the SQL editor is the Base Query.
- **Table**: You can view the SQL results in tabular form. By selecting this view, a Column drop zone appears which enables you to drag and drop selected columns from the Table browser. By dropping the selected columns in the drop zone, you can view only those columns in the Query Result tab. Select the cross mark beside the Column name to remove it from the drop zone.
- **Pivot**: You can view the results of the SQL query in pivot format. By selecting this format, a Columns, Rows, and Values drop zone appears where you can drag and drop the selected columns, rows or values from the Tables browser.

> **Note:**
>
> Values must be a NUMBER type.

- **Chart**: You can view the SQL results in the form of a chart. By selecting this view an X-axis and Y-axis drop zone appears. Drag and drop selected columns from the Table browser to the drop zone. You must ensure that only the columns with NUMERIC data type can be dropped in the Y axis. Otherwise, the display result would fail with a `Must be a NUMBER type` error. You can add multiple values to the Y-axis. To view the results in the chart view of only a particular y axis, select the Y axis value from the drop-down.

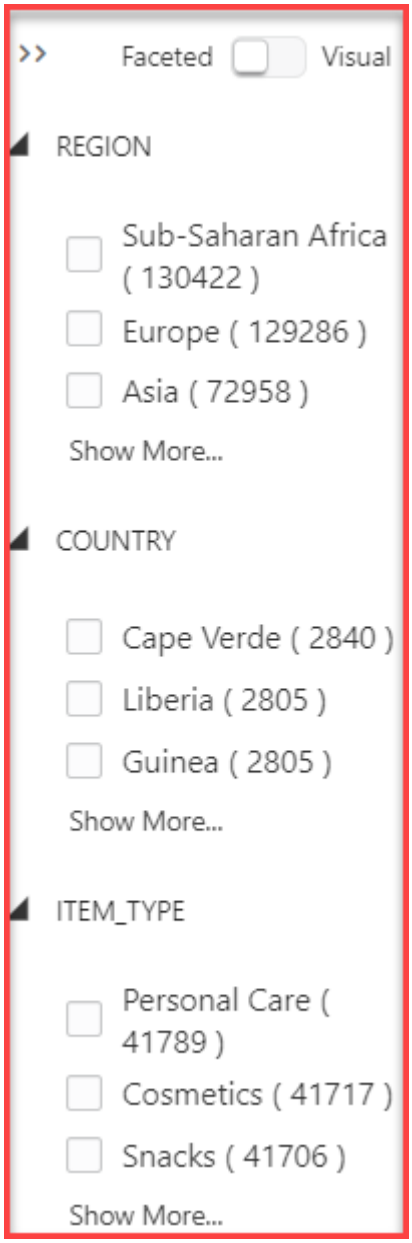
5. **Modes of visualization of reports generated on an Analytic View:** You can select any of the three modes to visualize the results of the report you generate on an Analytic View.



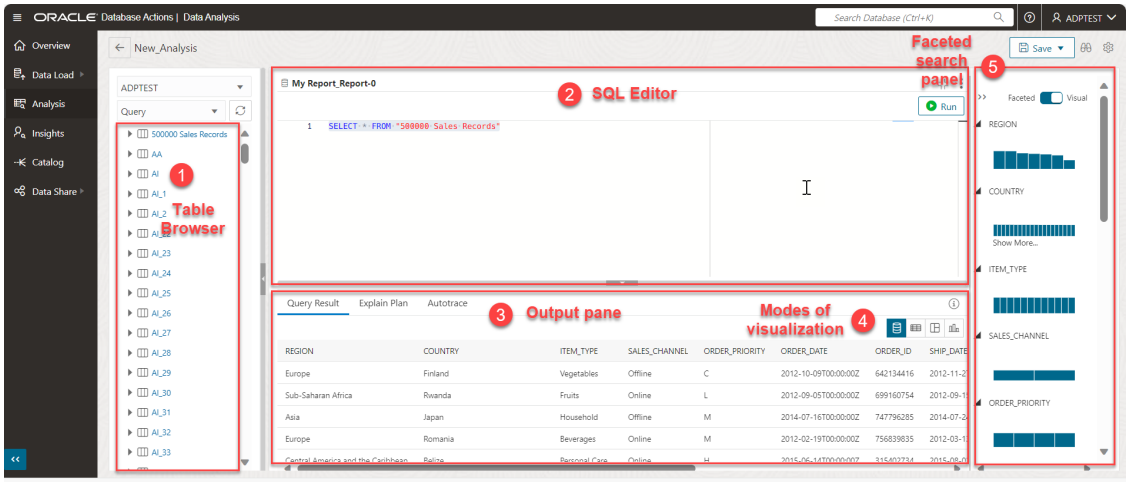
The three modes of visualization when you view the reports generated on an Analytic View are:

- **Table:** You can view the SQL results in tabular form. By selecting this view, a Rows and Filters drop zone appears which enables you to drag and drop selected Hierarchies and Measures from the Analytic View browser. This way you can view the report results that consist of the selected hierarchies and measures.
- **Pivot**: You can view the results of the Analytic View report in the pivot format. By selecting this format, a Columns, Rows,Values and Filters drop zone appears where you can drag and drop the selected hierarchies and measures from the Analytic View browser.Note: Values must be a NUMBER type.
- **Chart**: You can view the report you generate on an Analytic View in the form of a chart. By selecting this view an X-axis, Y-axis, and Filters drop zone appears. Drag and drop selected columns from the Table browser to the drop zone. You must ensure that only the columns with NUMERIC data type can be dropped in the Y axis. Otherwise, the display result would fail with a `Must be a NUMBER type` error. You can add multiple values to the Y-axis. You can select Horizontal and Vertical from the drop-down to view Horizontal and Vertical Charts respectively. You also have the option to select Area Chart, Bar Chart, Line Chart, and Pie Chart from the drop-down.

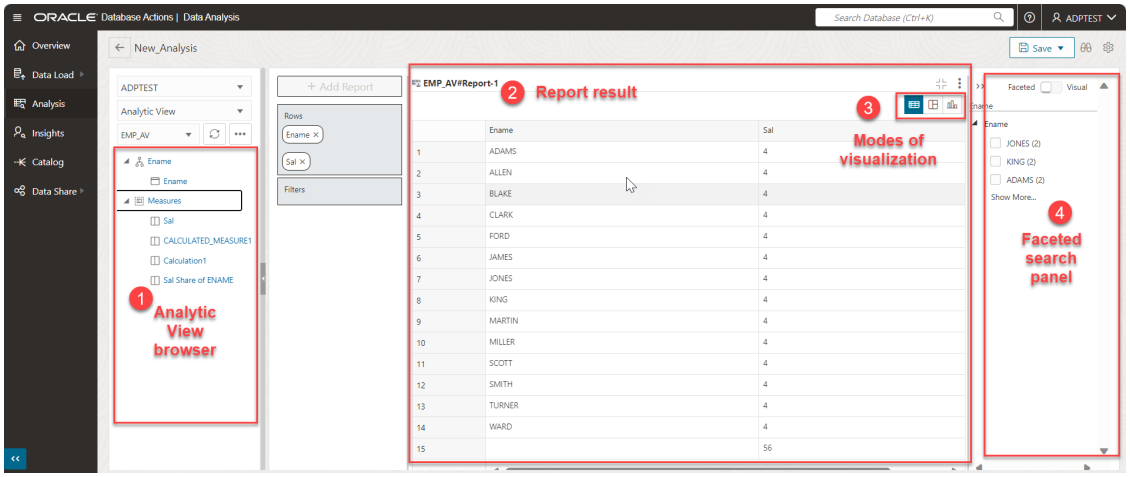
6. **Faceted search panel:** For the reports you generate on a SQL query and an Analytic View, you can view a Faceted search column. For a SQL report, this panel allows you to add filters to the report. The tool generates a filter for each value in the column that is retrieved from the query result. You can filter different columns on the faceted search panel and view the results in the Query result to get only the data you wish to view. You can view the data retrieved from the SQL query in either text or visual format. For reports you generate on an Analytic View, select **Faceted** from the radio button. This filter behaves differently than the Faceted search you generate on an SQL Query. See *Adding filters to a report you generate on an Analytic View*.



The Analyses page when you create a report on an SQL query looks like this:



The Analyses page when you create a report on an Analytic View looks like this:



The following topics describes how to create a report and access the Analyses page:

Creating Reports

1. You can create Reports using either of the following ways:
   - On a Query
   - On an Analytic View
2. Adding filters to a report you generate on an Analytic View

Creating Reports on a Query

This section describes the steps to create reports on an SQL query.

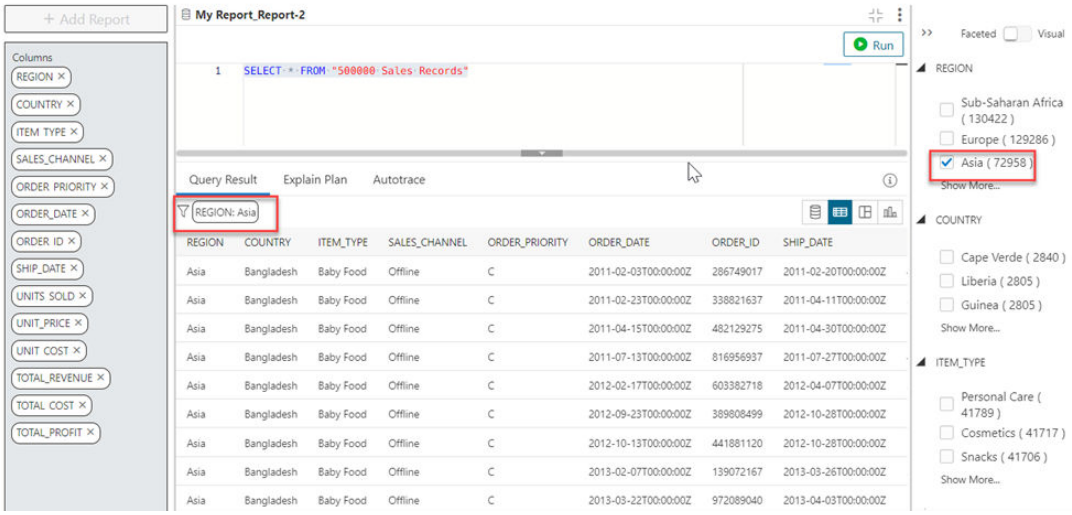
1. From the Analysis home page, select any of the Tables you want to create a report on. You will view the Analyses page with a default query displayed on the SQL editor.

> **Note:**
>
> By default, you will view “Select * from the <Tablename> you select.

2. Click **Run** to run the SQL statement.

   The Query Results tab displays the result in whichever mode you select. The default view is Base Query.

3. Add a filter to the report by displaying the Sales Records of only **Asia** region. Select **Asia** from the faceted search panel.



4. You will view a funnel icon in the Query Result tab which displays a filter with the **REGION** column as **Asia**. The Query Result will display only the records with **REGION** as **Asia**.

Creating Reports on an Analytic View

This section describes the steps to create reports on an Analytic View:

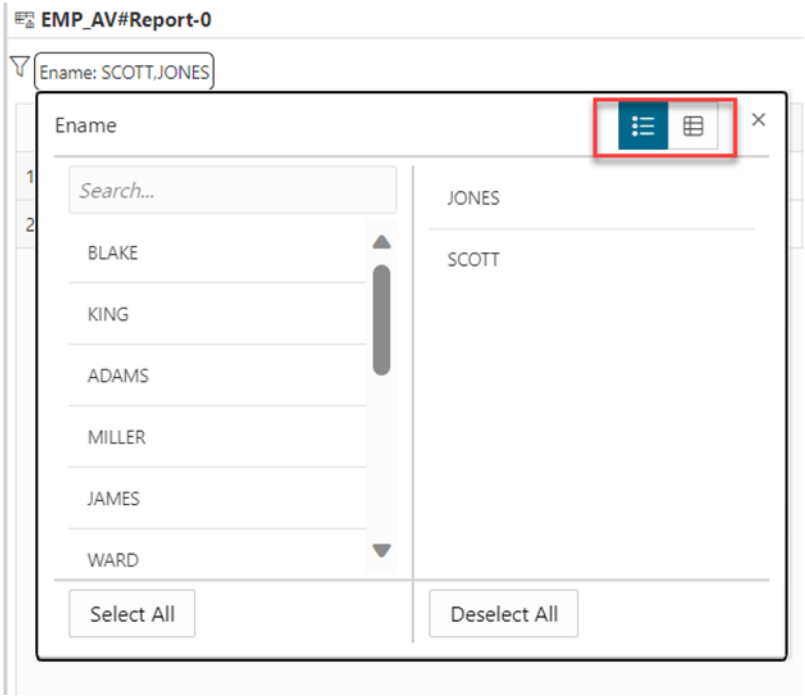
1. From the Analysis home page, under the Analytic View section, select any of the Analytic Views you want to create a report on. You will view the Analyses page with a default report displayed as the output.
2. Drag and drop hierarchies and measures from the Analytic View browser to edit the results you view in the output. For more information, refer to Working with Reports.
3. Click **Expand Report** to expand the view of the report and click **Collapse Report** to minimize the view of the report. The default view of the report you generate on an Analytic View is Pivot.

Adding filters to a report you generate on an Analytic View

Let us add a filter to the report you generated on an Analytic View. Let's say you wish to view the salary of an employee named **SCOTT**.

1. From the Analyses page which displays the report you generated on an **Employee Analytic View**, select “**Scott** from the faceted search panel. You will view the report displaying the result of employees named **Scott**.
2. You can select more filters by selecting the values from the faceted search panel, or by clicking the funnel.

3. Clicking the funnel icon displays all the values of the Employee name column. Select **Jones** to filter the report results further displaying the salary of employees named **Scott** and **Jones**. You can view the values in a list view or a multi select view.



4. Select **Deselect All** to remove all filters.

You can now view the original report result that does not consist of any filters.

Run Natural Language Query in the Data Analysis Tool

You can query the Oracle Autonomous Database by using Natural Language Query rather than having to write SQL Query.

Prerequisites

To utilize the Natural Language Query, you must have:

- An OpenAI, or Cohere or An Azure OpenAI account service with credit
- An access to `DBMS_CLOUD_AI` package.

> **Note:**
>
> - You must set your AI profile using the Natural Language Query feature. Follow the steps mentioned in this chapter to Use Select AI to Generate SQL from Natural Language Prompts.
> - After you have created and configured your AI profile, set your AI profile in Data Studio Preferences wizard on the Connections page to use AI features such as generating SQL from Natural Language in the Data Studio tool.

You can run Natural Language Queries on both Analytic Views and Tables.

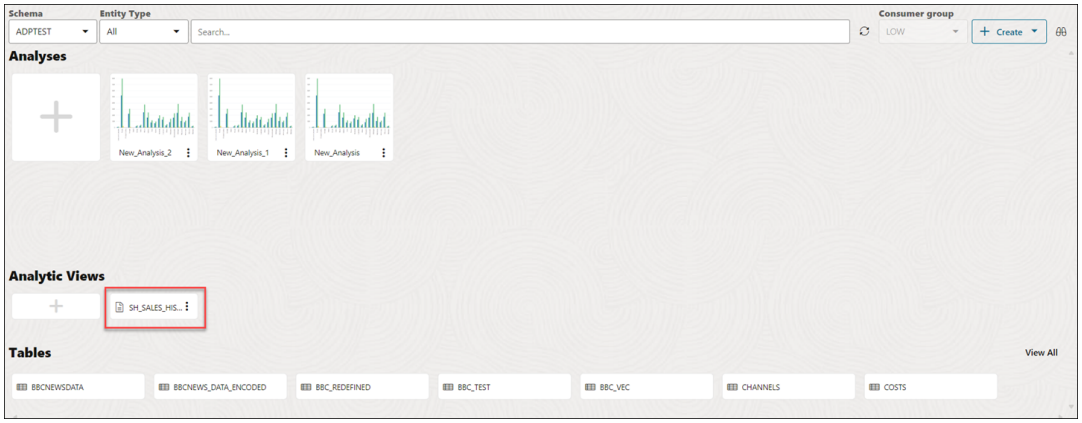
- Generate SQL Queries From Natural Language on Analytic Views
  You can create SQL queries on Analytic Views.
- Generate SQL Queries From Natural Language on Tables
  You can create SQL queries on Tables.
- Generate SQL Queries From Natural Language on Analytic Views
  To run a natural language query from SQL query on Analytic Views, perform the following steps:
- Generate SQL Queries from Natural Language on Tables
  Follow the procedure mentioned below to generate SQL queries from Natural Language on Tables:

Generate SQL Queries From Natural Language on Analytic Views

To run a natural language query from SQL query on Analytic Views, perform the following steps:

Let's say you wish to view sales amount in a categorized way.

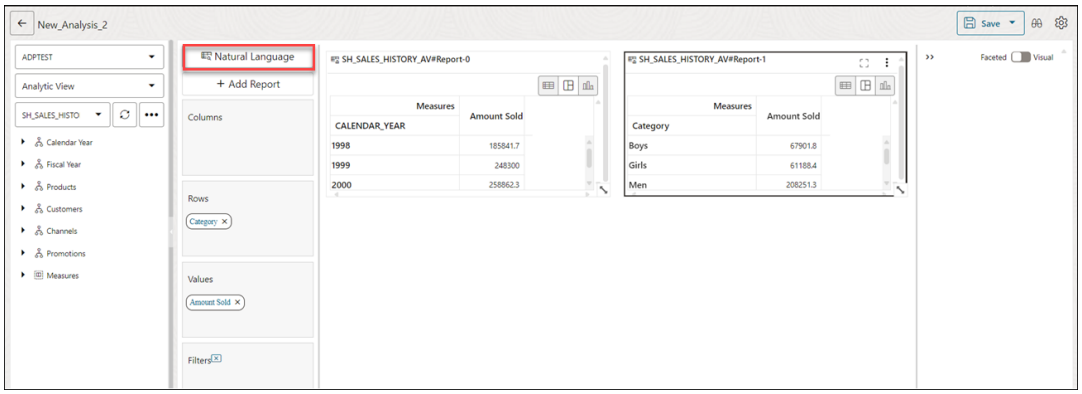
1. On the Data Analysis home page, click on any of the Analytic Views you wish to query.



In this example, you will query the `SH_SALES_HISTORY` Analytic View.

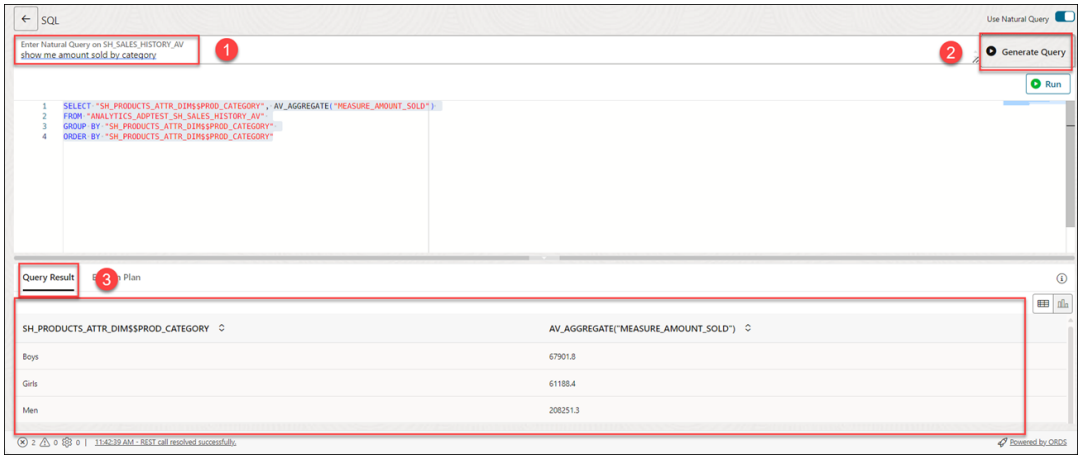
This opens the Analyses page.

2. Select **Natural Language**. Selecting Natural Language opens the SQL worksheet area with a predefined SQL query on the worksheet area.



> **Note:**
>
> You can view **Natural Language** option only after you have configured and set AI profile using the Use Select AI to Generate SQL from Natural Language Prompts procedure and set the Data Studio Settings wizard on the Connections page.

3. Enter the following natural Query you wish to run on `SH_SALES_HISTORY` in the Natural Query field: `show me amount sold by category`.
4. Click **Generate Query**.



5. After the tool generates the query, the **Query Result** tab displays the result of the query. You can also view the graphical representation of the contents of the `PLAN_TABLE` in the **Explain Plan** tab.

   You can alternatively view the query result in chart view by switching the display mode in **Chart View or Diagram View**.

Generate SQL Queries from Natural Language on Tables

Follow the procedure mentioned below to generate SQL queries from Natural Language on Tables:

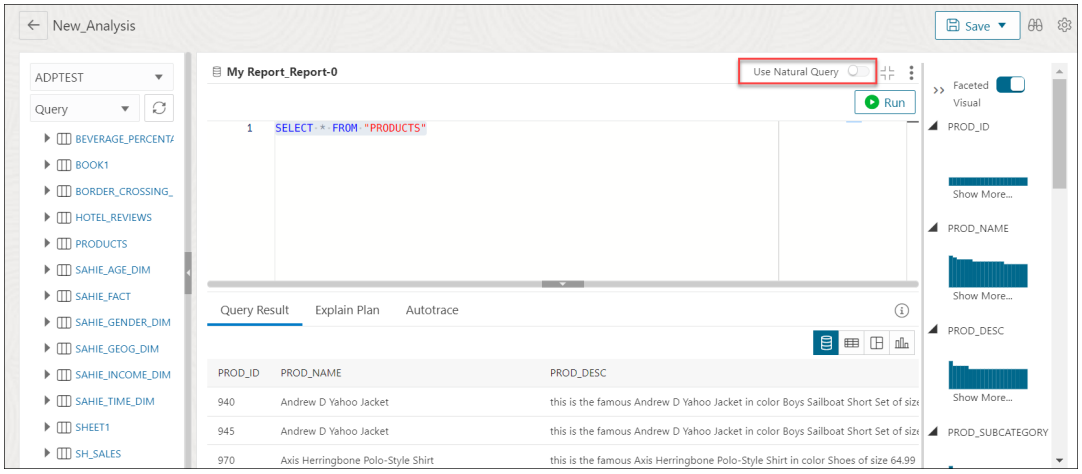
Let's say you want to view the product with the minimum product price.

1. On the Data Analysis home page, click on any of the Tables you wish to query. In this example, we select the `PRODUCTS` table.



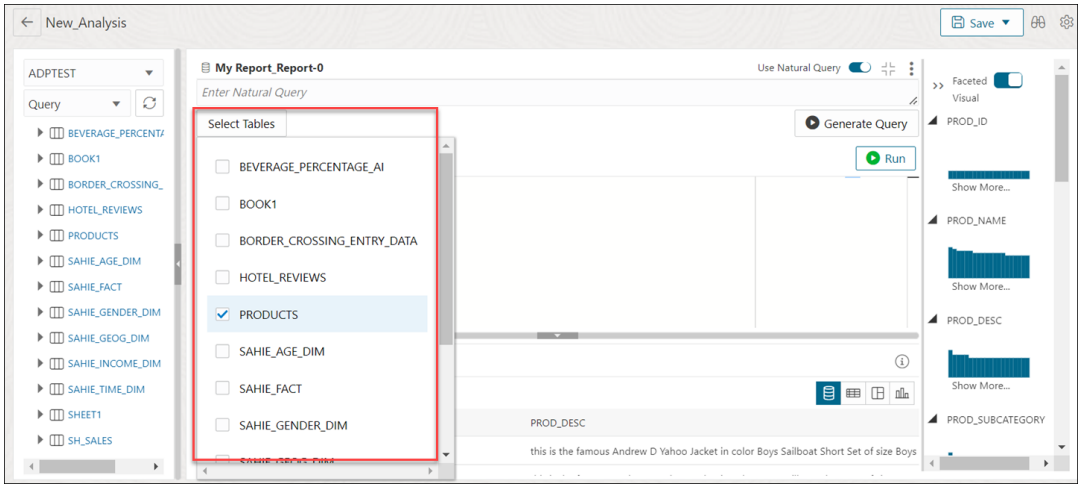
This opens the Analyses page with a query that retrieves all the columns from the selected table `PRODUCTS`.

2. Select **Use Natural Query**.



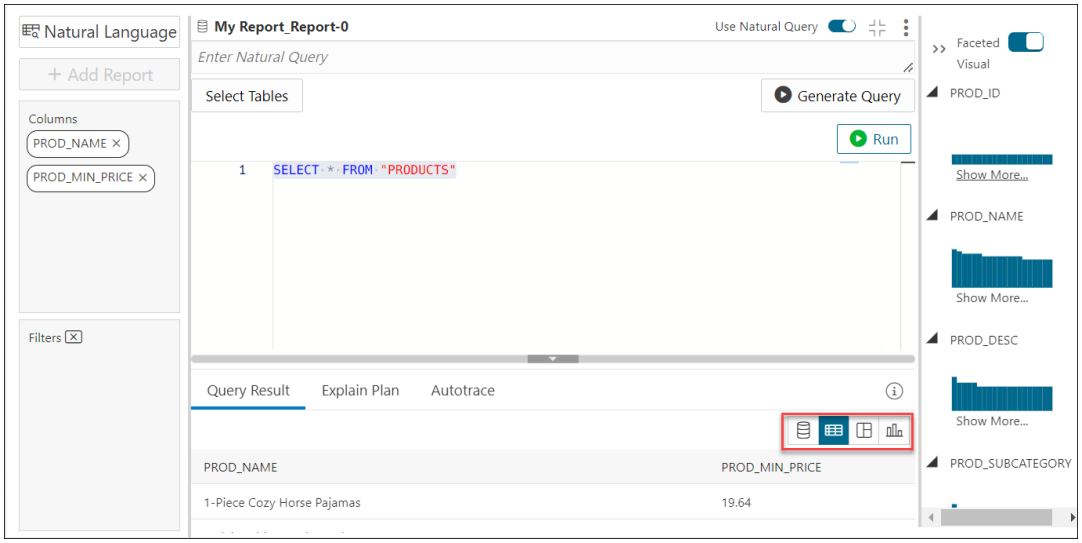
3. Selecting **Use Natural Query** lets you select multiple tables from the **Select Tables** icon.

> **Note:**
>
> You can view **Natural Language** option only after you have configured and set AI profile using the Use Select AI to Generate SQL from Natural Language Prompts procedure and set the Data Studio Settings wizard on the Connections page.

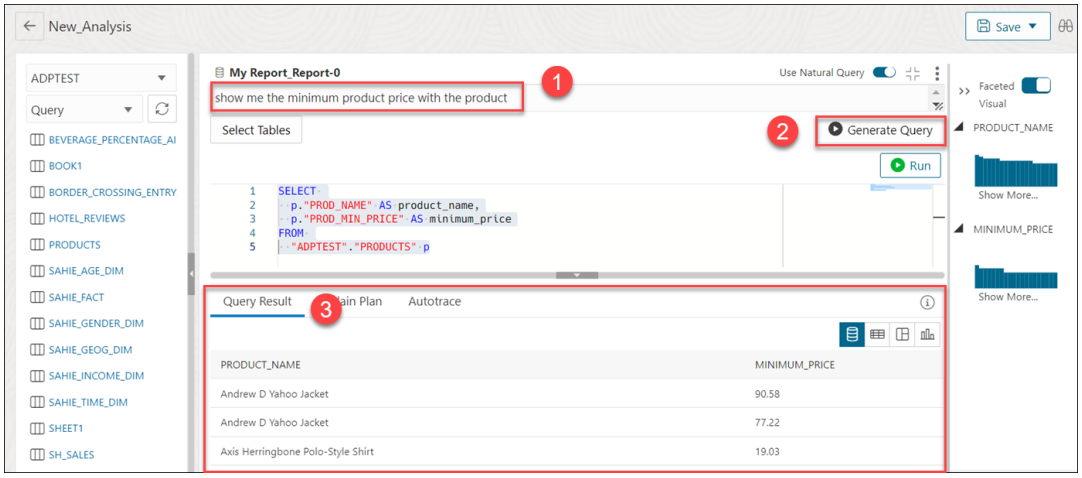


Select the tables you wish to generate the SQL query from. You can use the columns from the selected table in the text field where you enter the Natural Query.

> **Note:**
>
> You can click Tabular View, drag and drop columns from the navigator to the Columns and Filter drop area to select the intended columns to query.



4. Click **Base Query** mode of visualization to enter Natural Language.
Enter the following in the Natural Query Text field: `show me the minimum product price with the product.`



Click **Generate Query**.

Selecting **Generate Query** converts the Natural language to the equivalent SQL query and displays results under the Query Results tab.

You can alternatively view the query result in tabular view, pivot, and chart view. You can drag and drop rows, columns, and filters from the Tables Browser to the drop area.

> **Note:**
>
> From the Chart view you can view the result in horizontal and vertical sheet.

Generate Natural Language explanation from SQL Query using AI Explain

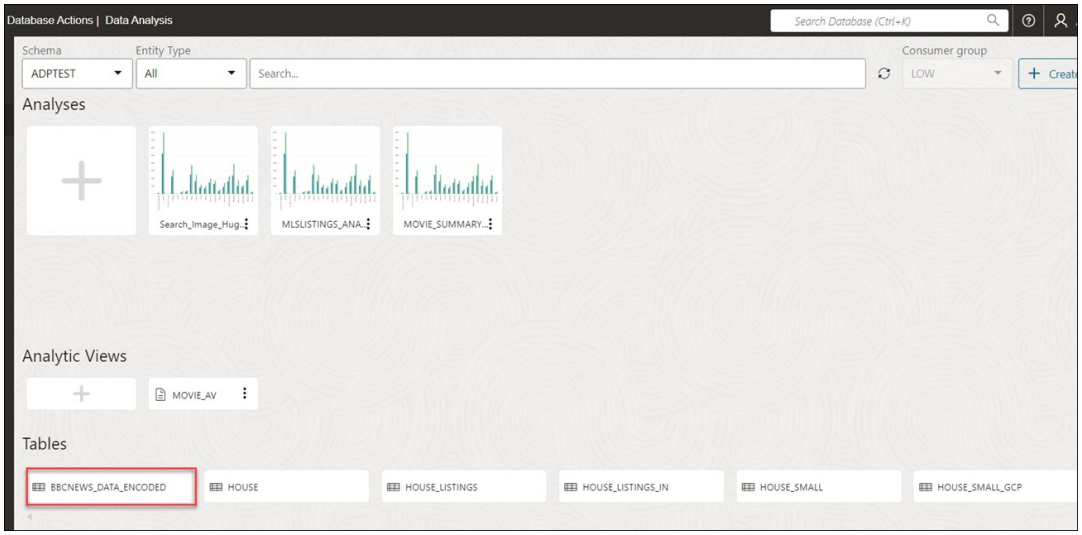
You can translate SQL queries to natural language that is understood by you.

> **Note:**
>
> This feature is available for SQL based reports on Queries and not on reports on Analytic Views.

To run an SQL query and view it's Natural Language explanation, perform the following steps:

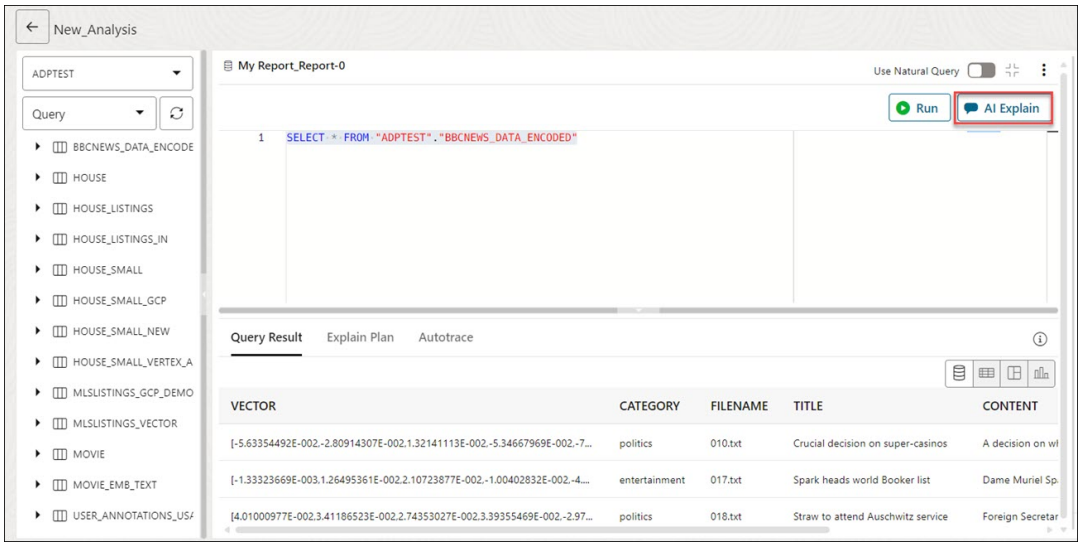
Let's say you wish to view the explanation of a query that retrieves all records of a table.

1. On the Data Analysis home page, click on any of the Tables you wish to query.

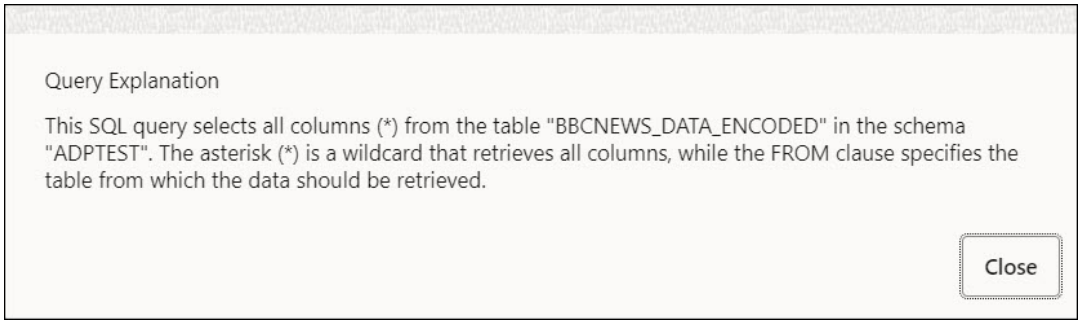


In this example, you will query the `BBCNEWS_DATA_ENCODED` table. This opens the Analyses page.

2. You can view the default query displayed on the worksheet area that retrieves all the records of the table along with the columns displayed in the Query Results tab of the output pane.



3. Click **AI Explain**.



You can view a Query Explanation dialog box that explains the query displayed on the worksheet area.

Click **Close** to close the dialog box.

Vector Search in the Data Analysis Tool

You can use Oracle AI Vector Search capabilities to search for relevant text from your source data on a specific column.

The Vector search helps quickly locate and match data that is similar to the text you provide as your search input.

Prerequisites

For using the vector search feature in the Data Analysis tool, you need the following:

- You must have 23ai Database with Vector support.
- The ORDS REST APIs must support VECTOR data type.
- You must use the Data Studio Settings page to select the same credential you use for vector search.
- The source data must already have a non-numeric column to be embedded first.
- To utilize the Vector Search in the Data Analysis tool, you must have:
  - An OpenAI, or Cohere or An Azure OpenAI account service with credit
  - An access to `DBMS_VECTOR` and `DBMS_VECTOR_CHAIN.UTL_TO_EMBEDDING` package.

> **Note:**
>
> - You must set your AI profile. Follow the steps mentioned in this chapter to Use Select AI to Generate SQL from Natural Language Prompts.
> - After you have created and configured your AI profile, set your AI profile in the Data Studio Settings wizard on the Data Studio navigation menu to use AI features such as generating SQL from Natural Language and Vector Search in the Data Studio tool.

The Data Analysis tool utilizes the Vector Utility PL/SQL package `DBMS_VECTOR` and `DBMS_VECTOR_CHAIN.UTL_TO_EMBEDDING` to provide the third-party REST APIs that let you interact with external embedding models such as Cohere, Google AI, Hugging Face, Oracle

Cloud Infrastructure (OCI) Generative AI, OpenAI, or Vertex AI. You must understand the terms of using third-party embedding models.

> **Note:**
>
> Certain features of the database may allow you to access services offered separately by third-parties, for example, through the use of JSON specifications that facilitate your access to REST APIs.
>
> Your use of these features is solely at your own risk, and you are solely responsible for complying with any terms and conditions related to use of any such third-party services. Notwithstanding any other terms and conditions related to the third-party services, your use of such database features constitutes your acceptance of that risk and express exclusion of Oracle's responsibility or liability for any damages resulting from such access.

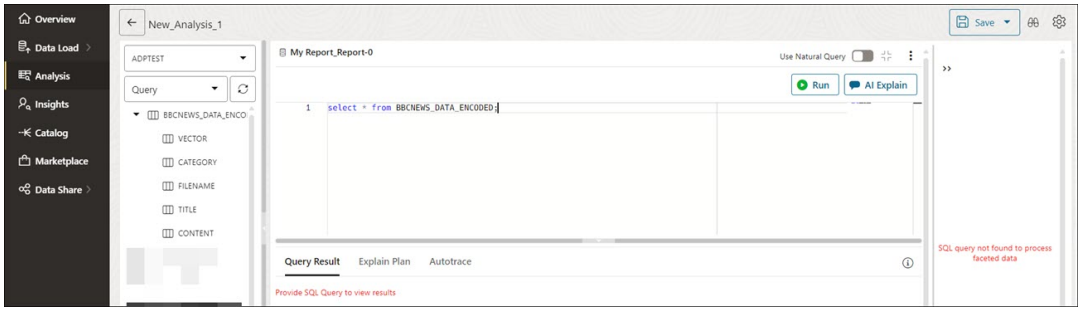
Suppose you are a finance enthusiast and you want to learn about finance through news to gain better sense of financial products and investments. In this example, you can use the Vector Search feature to find news related to the word `Bank`.

Consider a table that has the columns `CATEGORY`, `FILENAME`, `TITLE`, `CONTENT` and a different column labeled `VECTOR`, that contains the vector representation of the `CONTENT` column. Since we can use vector search on source columns with non-numeric data (image or text), any of the non-numeric column can be embedded and stored as a vector column. But in this example, we need to search for relevant text from `CONTENT` column.

You can use the Data Analysis tool to perform vector search.

To indicate which column's embedding is kept in the vector column, we must first map the appropriate column with the appropriate vector using the Data Analysis tool, search for similar text, and then run the query. The Query Result tab will display the content similar to `Bank`.

1. From the Data Analysis tool, click **+** under Analyses to create a new Analysis.
2. Select the Schema from **Schema** drop-down and Query from the **Type** drop-down.
3. Select the table you want to query from the table browser. In this example, you will select `BBCNEWS_DATA_ENCODE`.

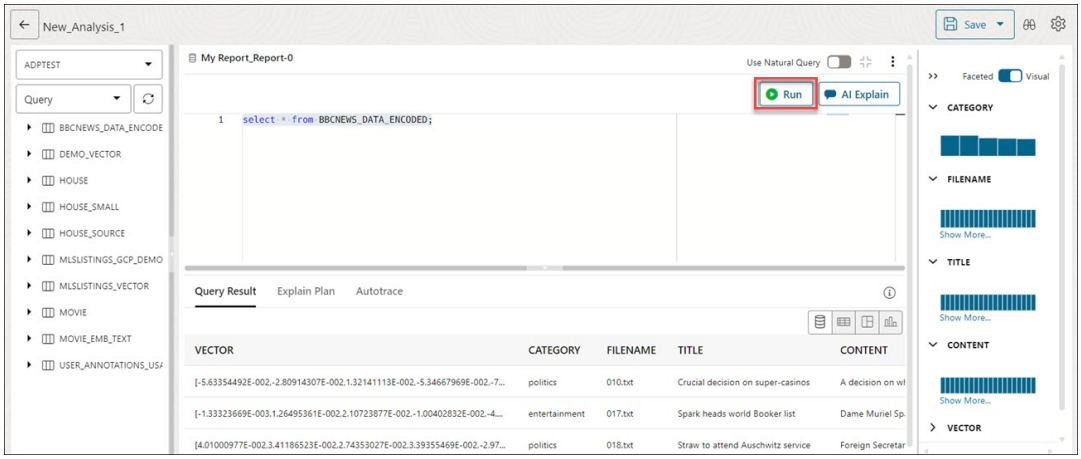


4. Enter the Select statement in the SQL Worksheet edit area to retrieve all the columns from the `BBCNEWS_DATA_ENCODE` table.

Specify the following command:

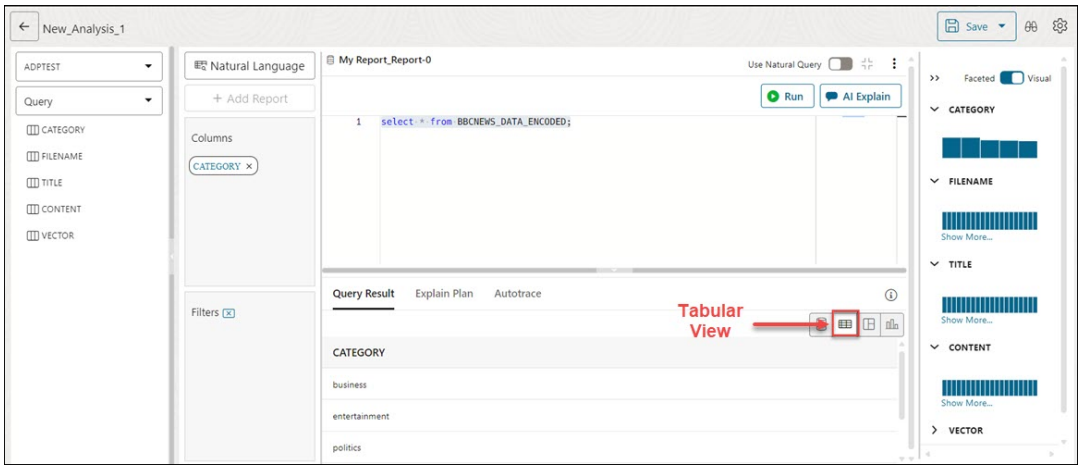
```
Select * from BBCNEWS_DATA_ENCODE;
```

5. Click **Run**.

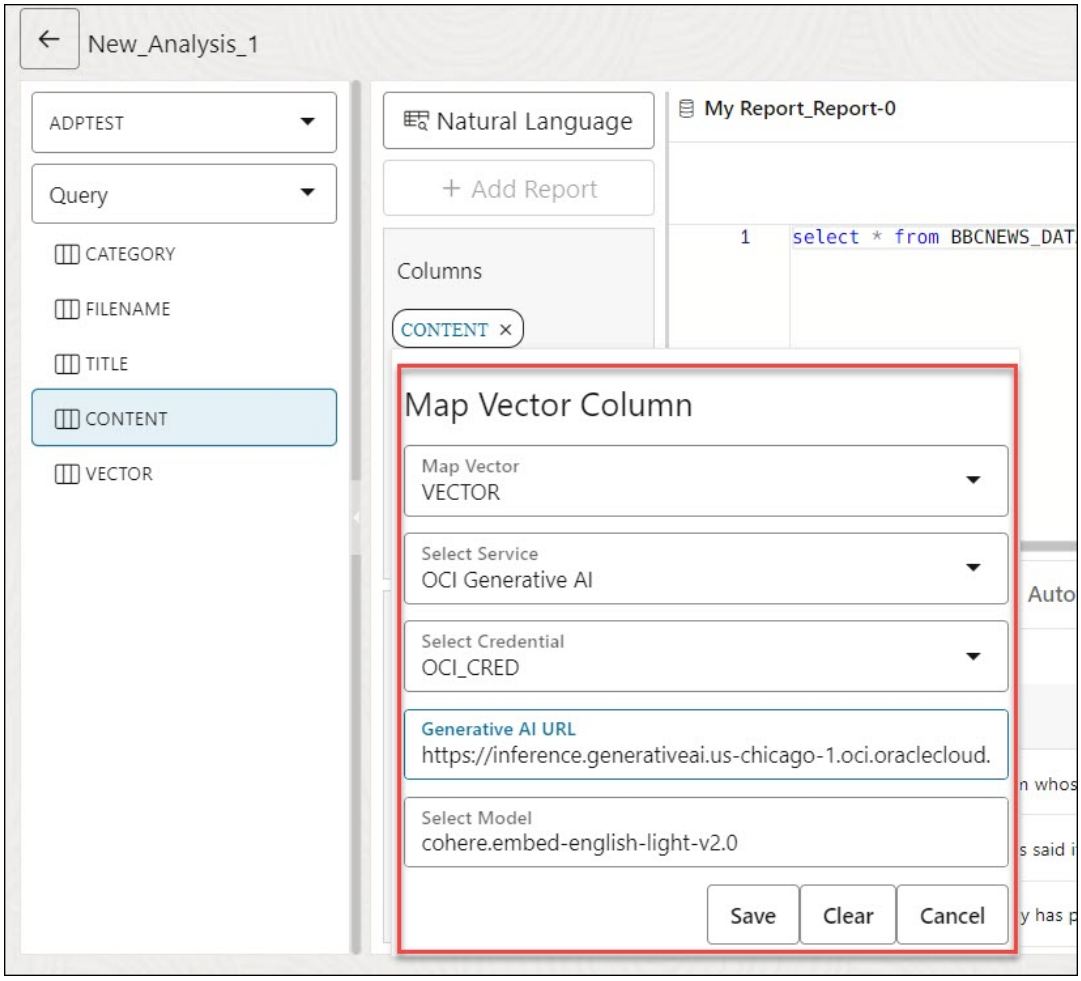


You can view all the columns of the `BBCNEWS_DATA_ENCODE` table in the Query Results tab.

6. Select tabular view of mode from the icons that display various modes of visualizing the query result.

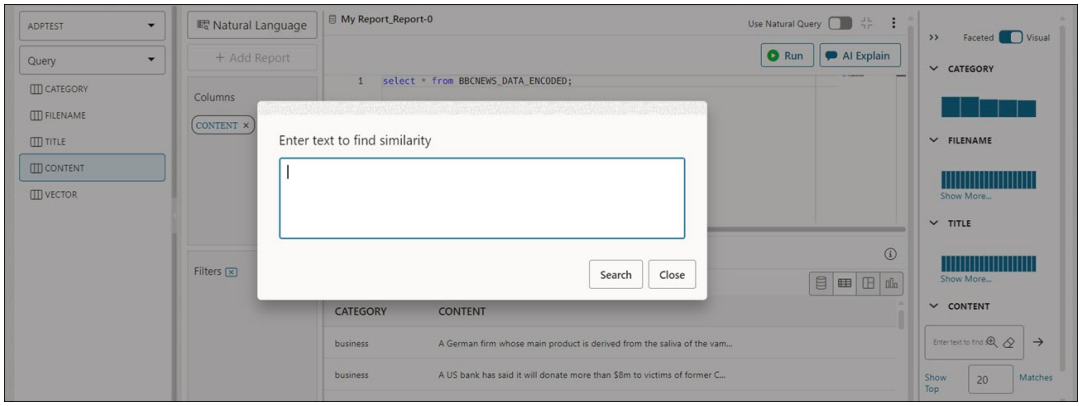


7. Drag and drop the `CONTENT` column to the Column drop area and click it to map the vector column.



8. On the Map Vector Column wizard, specify the following field values:
   - **Map Vector**: Select the column that consists of the vector embeddings of the selected column from the drop-down list. In this example, it is *VECTOR*.
   - **Select Service**: Select the AI service you will use to utilize this feature of the tool. In this example, it is *OCI Generative AI*.
   - **Select Credential**: Enter the credential you create from the Data Studio Settings to access the third party service provider. You require these credentials to enable access during REST API calls to your chosen third-party service provider.
   - **Generative AI URL**: Specify the Generative AI URL you will use to call the third-party service provider, such as Cohere, Google AI, Hugging Face, Oracle Cloud Infrastructure (OCI) Generative AI, OpenAI, or Vertex AI.
   - **Select Model**: Specify the embed model you will use to generate embeddings from the text. In this example, you will enter *cohere.embed-english-light-v2.0*.

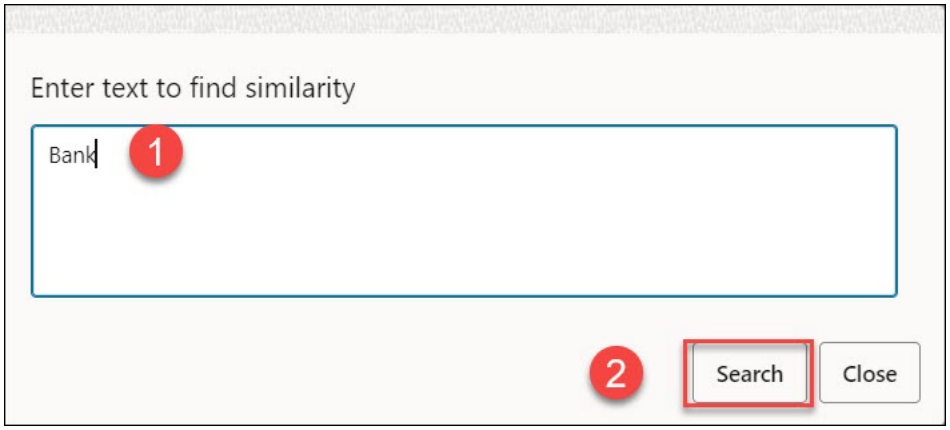
   Click **Save** to save the mapping. Select **Clear** to clear the field values you mention in this step.
9. On the Faceted Filters, you can view text box under the Content column. Click the **+** sign on it to expand the text box.



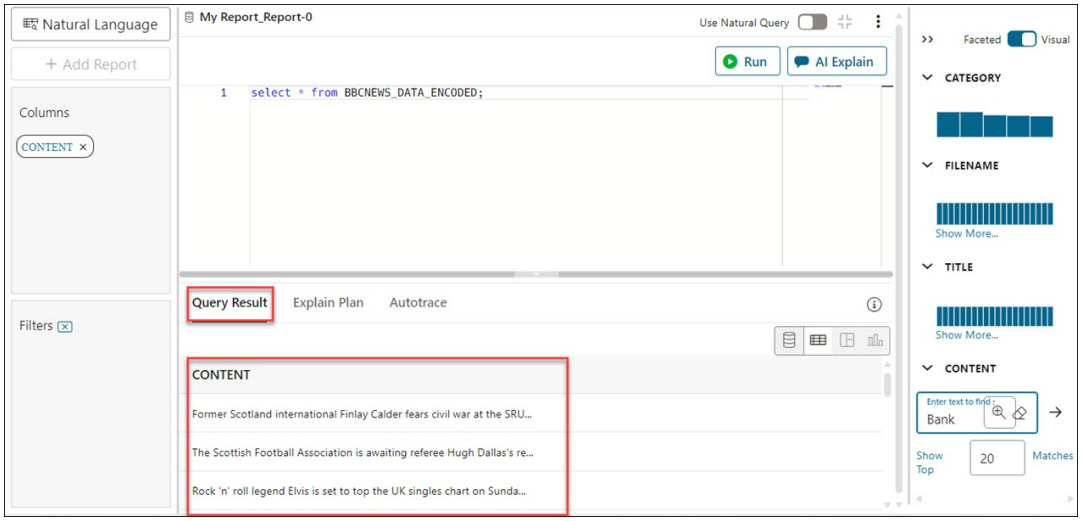
10. Enter the following text in the text box to content similar to the word `Bank` in the `CONTENT` column:

    ```
    BANK
    ```

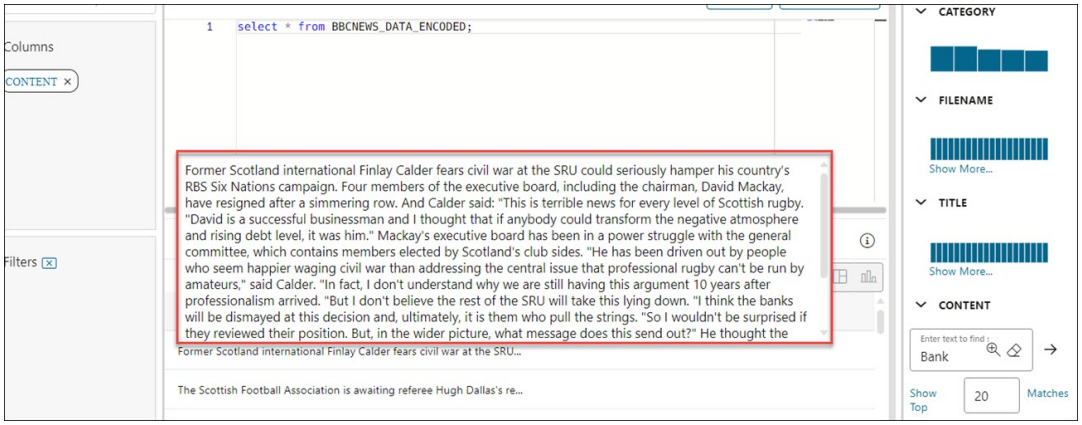
    Click **Search**.



11. The Query Results tab of the Output panel displays the content column that consists of content related to `Bank`.



You can hover over each content to view it's expanded form.



Using Calculation Templates

The Data Analysis tool provides templates for all of the calculations typically in demand for business intelligence applications.

The following topics describe the types of calculations available as calculation templates in the tool.

- Cumulative Aggregates
- Prior and Future Period
- Period to Date
- Parallel Period
- Moving Aggregates
- Share
- Rank

- Cumulative Aggregates
  Cumulative calculations start with the first time period and calculate up to the current member, or start with the last time period and calculate back to the current member.
- Prior and Future Period
  The Data Analysis tool provides several calculations for prior or future time periods.
- Period to Date
  Period-to-date functions perform a calculation over time periods with the same parent up to the current period.
- Parallel Period
  Parallel periods are at the same level as the current time period, but have different parents in an earlier period. For example, you may want to compare current sales with sales for the prior year at the quarter and month levels.
- Moving Aggregates
  Moving aggregates are performed over the time periods surrounding the current period.
- Share
  Share calculates the ratio of a measure's value for the current dimension member to the value for a related member of the same dimension.
- Rank
  Rank orders the values of a dimension based on the values of the selected measure. When defining a rank calculation, you choose the dimension, a hierarchy, and the measure.

Cumulative Aggregates

Cumulative calculations start with the first time period and calculate up to the current member, or start with the last time period and calculate back to the current member.

The tool provides several aggregation methods for cumulative calculations:

- **Cumulative Average**: Calculates a running average across time periods.
- **Cumulative Maximum**: Calculates the maximum value across time periods.
- **Cumulative Minimum**: Calculates the minimum value across time periods.
- **Cumulative Total**: Calculates a running total across time periods.

You can choose the measure, the time dimension, and the hierarchy. For selecting the time range see "Choosing a Range of Time Periods" in *Oracle OLAP User's Guide*.

Cumulative Calculation Example

This template defines a calculated measure using Cumulative Minimum.

```
Cumulative minimum of SALES in the TIME dimension and TIME.CALENDAR hierarchy within
ancestor at level TIME.CALENDAR_YEAR. Total from beginning to current member.
```

These are the results of a query against the calculated measure, which displays values for the descendants of calendar year 2021. The minimum value for quarters begins with Q1-21 and ends with Q4-21, and for months begins with Jan-21 and ends with Dec-21.

```
TIME     TIME_LEVEL                SALES  MIN_SALES
-------- -------------------- ---------- ----------
Q1.21    CALENDAR_QUARTER       32977874   32977874
Q2.21    CALENDAR_QUARTER       35797921   32977874
Q3.21    CALENDAR_QUARTER       33526203   32977874
Q4.21    CALENDAR_QUARTER       41988687   32977874
JAN-21   MONTH                  11477898   11477898
```

```
FEB-21   MONTH                 10982016   10982016
MAR-21   MONTH                 10517960   10517960
APR-21   MONTH                 11032057   10517960
MAY-21   MONTH                 11432616   10517960
JUN-21   MONTH                 13333248   10517960
JUL-21   MONTH                 12070352   10517960
AUG-21   MONTH                 11108893   10517960
SEP-21   MONTH                 10346958   10346958
OCT-21   MONTH                 14358605   10346958
NOV-21   MONTH                 12757560   10346958
DEC-21   MONTH                 14872522   10346958
```

Prior and Future Period

The Data Analysis tool provides several calculations for prior or future time periods.

Here are the calculations used for for prior or future time periods:

- **Prior Period**: Returns the value of a measure at an earlier time period.
- **Difference From Prior Period**: Calculates the difference between values for the current time period and an earlier period.
- **Percent Difference From Prior Period**: Calculates the percent difference between the values for the current time period and an earlier period.
- **Future Period**: Returns the value of a measure at a later time period.
- **Difference From Future Period**: Calculates the difference between the values for the current time period and a later period.
- **Percent Difference From Future Period**: Calculates the percent difference between the values for the current time period and a later period.

When creating a calculation for prior or future time periods, you choose the measure, the time dimension, the hierarchy, and the number of periods from the current period.

Prior Period Example

This template defines a calculated measure using Prior Period:

```
Prior period for measure SALES in TIME dimension and TIME.CALENDAR hierarchy 1 period
ago.
```

These are the results of a query against the calculated measure. The `PRIOR_PERIOD` column shows the value of Sales for the preceding period at the same level in the Calendar hierarchy.

```
TIME     TIME_LEVEL           SALES PRIOR_PERIOD
-------- -------------------- ---------- ------------
2020     CALENDAR_YEAR         136986572    144290686
2021     CALENDAR_YEAR         140138317    136986572
Q1.20    CALENDAR_QUARTER       31381338     41988687
Q2.20    CALENDAR_QUARTER       37642741     31381338
Q3.20    CALENDAR_QUARTER       32617249     37642741
Q4.20    CALENDAR_QUARTER       35345244     32617249
Q1.21    CALENDAR_QUARTER       36154815     35345244
Q2.21    CALENDAR_QUARTER       36815657     36154815
Q3.21    CALENDAR_QUARTER       32318935     36815657
Q4.21    CALENDAR_QUARTER       34848911     32318935
```

Period to Date

Period-to-date functions perform a calculation over time periods with the same parent up to the current period.

These functions calculate period-to-date:

- **Period to Date**: Calculates the values up to the current time period.
- **Period to Date Period Ago**: Calculates the data values up to a prior time period.
- **Difference From Period to Date Period Ago**: Calculates the difference in data values up to the current time period compared to the same calculation up to a prior period.
- **Percent Difference From Period To Date Period Ago**: Calculates the percent difference in data values up to the current time period compared to the same calculation up to a prior period.

When creating a period-to-date calculation, you can choose from these aggregation methods:

- Sum
- Average
- Maximum
- Minimum

You also choose the measure, the time dimension, and the hierarchy.

Period to Date Example

This template defines a calculated measure using Period to Date.

```
Gregorian Year to date for SALES in the TIME dimension and TIME.CALENDAR hierarchy.
Aggregate using MINIMUM from the beginning of the period.
```

These are the results of a query against the calculated measure. The `MIN_TO_DATE` column displays the current minimum `SALES` value within the current level and year.

```
TIME     TIME_LEVEL               SALES MIN_TO_DATE
-------- -------------------- ---------- -----------
Q1.21    CALENDAR_QUARTER       36154815    36154815
Q2.21    CALENDAR_QUARTER       36815657    36154815
Q3.21    CALENDAR_QUARTER       32318935    32318935
Q4.21    CALENDAR_QUARTER       34848911    32318935
JAN-21   MONTH                   13119235    13119235
FEB-21   MONTH                   11441738    11441738
MAR-21   MONTH                   11593842    11441738
APR-21   MONTH                   11356940    11356940
MAY-21   MONTH                   13820218    11356940
JUN-21   MONTH                   11638499    11356940
JUL-21   MONTH                    9417316     9417316
AUG-21   MONTH                   11596052     9417316
SEP-21   MONTH                   11305567     9417316
OCT-21   MONTH                   11780401     9417316
NOV-21   MONTH                   10653184     9417316
DEC-21   MONTH                   12415325     9417316
```

Parallel Period

Parallel periods are at the same level as the current time period, but have different parents in an earlier period. For example, you may want to compare current sales with sales for the prior year at the quarter and month levels.

The Data Analysis tool provides several functions for parallel periods:

- **Parallel Period**: Calculates the value of the parallel period.
- **Difference From Parallel Period**: Calculates the difference in values between the current period and the parallel period.
- **Percent Difference From Parallel Period**: Calculates the percent difference in values between the current period and the parallel period.

To identify the parallel period, you specify a level and the number of periods before the current period. You can also decide what happens when two periods do not exactly match, such as comparing daily sales for February (28 days) with January (31 days).

You also choose the measure, the time dimension, and the hierarchy.

Parallel Period Example

This template defines a calculated measure using Parallel Period.

```
Parallel period for SALES in the TIME dimension and TIME.CALENDAR hierarchy 1
TIME.CALENDAR.QUARTER ago based on position from beginning to ending of period.
```

These are the results of a query against the calculated measure, which lists the months for two calendar quarters. The parallel month has the same position within the previous quarter. The prior period for `JUL-21` is `APR-21`, for `AUG-21` is `MAY-21`, and for `SEP-21` is `JUN-21`.

```
TIME     PARENT      SALES    LAST_QTR
-------- ---------- ---------- ----------
APR-21   CY2006.Q2    11356940   13119235
MAY-21   CY2006.Q2    13820218   11441738
JUN-21   CY2006.Q2    11638499   11593842
JUL-21   CY2006.Q3     9417316   11356940
AUG-21   CY2006.Q3    11596052   13820218
SEP-21   CY2006.Q3    11305567   11638499
```

Moving Aggregates

Moving aggregates are performed over the time periods surrounding the current period.

The Data Analysis tool provides several aggregation methods for moving calculations:

- **Moving Average**: Calculates the average value for a measure over a fixed number of time periods.
- **Moving Maximum**: Calculates the maximum value for a measure over a fixed number of time periods.
- **Moving Minimum**: Calculates the minimum value for a measure over a fixed number of time periods.
- **Moving Total**: Returns the total value for a measure over a fixed number of time periods.

You can choose the measure, the time dimension, and the hierarchy. You can also select the range, as described in "Choosing a range of time periods" in *Oracle OLAP User's Guide* , and the number of time periods before and after the current period to include in the calculation.

Moving Aggregates Example

This template defines a calculated measure using Moving Minimum.

```
Moving minimum of SALES in the TIME dimension and TIME.CALENDAR hierarchy. Include 1
preceding and 1 following members within level.
```

These are the results of a query against the calculated measure, which displays values for the descendants of calendar year 2021. Each value of Minimum Sales is the smallest among the current value and the values immediately before and after it. The calculation is performed over all members of a level in the cube.

```
TIME     TIME_LEVEL           SALES  MIN_SALES
-------- -------------------- ---------- ----------
Q1.21    CALENDAR_QUARTER       32977874   32977874
Q2.21    CALENDAR_QUARTER       35797921   32977874
Q3.21    CALENDAR_QUARTER       33526203   33526203
Q4.21    CALENDAR_QUARTER       41988687   31381338
JAN-21   MONTH                  11477898   10982016
FEB-21   MONTH                  10982016   10517960
MAR-21   MONTH                  10517960   10517960
APR-21   MONTH                  11032057   10517960
MAY-21   MONTH                  11432616   11032057
JUN-21   MONTH                  13333248   11432616
JUL-21   MONTH                  12070352   11108893
AUG-21   MONTH                  11108893   10346958
SEP-21   MONTH                  10346958   10346958
OCT-21   MONTH                  14358605   10346958
NOV-21   MONTH                  12757560   12757560
DEC-21   MONTH                  14872522   12093518
```

Share

Share calculates the ratio of a measure's value for the current dimension member to the value for a related member of the same dimension.

You can choose whether the related member is:

- **Top of hierarchy**: Calculates the ratio of each member to the total.
- **Member's parent**: Calculates the ratio of each member to its parent.
- **Member's ancestor at level**: Calculates the ratio of each member to its ancestor, that is, a member at a specified level higher in the hierarchy.

When creating a share calculation, you can choose the measure, dimension, and hierarchy. You also have the option of multiplying the results by 100 to get percentages instead of fractions.

Share Example

This template defines a calculated measure using `SHARE`:

```
Share of measure SALES in PRODUCT.PRIMARY hierarchy of the PRODUCT dimension as a ratio
of top of hierarchy.
```

These are the results of a query against the calculated measure. The `TOTAL_SHARE` column displays the percent share of the total for the selected products.

```
PRODUCT            PROD_LEVEL          SALES TOTAL_SHARE
------------------- --------------- ---------- -----------
Total Product        TOTAL             144290686         100
```

```
Hardware             CLASS          130145388          90
Desktop PCs          FAMILY          78770152          55
Portable PCs         FAMILY          19066575          13
CD/DVD               FAMILY          16559860          11
Software/Other       CLASS           14145298          10
Accessories          FAMILY           6475353           4
Operating Systems    FAMILY           5738775           4
Memory               FAMILY           5430466           4
Modems/Fax           FAMILY           5844185           4
Monitors             FAMILY           4474150           3
Documentation        FAMILY           1931170           1
```

Rank

Rank orders the values of a dimension based on the values of the selected measure. When defining a rank calculation, you choose the dimension, a hierarchy, and the measure.

You can choose a method for handling identical values:

- **Rank**: Assigns the same rank to identical values, so there may be fewer ranks than there are members. For example, it may return `1`, `2`, `3`, `3`, `4` for a series of five dimension members.
- **Dense Rank**: Assigns the same minimum rank to identical values. For example, it may return `1`, `2`, `3`, `3`, `5` for a series of five dimension members.
- **Average Rank**: Assigns the same average rank to identical values. For example, it may return `1`, `2`, `3.5`, `3.5`, `5` for a series of five dimension members.

You can also choose the group in which the dimension members are ranked:

- **Member's level**: Ranks members at the same level.
- **Member's parent**: Ranks members with the same parent.
- **Member's ancestor at level**: Ranks members with the same ancestor at a specified level higher in the hierarchy.

Rank Example

This template defines a calculated measure using Rank:

```
Rank members of the PRODUCT dimension and PRODUCT.PRIMARY hierarchy based on measure
SALES. Calculate rank using RANK method with member's parent in order lowest to highest.
Rank NA (null) values nulls last.
```

These are the results of a query against the calculated measure in which the products are ordered by `RANK`:

```
PRODUCT              SALES       RANK
-------------------- ---------- ----------
Monitors                4474150          1
Memory                  5430466          2
Modems/Fax              5844185          3
CD/DVD                 16559860          4
Portable PCs           19066575          5
Desktop PCs            78770152          6
```

Oracle Autonomous Database for Excel

The Oracle Autonomous Database for Excel integrates Microsoft Excel spreadsheets with the Autonomous Database to retrieve and analyze data from Analytic Views in the database. You can also directly run SQL queries to view their results in the worksheet.

> **Note:**
>
> - The **Download Microsoft Excel/Google Sheets add-in** is available to you under the **Downloads** menu of your Database Actions instance only if you have the `ADPUSER` role.
> - You require ADMIN privileges to grant roles.
> - To grant yourself an `ADPUSER` role, you must log into your Database Actions instance and enter the following command in the SQL worksheet area displayed in the **SQL** tab under **Development** tools present in the Launchpad:
>
>   ```
>   Grant ADPUSER to USER;
>   ```

- Install the add-in from Microsoft AppSource
  The **Oracle Autonomous Database for Excel** is available in **Microsoft AppSource**.
- Install the add-in from Database Actions
  You can install the add-in from Database Actions. This is accessed by users with Oracle Account.
- Uninstall the add-in
  The following section describes the steps to uninstall the Oracle Autonomous Database add-in.
- Using Oracle Autonomous Database for Excel
  After you install the add-in, a new ribbon tab, **Oracle Autonomous Database** appears in MS Excel.
- Connection management
  Each time you start the add-in for Excel, you must create a connection.
- Log Settings
  The Settings icon enables you to set the log levels using the Excel add-in.
- Run Direct SQL queries in an Excel worksheet
  The Oracle Autonomous Database for Excel lets you run Direct SQL queries to work with your data in an Excel worksheet.
- Read and Access Data Using Table Hyperlinks in Excel Sheet
  Table Hyperlinks in Oracle Autonomous Database provide secure, preauthenticated URLs that allow read-only access to data stored in tables, views, or the result of SQL queries.
- Reporting and Analysis in Excel
  You can view Reports and Analytic Views or visualizing data for analysis purpose using the **Analysis and Reports** menu.
- Query an Analytic View in an Excel worksheet
  The Query Wizard menu enables you to query an Analytic View and retrieve the results in an Excel Worksheet. Once the wizard retrieves the data, it becomes local to Excel. You can further edit the data in Excel but not write back to the Autonomous Database.

- Data Analysis in Excel Sheet
  The Data Analysis tool enables you to analyze data in the Autonomous Database by running SQL queries or querying an Analytical View using an intuitive drag-and-drop interface. You can also write custom queries to be run. You now can dynamically apply filters to the result set retrieved using the new faceted search capability on specific columns.
- Natural Language in Excel Sheet
  You can use Natural Language Query to query the Oracle Autonomous Database using the **Natural Language** menu in the **Oracle Autonomous Database for Excel** .
- FAQs for Troubleshooting errors with Excel Add-in
  If you experience issues with Oracle Autonomous Database for Excel, refer to frequently asked questions in this section to identify and resolve issues.
- Oracle Autonomous Database for Excel Privacy Policy Details
  This topic covers details for writing policies to control access to Autonomous Database resources.

Using Oracle Autonomous Database for Excel

After you install the add-in, a new ribbon tab, **Oracle Autonomous Database** appears in MS Excel.

You can connect to multiple Autonomous Databases, work with Analytic Views, tables, and Views, and view the data in the worksheet.

This ribbon provides buttons that let you connect to the Autonomous Database.

Click **Connections** to connect to an Autonomous Database. You must Refer to the *Connection management* for more details.

Click **Settings** to view the logging level settings of the Excel Add-in. You can also clear the logs or export the log files by copying the logging information to the clipboard.

Click **About** to view the Add-in and the supported Excel versions. The About window also displays whether the spreadsheet is connected to the database. It also shows version information for the database and Oracle Rest Data Services.

Click **Direct SQL** to run custom SQL queries. See

Click **Data Analysis** to query an existing Analytic View and run queries.

Click **Natural Language** to run Natural Language Queries on Tables.

Read and Access Data Using Table Hyperlinks in Excel Sheet

Table Hyperlinks in Oracle Autonomous Database provide secure, preauthenticated URLs that allow read-only access to data stored in tables, views, or the result of SQL queries.

They enable easy sharing and retrieval of database data without requiring additional authentication beyond the hyperlink itself. You can read and access data in a Microsoft Excel sheet using a Table Hyperlink through the Oracle Autonomous Database Add-in for Excel, which allows direct retrieval and viewing of tabular data from the Autonomous Database within Excel.

You can use **Table Hyperlinks** with the add-in.

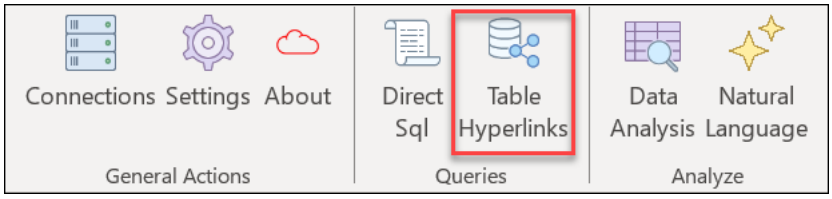
For more information of Table Hyperlinks, refer to the Use Table Hyperlinks for Read Only Data Access on Autonomous Database Chapter.

> **Note:**
>
> This feature allows read-only access to data through secure, table hyperlink URLs without requiring an additional database connections or an Oracle Autonomous Database account.

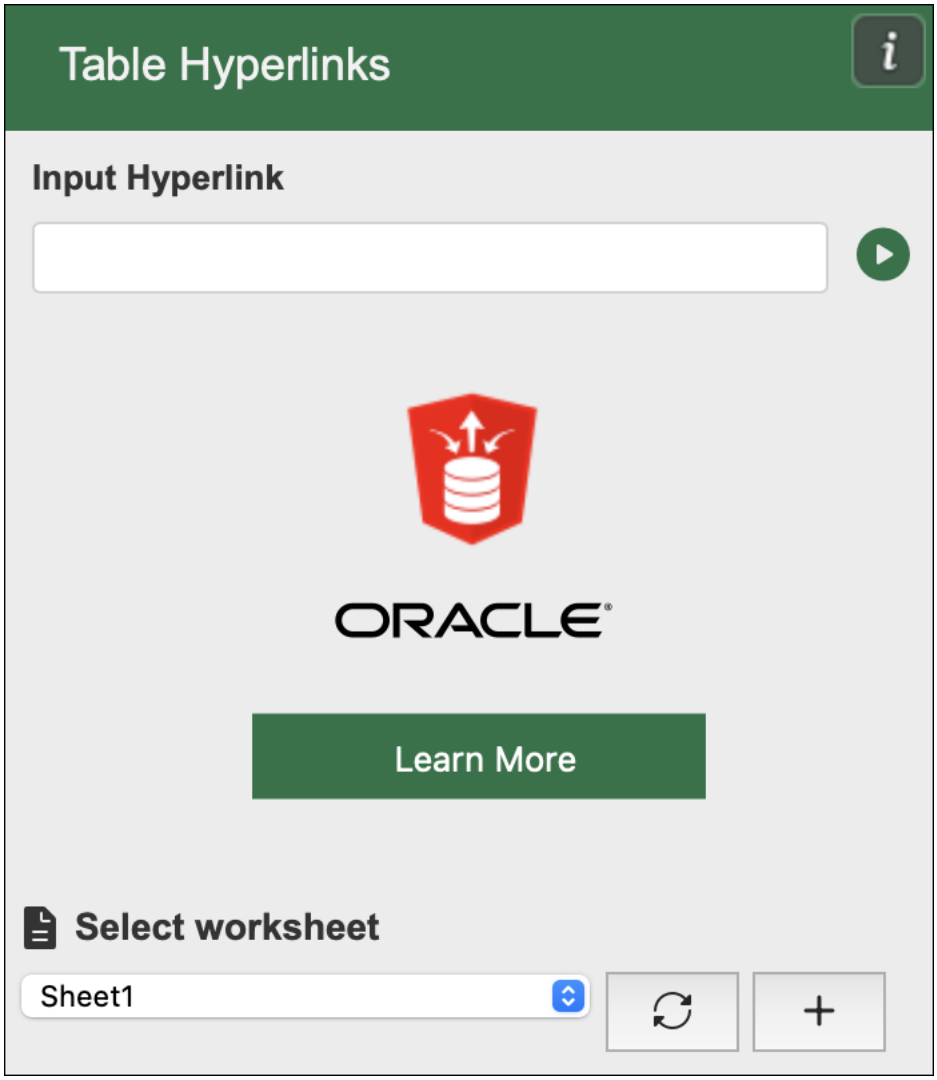
Use **Table Hyperlinks** to read and access data and view the data in a tabular format.

To return data from the add-in using Table Hyperlinks:

1. On the Excel sheet, select the menu **Autonomous Database**.
2. Select **Table Hyperlinks**.



3. Selecting Table Hyperlinks opens a **Table Hyperlinks** panel.



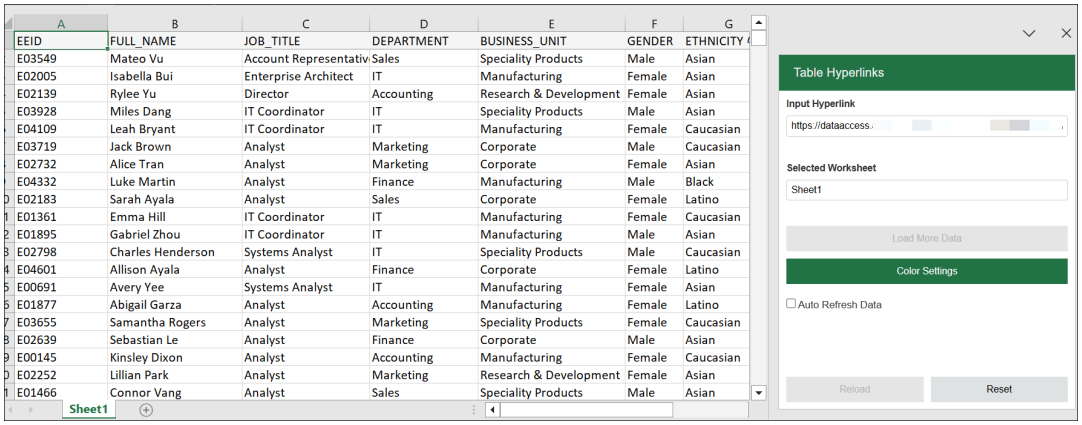
4. On the **Input Hyperlink** field, enter the input Table Hyperlink the producer shares with you. A Producer generates a PAR URL in their database and provides it to you. For example, *https://dataaccess.adb.eu-frankfurt-1.oraclecloudapps.com/adb/p/F5Sn....*
You can create Table Hyperlinks in Oracle Autonomous Database Serverless using the `DBMS_DATA_ACCESS` PL/SQL package and share them with the consumers.

   See DBMS_DATA_ACCESS Package for more information on this topic.

   Select Worksheet: Enter the name of the sheet where you want to display the results. For example, *Sheet1*.

   You can click the + sign beside the Select worksheet drop-down to display the results in a new worksheet.

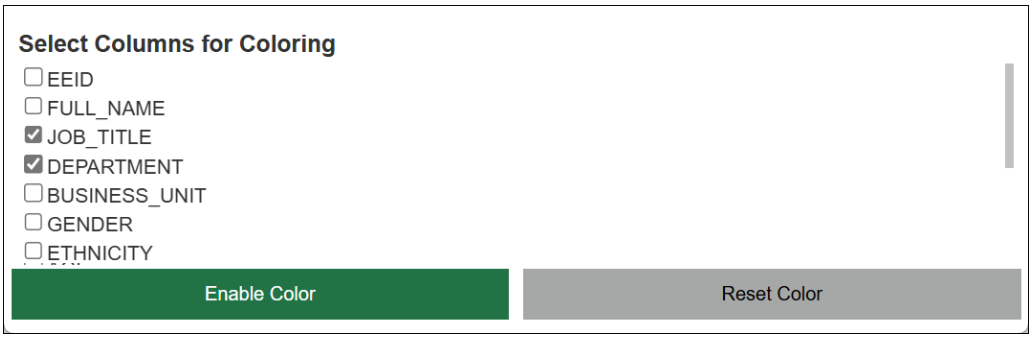
5. Click the **Play** button to load the Autonomous Data table in the worksheet.
The result data gets retrieved in the worksheet. You can view only first 10000 records in the first load.



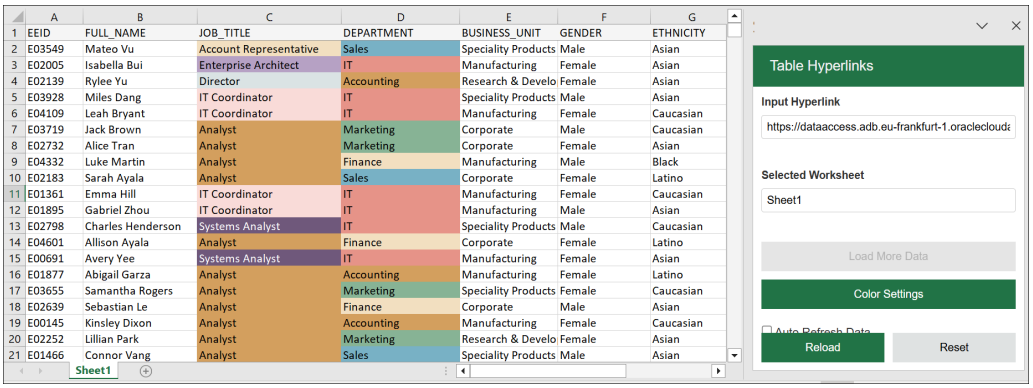
After selecting the play icon, along with the results retrieved in the worksheet, you can also see the following generated fields:

- **Load More Data**:
  Select **Load More Data** to retrieve the next 10000 rows from the table.
- **Color Settings**:
  Select this icon to enable color in the column from a predefined pool of colors. If the feature runs out of colors, the pool is refilled with the same colors for the new value keys.

  Click the icons you want to add color to. In this example, select *JOB_TITLE* and *BUSINESS_UNIT*.



- Select **Enable Color** to color the selected columns in the data.



You can select **Reset Color** to reset your current selection, remove colors from all the selected columns and choose different columns you wish to enable color to.

- **Auto Refresh Data**:
  Enable this option to automatically update data at regular intervals with the latest data from its source database. This automatic refresh keeps the tabular data in worksheet in sync with the current data within a specific refresh frequency.

  You can select the frequency from the frequency slider where the selection varies from a range of 5 minutes to 1 day.
- **Reload**: Select **Reload** to reload the preloaded table hyperlink data.

> **Note:**
>
> This button loads a maximum of 50,000 records.

- **Reset**: Select **Reset** to clear all data from the sheet, reset your selection of all the current options and enter the **Input hyperlink** again.

Install the add-in from Microsoft AppSource

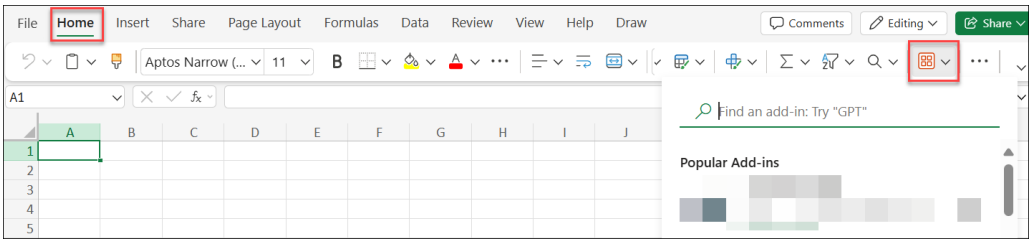
The **Oracle Autonomous Database for Excel** is available in **Microsoft AppSource**.

Prerequisites

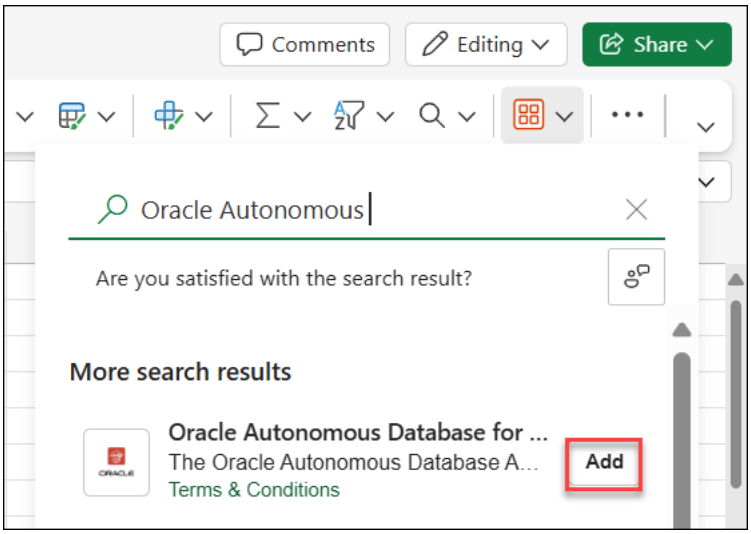
- You must have an active Microsoft account or an account with access to Microsoft 365.
- You can access the web version of Excel or the (Licensed version) from desktop.
- You must verify within your organization whether external apps from AppSource are permitted to be installed in your Microsoft 365 accounts. If not, you can Install the add-in from Database Actions.

Follow these steps to install the add-in:

1. **Start Excel and open a new or existing workbook**:
   - You can access the web version of Microsoft Excel with your Microsoft AppSource account:
     - Go to your Microsoft 365 Copilot (earlier known as MS Office) app link:https://m365.cloud.microsoft
     - Login with your account.
     - Select Excel Workbook from the Microsoft 365 Copilot home page.
2. **Install the Add-in from Office Add-ins**:
   - From the Home menu, go to **Add-ins** and search for "*Oracle Autonomous*" in the search bar.



- You can view the Oracle Autonomous Database for Excel in the search results.



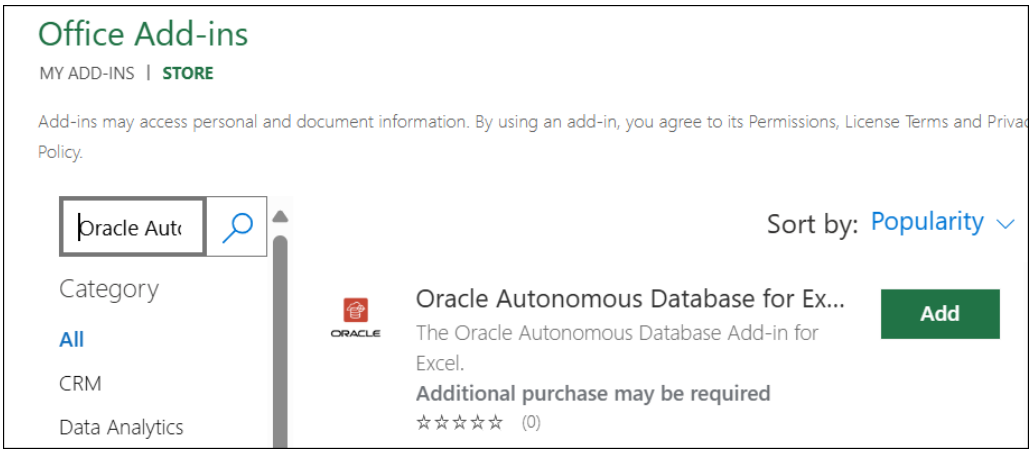
- Click **Add** to install the add-in.

OR

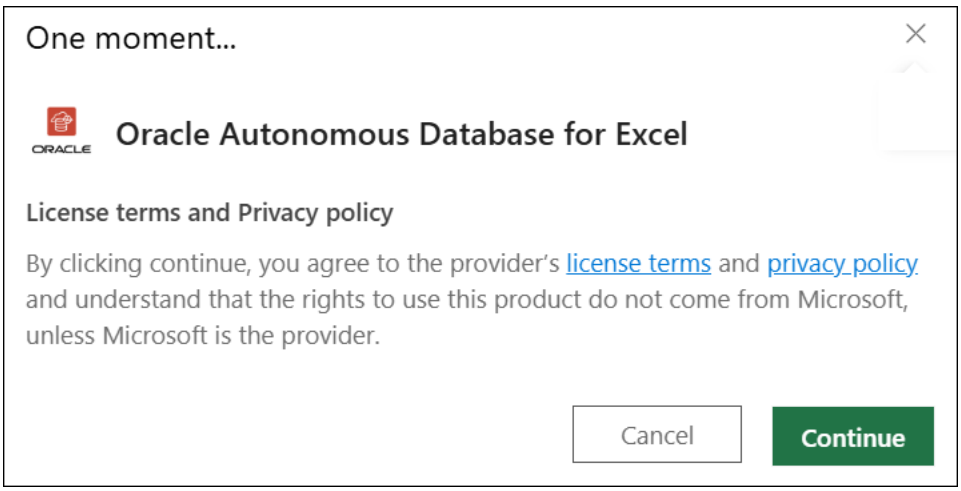
- You can alternatively click the **Add-ins** icon from the **Home**menu and select **+ More Add-ins** from the bottom of the Add-ins drop-down list.



- Enter "*Oracle Autonomous*" in the search field of the **STORE** tab of **Office Add-ins** and click the **search** icon.

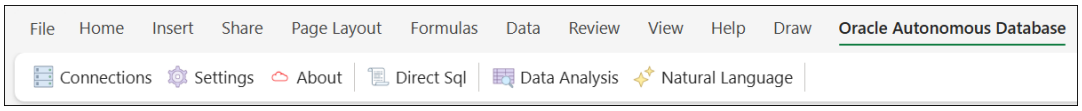


- You can view the Oracle Autonomous Database for Excel in the search results.
- Click **Add** to install the add-in in Excel.
- You will view a dialog to accept Oracle's license terms and Oracle Autonomous Database for Excel Privacy Policy Details. Click **Continue** to agree to it.



You will view a message confirming the add-in loaded successfully.

A new **Oracle Autonomous Database** ribbon tab appears in MS Excel.



After you have installed the add-in, you can now connect to the Autonomous Database to analyze and query the data. Refer to the Connection management chapter for more details on this topic.

Refer to the Uninstall the add-in chapter to uninstall the add-in you installed from Microsoft AppSource.

Install the add-in from Database Actions

You can install the add-in from Database Actions. This is accessed by users with Oracle Account.

Prerequisites:

- You must have an Oracle Cloud Account.
- You can install the add-in only from the licensed desktop version of Microsoft Excel.

The following instructions outline how to install the add-in depending on your operating system, specifically for Mac and Windows.

- Install the add-in on Mac
- Install the add-in on Windows

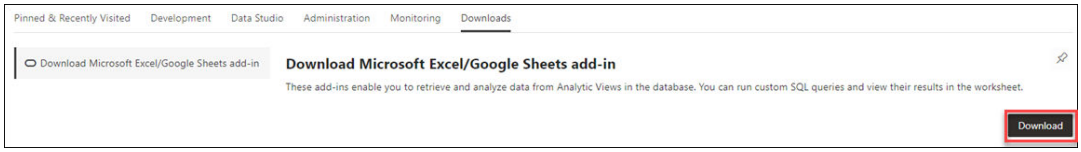
- Install the add-in on Mac
  The Oracle Autonomous Database for Excel is supported on Mac OS running Microsoft Office 365.
- Install the add-in on Windows
  The Oracle Autonomous Database for Excel is supported on Windows 10 and Windows 11 operating systems running Microsoft Excel 365.

Install the add-in on Mac

The Oracle Autonomous Database for Excel is supported on Mac OS running Microsoft Office 365.

To install the Autonomous Database for Excel, run the installer script file from your Autonomous Database instance by following the steps below:

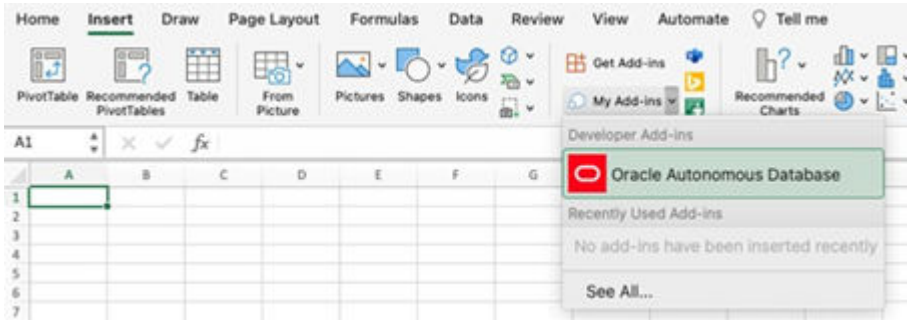
- Open the Database Actions Launchpad.
- On the **Downloads** tab of the Database Actions page, click the **Download Microsoft Excel/Google Sheets Add-in** pane.



- Click **Download**.
- Click the Microsoft Excel tab and select the **Download Add-in** button to download the **oracleplugin.zip** file.
- You can now view the zip file in the Downloads folder.
- Create a new folder named **Add-in** on your system.
- Extract the contents of the zip file in the **Add-in** folder.

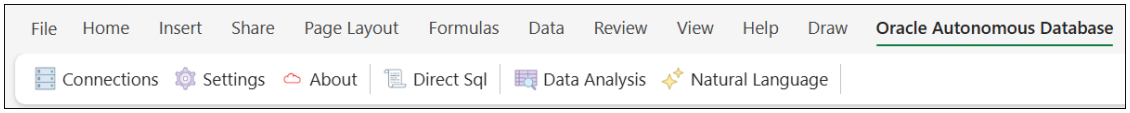
Follow these steps to install the add-in.

1. Quit Excel before you run the installer.
2. Navigate to the *install.sh* file in the **Add-in** folder.
3. Right-click *install.sh* and select the following options as shown : **Open With** -> **Other…** -> **Enable: All Applications** ->**Utilities**->**Terminal**
4. On completion, close the Terminal window.
5. Start Excel and open a new or existing workbook.
6. From the **Insert** menu in the Excel ribbon, select the drop-down menu of **My Add-ins**. A new **Oracle Autonomous Database** entry appears under the **Developer Add-Ins** dialog box.



7. Select **Oracle Autonomous Database**.

A new **Oracle Autonomous Database** ribbon tab appears in MS Excel.



Install the add-in on Windows

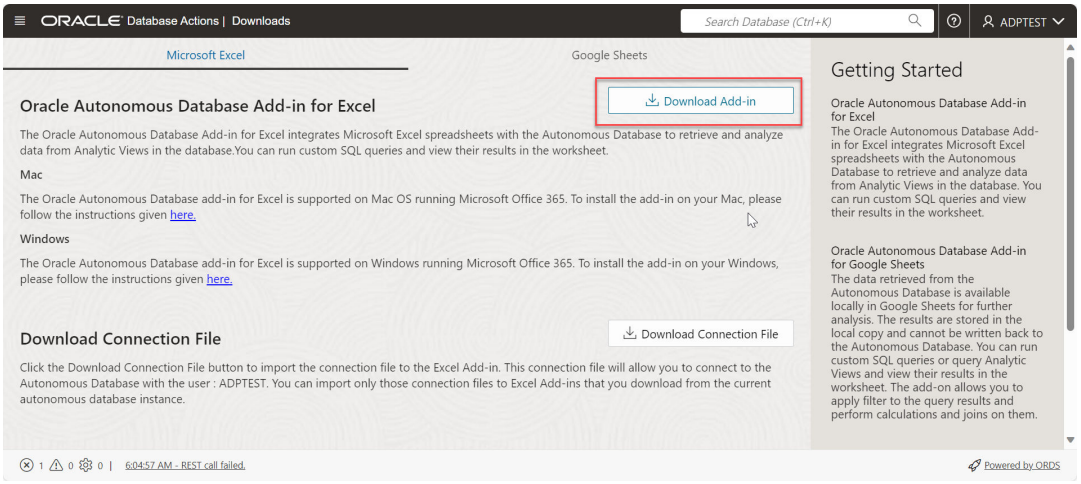
The Oracle Autonomous Database for Excel is supported on Windows 10 and Windows 11 operating systems running Microsoft Excel 365.

To install the Oracle Autonomous Database for Excel, download the *oracleplugin.zip* file and extract it to get the *install.cmd* script file from your Database Actions instance.

- Open the Database Actions Launchpad.
- On the **Downloads** tab of the Database Actions page, click the **Download Microsoft Excel/Google Sheets Add-in** pane.



- Click **Download**.
- Click the **Download Add-in** icon in the Microsoft Excel tab to download the Oracle Autonomous Database Add-in for Excel.



- Extract the *oracleplugin.zip* file to a new folder in the Downloads of your system. The extracted folder consists of an installer (install.cmd file), a manifest.xml file and a readme.txt file.

Follow these steps to install the add-in.

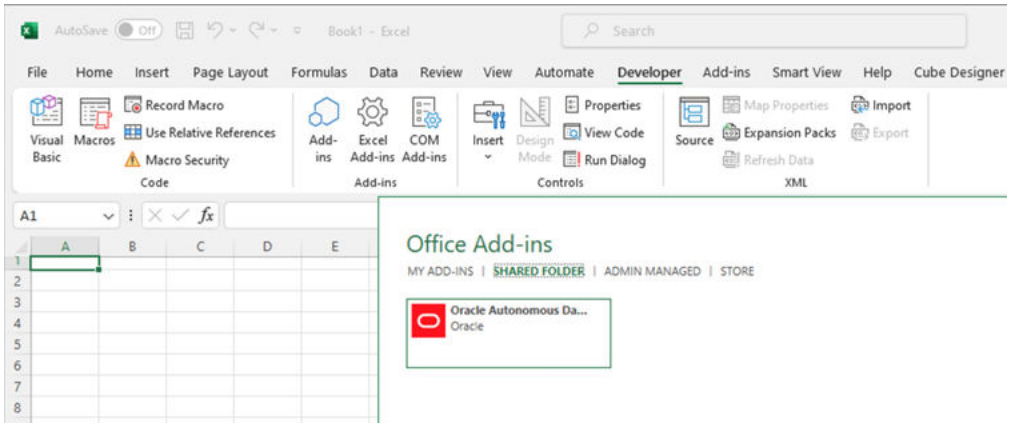
1. Quit Excel before you run the installer.
2. Right-click the *install.cmd* file that you downloaded.

> **Note:**
>
> After running the installer on Windows, the add-in automatically creates a network share and adds the shared location as a trusted catalog location for Office add-ins. A catalog is used to store the manifest for the Excel Add-in. It enables publishing and management of the Excel add-in as well as other add-ins that are available in the Office Store and licensed for corporate use. You can acquire the Excel add-in by specifying the shared manifest folder as a trusted catalog.

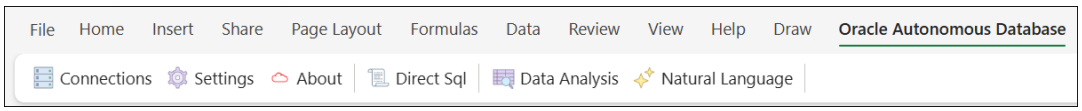
3. Select **Run as administrator**.

> **Note:**
>
> You must have Administrator privileges to install the Excel add-in for Oracle Autonomous Database successfully.

4. Start Excel and open a new or existing workbook.
5. From the **Developer** menu in the Excel ribbon, click **Add-ins**, select the **SHARED FOLDER** tab on the pop-up window and select **Oracle Autonomous Database**.



6. After you install the add-in, a new **Oracle Autonomous Database** ribbon tab appears in MS Excel.



> **Note:**
>
> You can re-run the installer after the initial installation. Re-run the installer and choose the option of your preference. You can repair your existing installation by deleting it, selecting the installed trusted catalog or adding another manifest to the working installation.

Uninstall the add-in

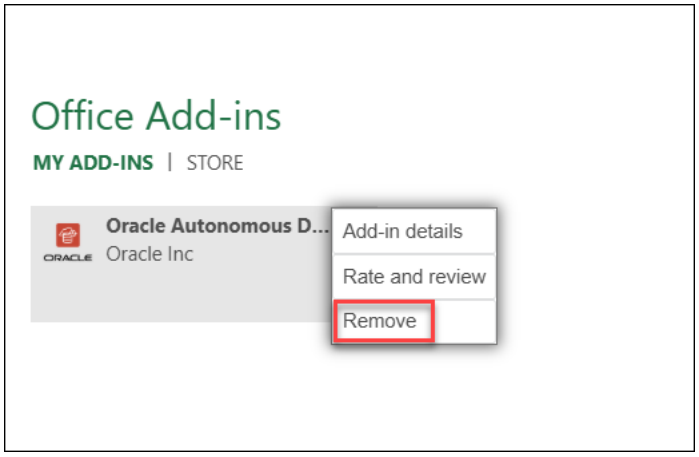
The following section describes the steps to uninstall the Oracle Autonomous Database add-in.

If you have installed **Oracle Autonomous Database for Excel** for Windows using the Microsoft AppSource, you can uninstall the add-in by following these steps:

- Click the **Add-ins** icon from the **Home**menu and select **+ More Add-ins** from the bottom of the Add-ins drop-down list.

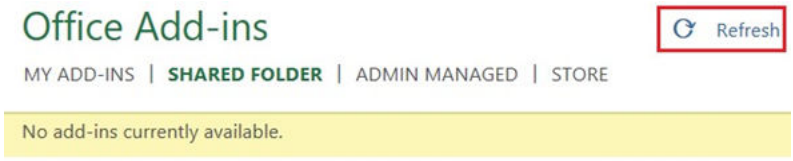


- Click **MY ADD-INS** from **Office Add-ins** and right-click the **Oracle Autonomous Database for Excel** add-in.
- Click**Remove** to uninstall the add-in from Excel.



If you have installed **Oracle Autonomous Database for Excel** for Windows using Database Actions, you can uninstall the add-in by following these steps:

- Delete the *manifest.xml* file from the folder located in *%LOCALAPPDATA% \Oracle\Autonomous Database\manifest* .
- Click **Refresh** in the Office Add-ins window to remove the Autonomous Database tab from MS Excel.



No add-ins will now be available from the Shared folder of the Office Add-ins window.

> **Note:**
>
> After uninstalling the Add-in, if you re-install it from a different Autonomous Database (ADB), the add-in attempts to load the old ADB. You must check if the shared manifest folder's location (share path) points to the correct location. For more details, refer to *Configuring the Excel Trusted Add-in Catalog* in FAQs for Troubleshooting errors with Excel Add-in.

Uninstall the add-in for Mac

To uninstall the Oracle Autonomous Database for Excel for Mac:

- Enter the following command in the terminal to remove the *manifest.xml* file.

```
rm ~/Library/Containers/com.microsoft.Excel/Data/Documents/wef/manifest.xml
```

The Oracle Autonomous Database for Excel is uninstalled from Mac.

Connection management

Each time you start the add-in for Excel, you must create a connection.

The connections feature lets you manage and connect to multiple Autonomous Databases with a single add-in. Multiple connections can be created. However, only one connection can remain active.

The connection panel lets you connect to the Autonomous Database through a connection where you provide the login credentials and access the Autonomous Database.

Within the Connections panel, you can:

- Create or delete multiple connections using a single add-in.
- Share connection information by exporting and importing connection information to a file.
- View existing connections.

Selecting **Connections** opens the Connections wizard.

> **Note:**
>
> This is an implicit type of connection. Refer to Authenticate using Implicit connection to understand more about implicit connection.

The Connections wizard consists of **Manage Connections** icon that has the following options available:

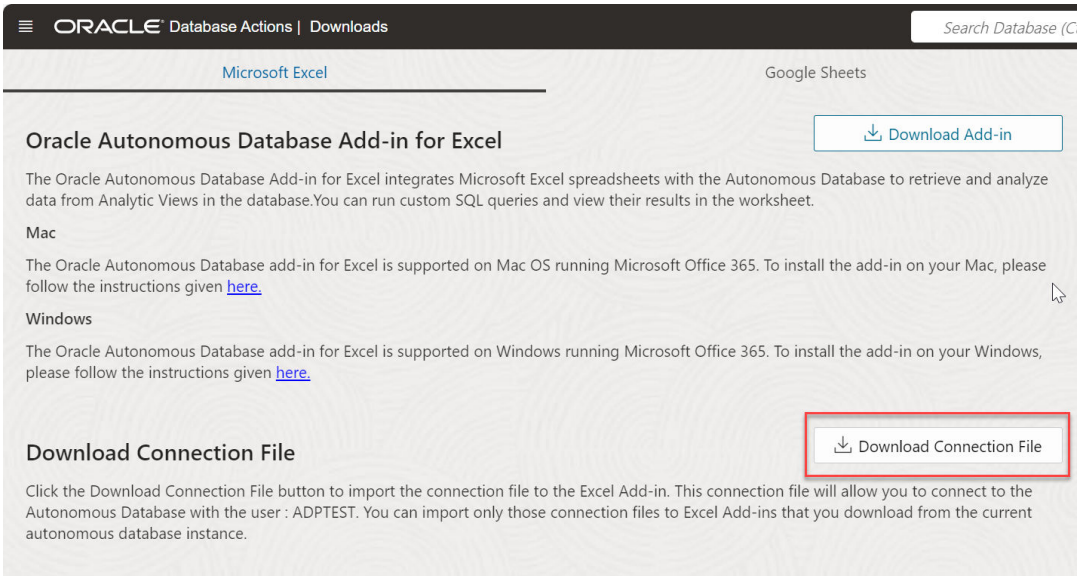
- **Add Connection**: Select **Add Connection** to Add a Connection. Refer to Add a connection section for more details.
- **Export Connection**: Select **Export Connection** to export connections. Refer to the Share a connection section for more details.
- **Import Connection**: Select **Import Connection** to launch the import wizard to choose a connection file. These files are in JSON format. Refer to the Share a connection section for more details.
- Import a Connection
  You can import a connection file that you can download from the Database Actions launchpad. This file is in JSON format.
- Add a connection
  You can manually create a connection to an autonomous database.
- Share a connection
  You can import or export a connection using the **Import Connection** and **Export Connection** menus from the **Manage Connections** drop-down field on the Connections panel.

Import a Connection

You can import a connection file that you can download from the Database Actions launchpad. This file is in JSON format.

Follow the below steps from the Database Actions instance, to download the connection file.

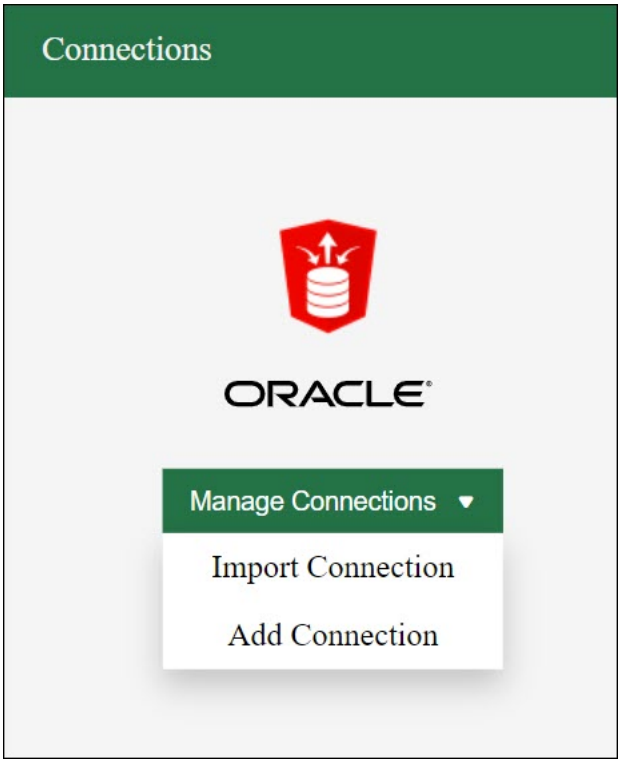
1. From the Downloads page of the Database Actions instance, click the **Download Connection File** icon to download the connection file in your system.



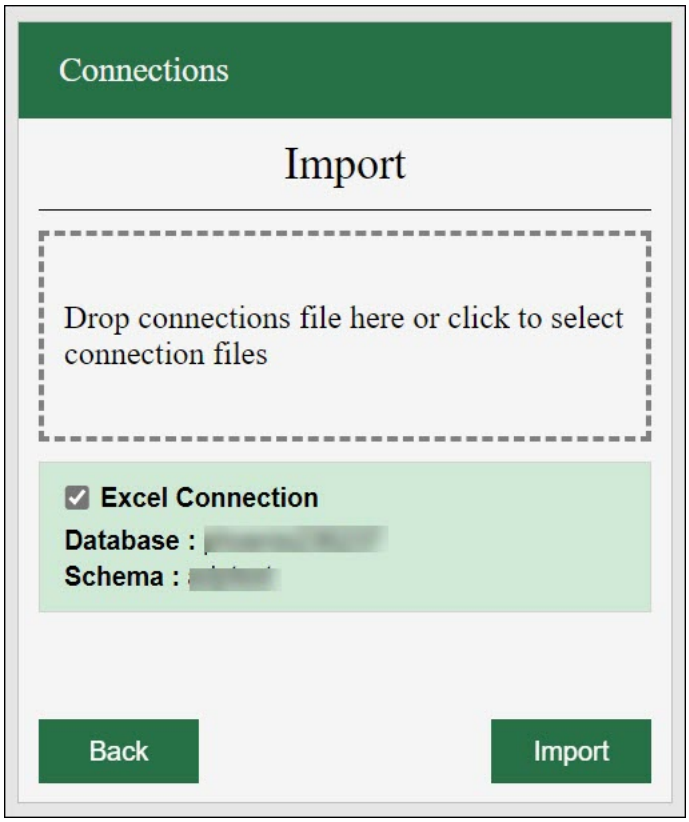
2. The file is downloaded onto your system.

To import a connection:

1. Select **Connections** from the Autonomous Database menu in the Excel sheet. This opens the Connections wizard that consists of a **Manage Connections** icon.
2. Click **Import Connection** to import the connection file you downloaded from the Database Actions instance.



3. Click and drop the connection file from your system to the drop area of the wizard. After the connection file loads, select the check box beside the connection file you want to import from the file.



4. Click **Import** to proceed.
5. Click the three vertical dots beside the connection file and select **Connect**.

> **Note:**
>
> If you view a red icon beside the connection file, the connection is inactive or incorrect. Click the three vertical dots beside the connection file and select **Edit** to update the connection file. Ensure the Autonomous Database URL is correct and click **Save**. An example of a correct URL is “https://<hostname>-<databasename>.adb.<region>.oraclecloudapps.com/"

6. Specify the schema name in the username field and the corresponding password in the credentials screen you view. Click **Sign in** to log in to the autonomous database.

Once signed in, you will view the connection highlighted in green. This means the connection is Active.

Add a connection

You can manually create a connection to an autonomous database.

> **WARNING:**
>
> This feature of adding connection manually is not recommended for this release.
>
> Please use the Import a Connection feature to import a connection file from the Downloads menu of the Database Actions launchpad.

Share a connection

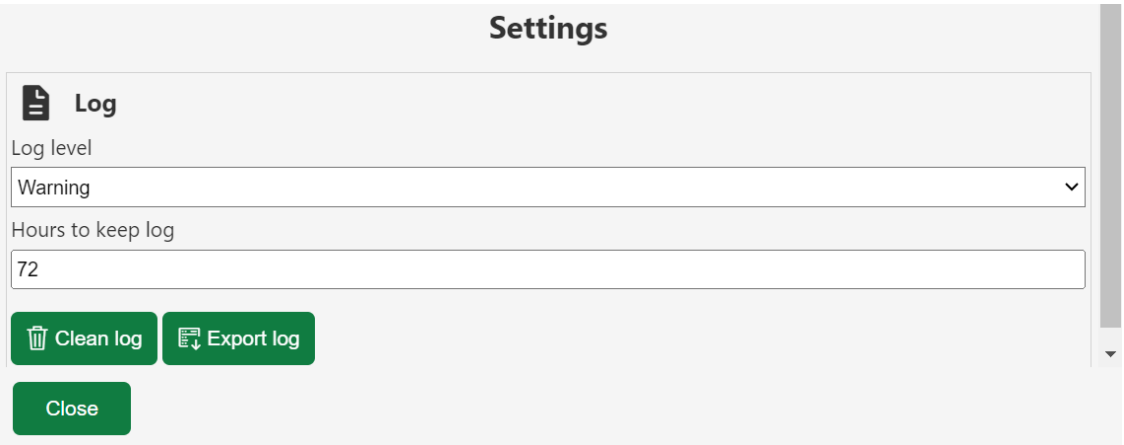
You can import or export a connection using the **Import Connection** and **Export Connection** menus from the **Manage Connections** drop-down field on the Connections panel.

- **Import Connection**: Click **Import Connection** from the **Manage Connections** drop-down icon and select a connection file from your local device. Once you import the connection file, you can view the connection in the Connection panel with a check box beside it. Select the check box and click **Import** to add the connection to the list of connections in the panel. The add-in copies the connection information you can use as a new connection. With the import feature, you do not have to enter the connection information to add a new connection.
- **Export Connection**: The export connection button exports an existing connection which you can import later. Clicking **Export Connection** opens a check box beside each connection in the connections list. You can select the connection you wish to export. Multiple selections are allowed. After you select the connection, click **Export**. Once the connection file is exported, you can view that the add-in for Excel downloads the connection file (*.json file) to your local device. The exported connection file is named spreadsheet_addin_connections.json.

Log Settings

The Settings icon enables you to set the log levels using the Excel add-in.

You can export or save the log. The log can be sent to the support team for diagnosing errors. The Log helps you troubleshoot errors. This field allows you to set the time limit for logging information that is stored in the log file. This setting automatically clears log information older than the entry made in this field. The default is 72 hrs. If the records are older than 72 hours, they will be deleted automatically.



Use Log Levels

Each log message is associated with a severity level. The level gives a rough guide to the importance and urgency of a log message. The log levels are specified with the following values ranging from one having lowest impact (left-most) to the one having highest impact (right-most):

```
STOP, ERROR, WARNING, INFO, VERBOSE
```

The following table lists the severity levels of log messages:

Log Level	Description
STOP	Records no information.
ERROR	Records events that may require corrective action, for example, a component is unavailable.
WARNING	Records issues that may lead to an error or require corrective action in the future. This is the default log level.
INFO	Records completed actions or the current state of a component, for example, the component is initializing.
VERBOSE	Displays detailed processing information along with errors, warnings and info types of severity levels.

Each logging level is cumulative, that is, each level contains all the data generated by the higher levels. For example, Warning logs contain all the data generated by the Error logs, plus the events that are specific to the Error category.

Log message format

The Log Settings writes a message to the specified log file consisting of a module, Timestamp, Severity, Server Name, Connection, Thread ID or User ID or Transaction ID, Message ID, and the Message, along with a stack trace if any. Each attribute is contained between angle brackets.

The following is an example of a message in the server log file (split for readability; in practice, the message might be on one line):

```
{"module":"ServerHelper","logMessage":"connection not is
      avialble","timestamp":1688901058020,"level":3,"payload":
{"conAlias":"Excel
      Connection","status":400},"id":161,"timeString":"9 Jul 2023
      16:40:58"},{"module":"ServerHelper","logMessage":"connection not is
      avialble","timestamp":1688901059724,"level":3}
```

Hours to keep a log

This field has the automatic cleaning operation. The log settings automatically clears old logs. If the logs are older than 72 hours which is the default value, they will be deleted automatically.

Clean Log

This button deletes all the logs.

Export Log

The Export Log button exports or saves logs to a file. Selecting the Export Log icon in Windows enables the add-in in downloading the exported log to “Downloads” folder.

Selecting the Export Log icon in Mac enables the add-in in copying the log information to clipboard. You can paste it to a file and save the log. The log you export will be in JSON format. The JSON file might contain sensitive information like your password, date of birth or any other data you want to secure.

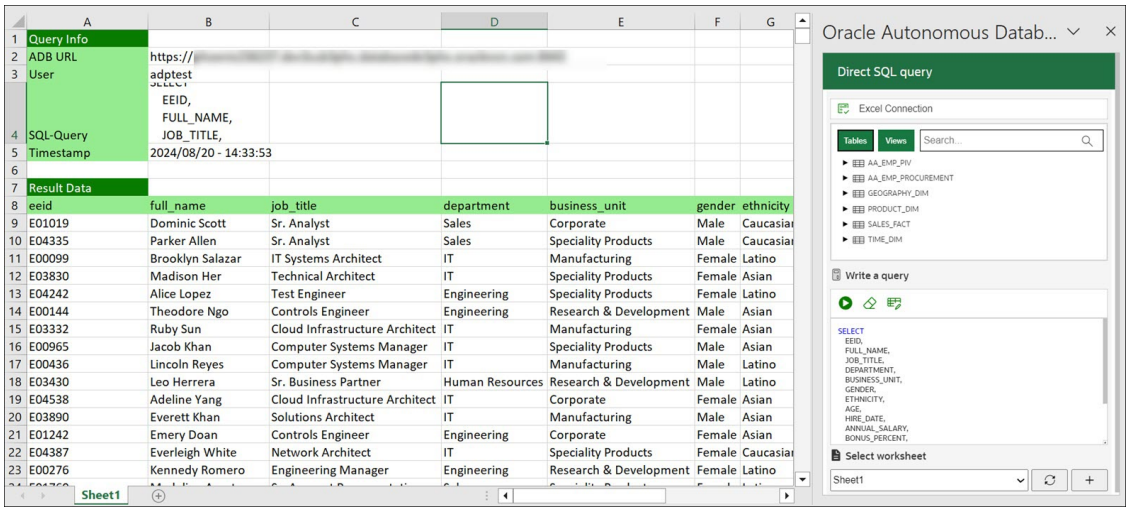
Run Direct SQL queries in an Excel worksheet

The Oracle Autonomous Database for Excel lets you run Direct SQL queries to work with your data in an Excel worksheet.

With the add-in, you can create a table and insert, update and delete rows from the existing tables or views. You can view the results in the current worksheet or different worksheets.

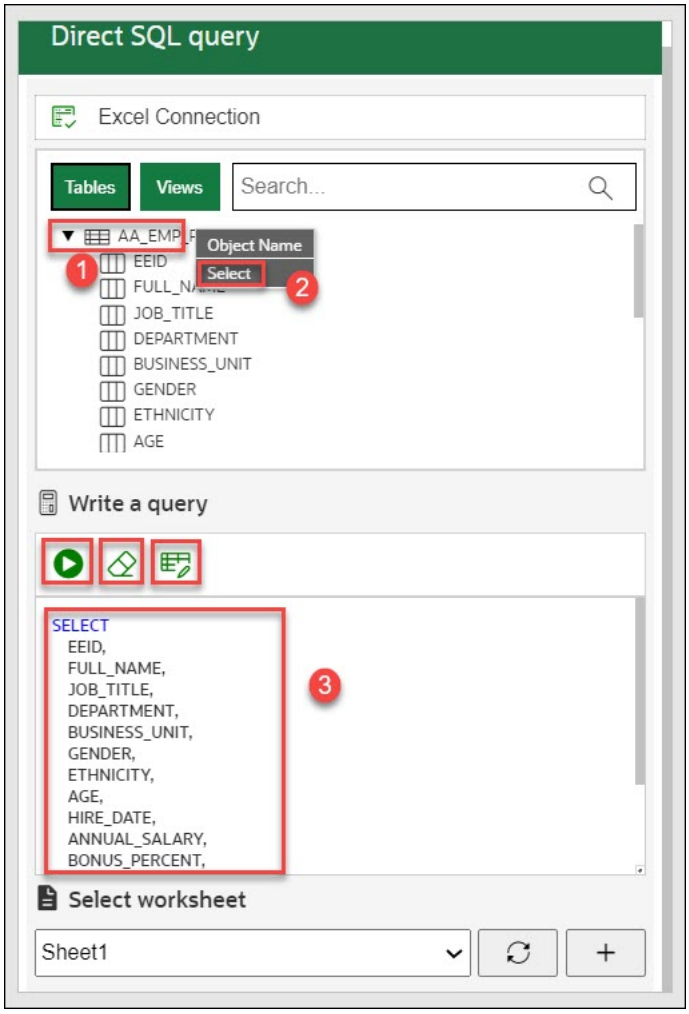
The following image shows your data retrieved from the Autonomous Database and displayed in the worksheet. The Query Info section comprising the Timestamp, User name and SQL Query are shown in Excel. You can edit custom queries and run them. The worksheet displays the results of queries from the retrieved data in tabular format.

The add-in maintains a live connection with the database. However, the data retrieved is local to Excel. In case of inactivity, the connection times out, and you must log in again. You can change the active connection from the connections panel. The image shows the results from a single query, but you can insert many queries in a single workbook.



To run a query using the add-in, run Excel and create a blank workbook using the standard Excel workbook file format.

1. In the Excel ribbon, select the **Autonomous Database**.
2. Click the **Direct SQL** icon from the ribbon. This opens an **Oracle Autonomous Database** dialog box in the Excel Task Pane with **Tables** and **Views** icons and a search field beside them.



3. Select **Table** to view all the existing tables in the schema. Click **Views** to see the current views in the schema.
4. You can right-click the table whose data you want to query and choose **Select** to view all the table's columns. The column names will be displayed in the Write a Query section. You can click on the table and view individual columns as well. Click the **Run** button to run the SQL query in the query editor. The query results will be displayed in the worksheet you select.

> **Note:**
>
> You will view an error message if you click the **Run** icon while the query editor is empty.

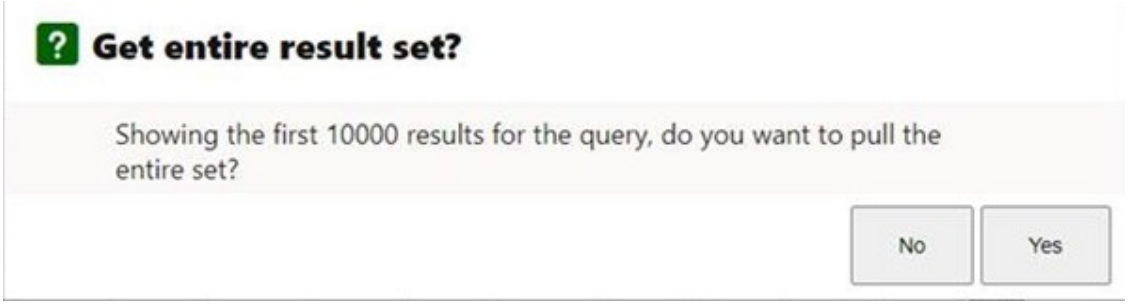
5. You can click **+** sign beside the **Select worksheet** drop-down to display the results in a new worksheet.

6. The worksheet also displays the timestamp, the user who creates and runs the query and the autonomous database URL.

To run another query follow these steps:

1. Click the eraser icon to clear the previous query from the SQL editor and write a new query.
2. Click **Retrieve query from Sheet** to import the SQL query from the existing worksheet and generates that query in the SQL editor.
3. In the **Select worksheet** drop-down, select a new sheet, *Sheet 2*, in this case. The Add-in adds a sheet for the user. If you work on the same sheet, the Add-in refreshes the data in the existing worksheet.
4. Click the **Run** icon to display the query results.

The worksheet displays the result of the query at a go. While this behavior works for most scenarios, sometimes, for large data sets, the query result might exceed 10K rows. Although you can view the 10K rows, a confirmation window asks if you want to view the rest of the results.



Select **Yes** to view the entire result set. Loading all the data may take a while. You must fetch all data before working with Pivot tables, or it will lead to incorrect results from aggregation in Pivot tables.

Close the Query Wizard panel to cancel the operation of fetching the result.

> **Note:**
>
> Close the Query Wizard panel to cancel the operation of fetching the result.

Reporting and Analysis in Excel

You can view Reports and Analytic Views or visualizing data for analysis purpose using the **Analysis and Reports** menu.

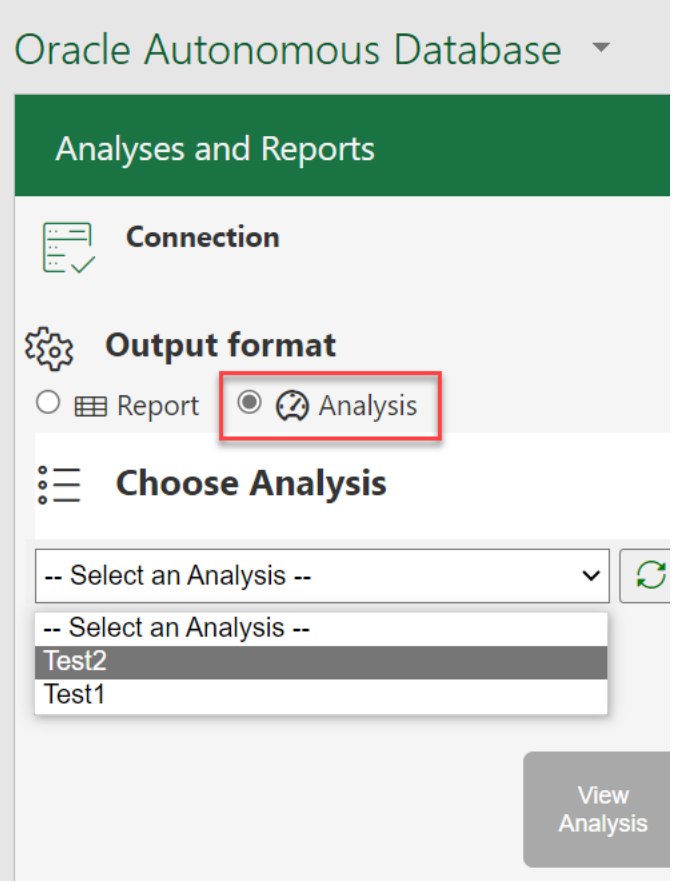
> **Note:**
>
> This menu is unavailable if you install **Oracle Autonomous Database for Excel** using Microsoft Store.

The reports and charts can be viewed in of different range of charts, namely, Bar Chart, Area Chart, Line Chart and Pie Chart. The Analyses and Reports icon enables you to retrieve Analyses and Reports from the Autonomous Database.

View Analysis

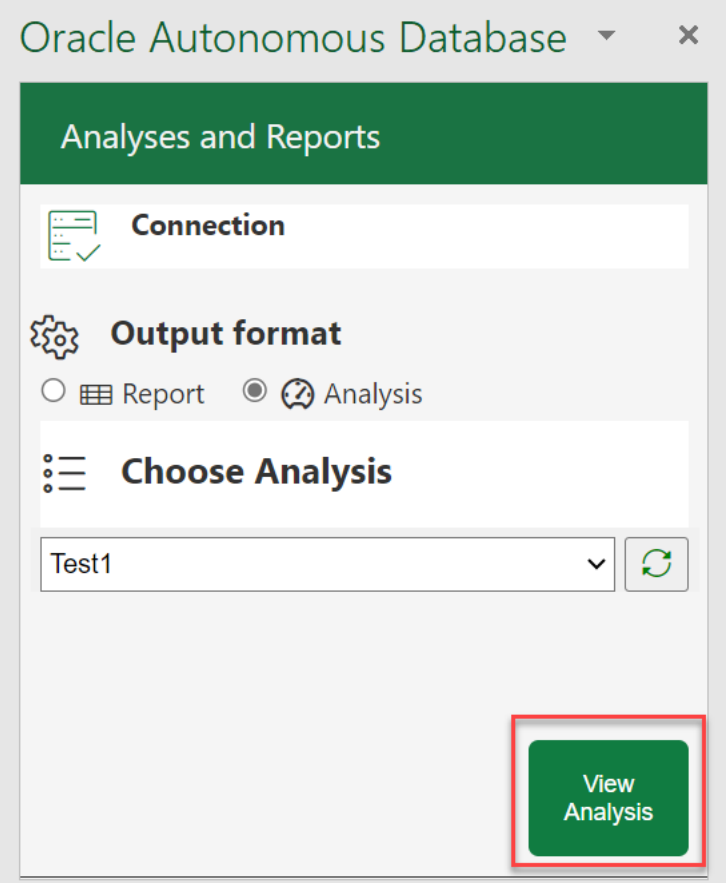
To view Analysis and explore the Analyses and Reports menu:

1. On the MS Excel ribbon, select the **Analyses and Reports** icon from the Autonomous Database menu.
2. Select **Analysis** under Output format.



3. Use the **Select an Analysis** drop down to choose the Analysis you want to view.

4. Click **View Analysis** to view analysis in a new excel sheet.

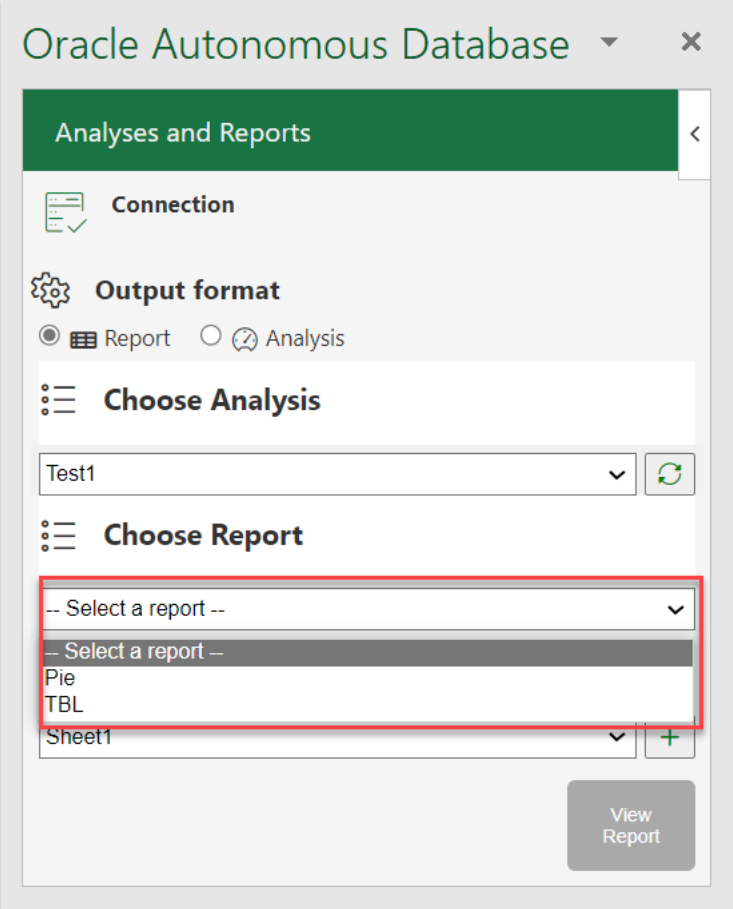


View Report

To view Reports :

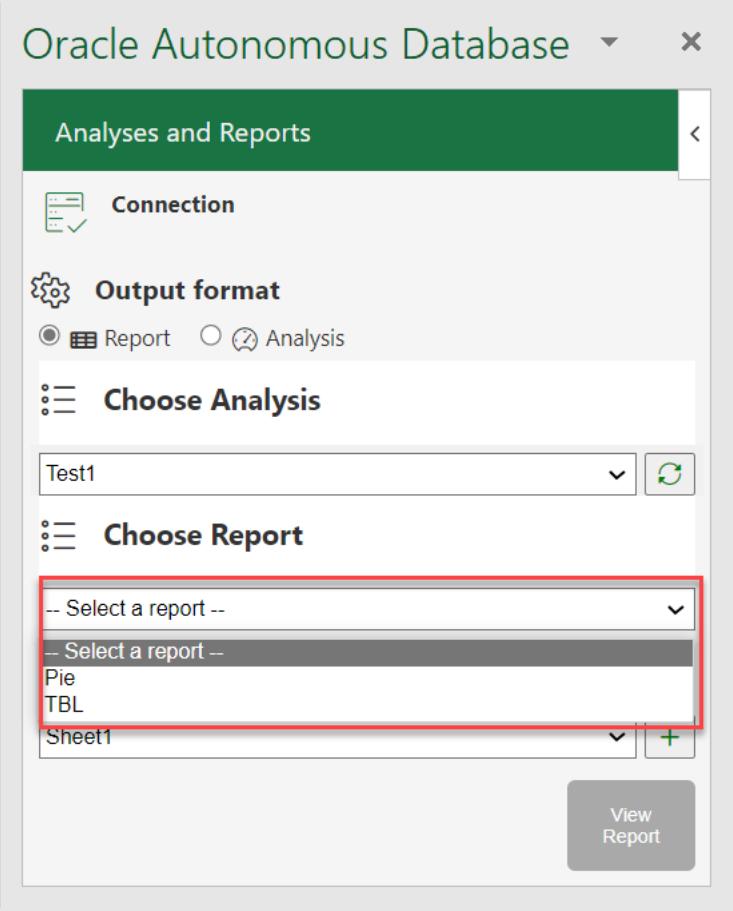
1. On the MS Excel ribbon, select the **Analyses and Reports** icon from the Autonomous Database menu.

2. Select **Report** under Output Format.



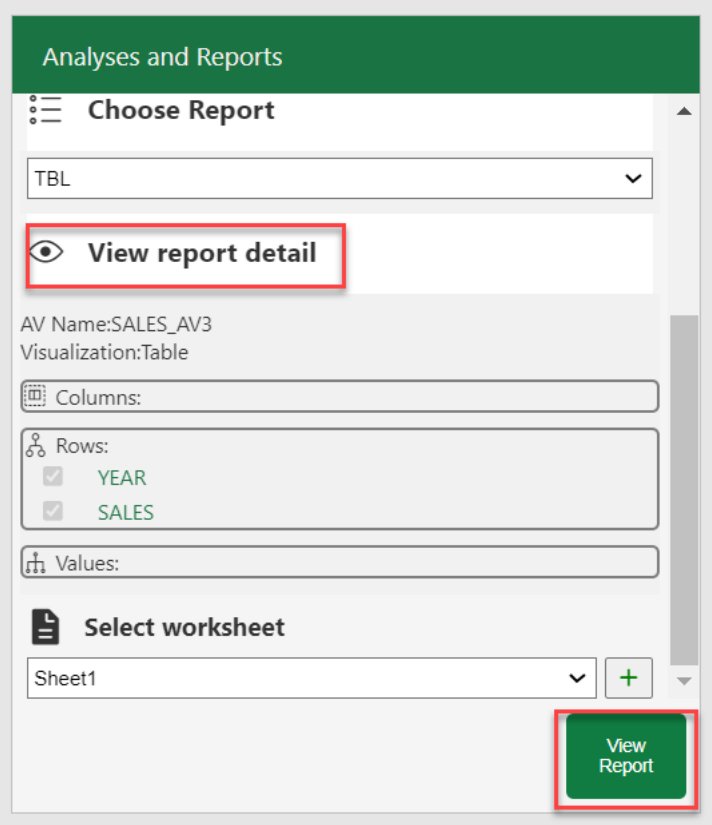
3. Use **Select an Analysis** drop down under **Choose Analysis** to choose the Analysis you want to view.

4. After you select the Analysis, to view the report present in the Analysis, click the **Select a report** drop-down and select the report you wish to view.



5. Click **View Report Detail** to view more information about the report namely, Analytic View Name, Type of visualization and rows, columns or values you select while creating the

report.



6. Select the worksheet from the drop-down where you would want to view the report.
7. Click **View Report** to view the report in the sheet you select from the previous step.

You can now view the report in the worksheet you select.

Query an Analytic View in an Excel worksheet

The Query Wizard menu enables you to query an Analytic View and retrieve the results in an Excel Worksheet. Once the wizard retrieves the data, it becomes local to Excel. You can further edit the data in Excel but not write back to the Autonomous Database.

> **Note:**
>
> This menu is unavailable if you install **Oracle Autonomous Database for Excel** using the Microsoft Store.

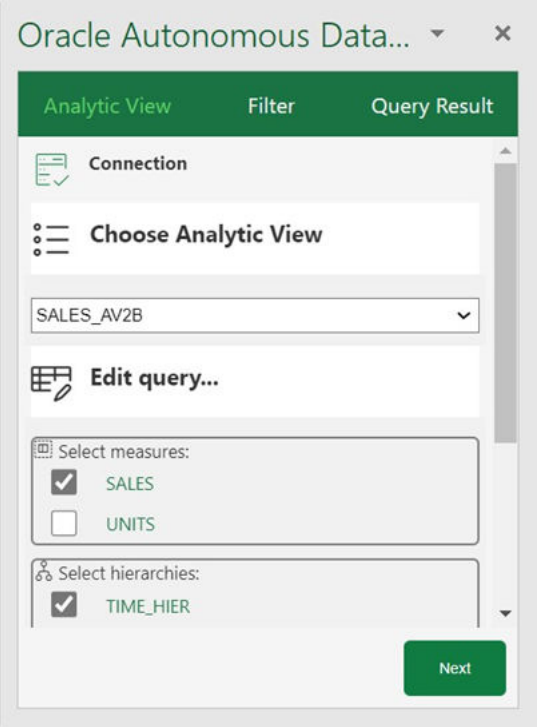
You can query an Analytic View to visualize the result data in the worksheet. You can search for the Analytic View, and select measures, hierarchies, and levels from the query. You can also add filters and calculated measures to the query and view the query result in the spreadsheet.

By default, the data is retrieved in tabular format. You also have the option to create an Excel pivot from this data.

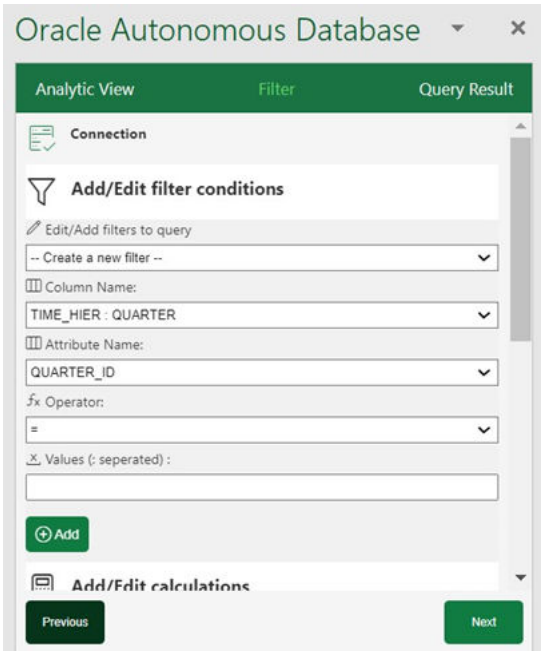
The Query Wizard has three panels:

1. **Analytic View panel:**The Analytic View panel contains a list of Analytic Views from which you build queries. You edit the query by selecting
   - measures
   - hierarchies,
   - and levels

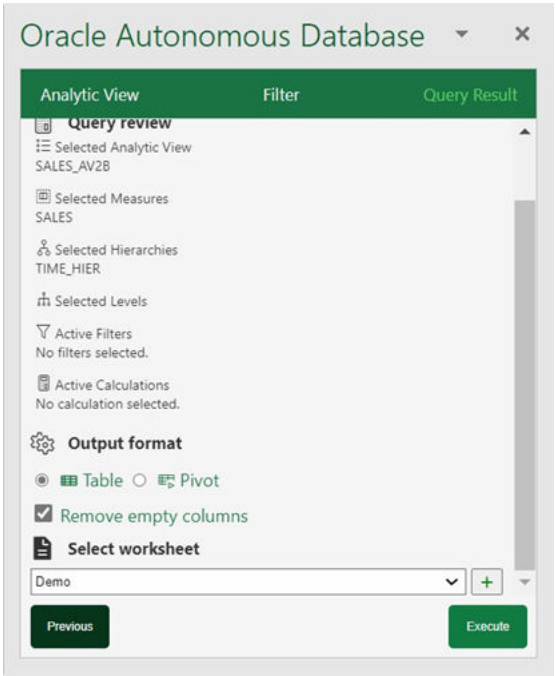
   and progress to the next panel.



2. **Filter panel:** The Filter panel displays to the right of the Analytic View panel when you click **Next** on the wizard. You can create filter conditions to filter the data and also add manual calculations to the Analytic View query in this panel.



3. **Query Result panel:**When you click **Next** on the wizard, the Query Result panel displays to the right of the Filter panel. You run the query once you select the filter criteria and determine what calculated measures to add to your query. You can view and revise the SQL query. After the SQL query runs, you view the query results in the worksheet. You can select the output format of the result here. You can view the results in tabular format or a Pivot table.



To query an analytic view and explore the Query Wizard menu in the MS Excel ribbon:

1. On the ribbon, select the Query Wizard icon.

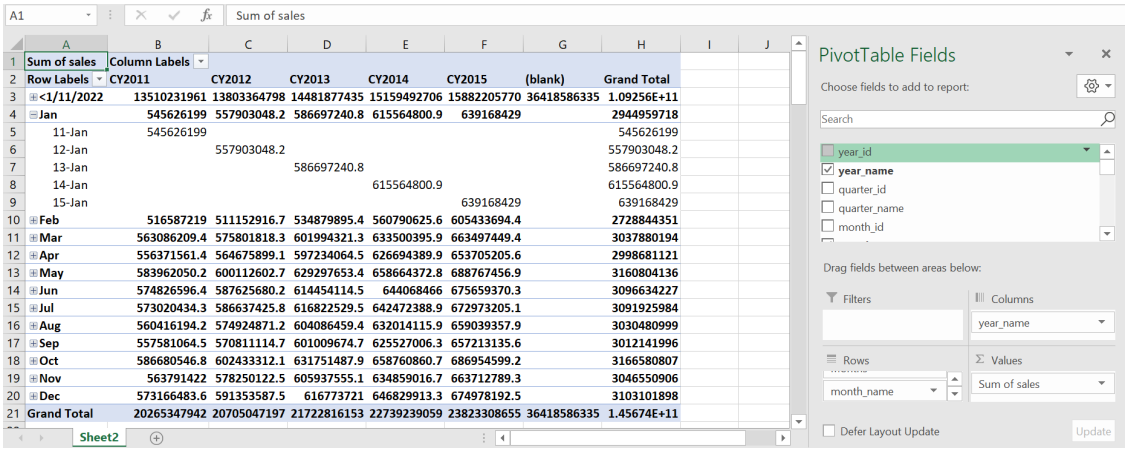
2. Selecting the Query Wizard opens an **Oracle Autonomous Database** dialog box in the Excel Task Pane.
3. Select an existing Analytic View from the drop-down in the Analytic View pane. As you select the Analytic View, it appears on the Analytic View field.
4. Select your choice of measures, hierarchies, and levels the available measures, hierarchies and levels associated with the Analytic View. Click **Next**.
5. The wizard window progresses to the **Filter** pane where you can add or edit filters to query.
6. Under Add or Edit filter conditions, do the following.
    - Select the column name and the attribute name from the drop-down- the values of the attribute change dynamically with the change in column names.
    - Select an operator in the Operator field to apply to the values that you specify in the Value field.
    - Specify a value or values from the list containing your selected column members. You need to enter the value into the Values field manually. For example, you can select > in the Operator field to use only values greater than the value that you select in the Value list. If you select 100,000 from the Value list, the filter uses values from the column greater than 100,000. You can use this information in an analysis to focus on products performing well. For multiple values use ":" as the separator.
    - Click **Add Filter** to add another filter condition.
7. Under Add or Edit Calculations, do the following.
    - Specify the column whose values you want to include in the group or calculated item.
    - On the Calc expression field, enter a custom calculated expression you want to perform on the column value. You can add functions or conditional expressions.
8. Click **Next** to progress to the **Query Result**.
9. You can view, edit, and review the query you have generated from the Query Review editor.
10. Select **Remove empty columns** to remove columns with no values returned in the result.
11. Select **Column per level** to retrieve all hierarchy levels in a single column.
12. Select the worksheet from the drop-down where you want to view the result.
13. Click **Execute** to run the query.
14. You can view the result of the query in the worksheet you select.
15. You can always modify the query in the Oracle Autonomous Database dialog box editor even after results are generated.
16. Select **Table** in the Query Result pane to view the results in the worksheet in a tabular format.
17. Select **Pivot** in the Query Result pane to view the results in a new worksheet in Pivot format.

View the results in Pivot tables

A Pivot table view is interactive and allows you to transpose rows and columns. A pivot table can summarize, sort, reorganize, count the total and perform an average of the result data. They are navigable and drillable.

Apart from tabular mode, to view the query results in pivot table mode, select the **Pivot Table** option in the Autonomous Database wizard. Click **Run** to view the query results in the Pivot table.

Clicking **Run** opens the query results in a new sheet with a PivotTable Fields wizard.

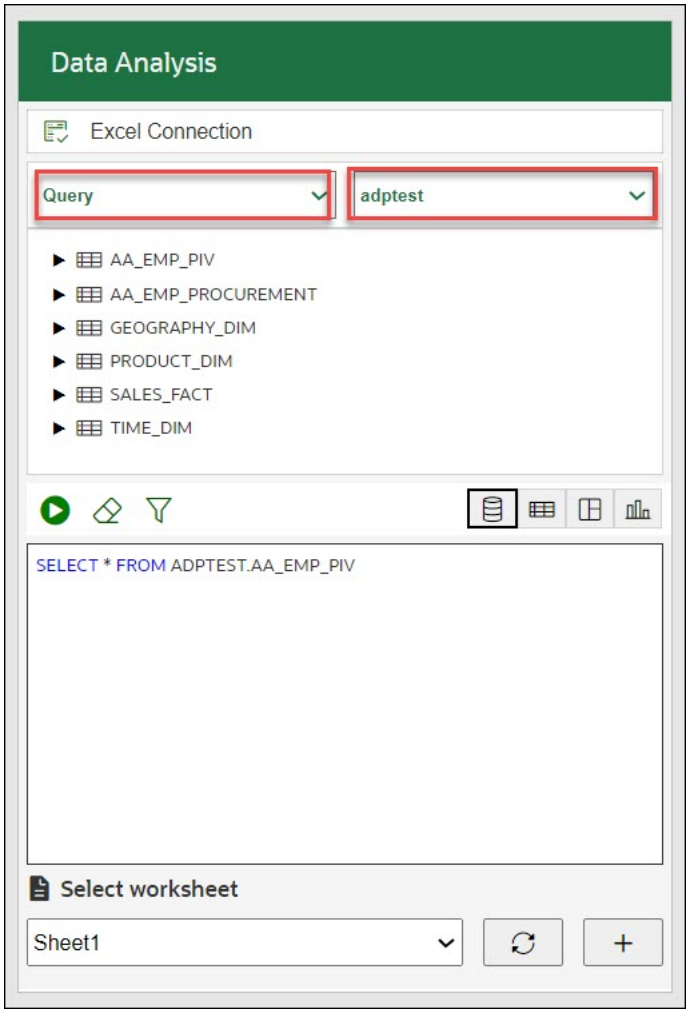


Data Analysis in Excel Sheet

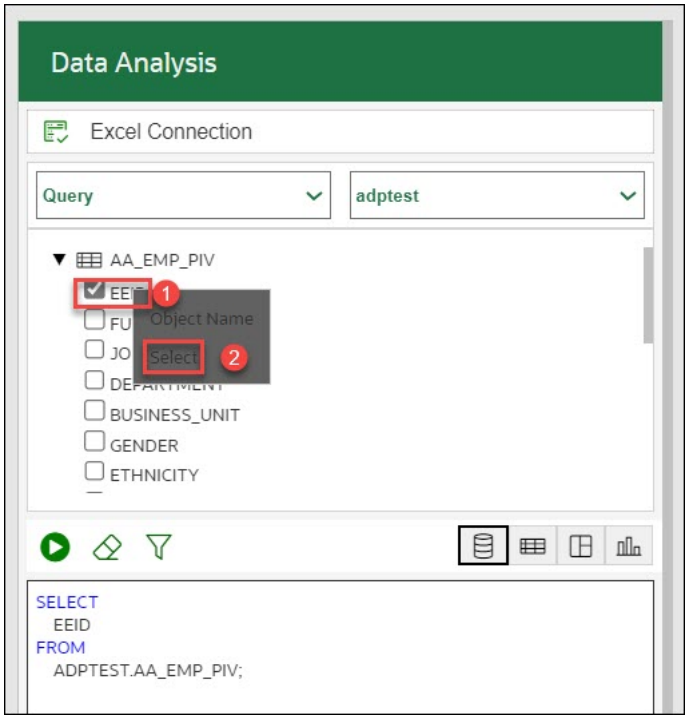
The Data Analysis tool enables you to analyze data in the Autonomous Database by running SQL queries or querying an Analytical View using an intuitive drag-and-drop interface. You can also write custom queries to be run. You now can dynamically apply filters to the result set retrieved using the new faceted search capability on specific columns.

Run query on database tables using the add-in:

1. On the Excel Sheet, select the menu item **Autonomous Database**.
2. Select **Data Analysis**. Selecting Data Analysis opens a Data Analysis panel. On the Data Analysis panel, select **Query** from the drop-down and the **schema** you want to use from the drop-down.



3. You can select a column of the table you want to query, right-click the column, and click **Select** to assist the add-in in forming a select query of the column from the table. Alternatively, you can drag and drop the selected column to the query area that enables the panel to produce a select query of the column in the query display area.



You will view the default query in the query editor area.

4. You can select any of the four modes to visualize the results of the SQL query report you generate.



- **Base Query**: This type of view is by default. The query written in the SQL editor is the Base Query. This query will be the base query of the other three modes.

> **Note:**
>
> All the numeric columns are displayed in different colours.

The Table browser now has all the columns from the base query.

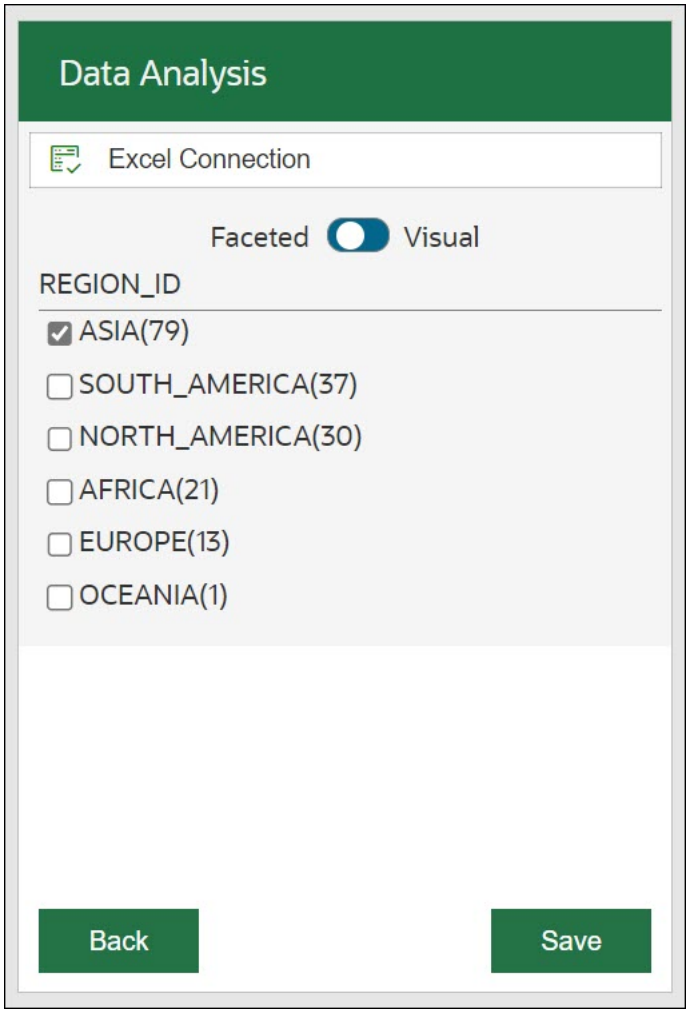
- **Table**: You can view the SQL results in tabular form. When you select this view, a column drop zone appears for selecting Rows, enabling you to drag and drop columns of the base query. Moving the selected columns in the drop zone allows you to view only those columns in the Result data generated in the worksheet. Select the cross mark beside the Column name to remove it from the drop zone.
- **Pivot**: You can view the SQL query results in pivot format. By selecting this format, an X and Y drop zone appears where you can drag and drop the selected columns from the Tables browser based on the Base Query to the drop area.

> **Note:**
>
> Only numeric values are allowed to be dropped in the values section.

- **Chart**: You can view *Area Chart*, *Bar Chart*, *Line Chart*, or *Pie Chart* when you select this option. The mappings displayed when you select one of the options are as follows:
  - **Orientation**: Choose between horizontal and vertical orientation types from the drop-down list.
  - **X axis label and Y axis label**: Optionally enter labels for X axis and Y axis.

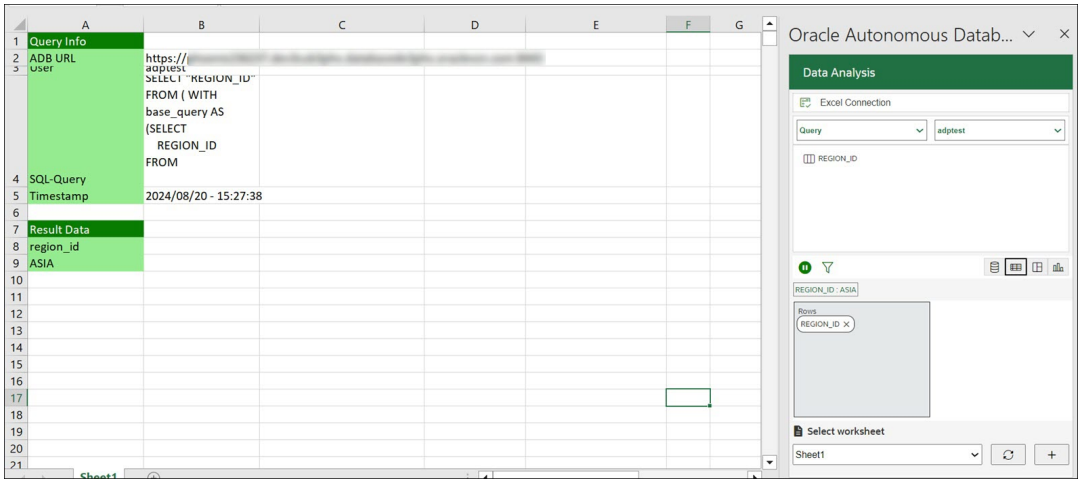
> **Note:**
>
> Only numeric values are allowed to be dropped in the Y axis drop zone.

5. Click the **funnel icon** (Faceted Filter) to add filters to the result. The panel generates a filter for each value in the column retrieved from the query result. You can filter different columns on the faceted filter panel and view the results in the worksheet to view only the data you wish. For example, to view the customer reports by Region, click the **faceted filter** and select **Asia** under `Region_ID`. The number in the brackets displays the count of the items with this property. In the example shown below, there are `79` records with the region as `ASIA`.



6. Click **Save** to view the results. Click **Back** to go back to the main panel.

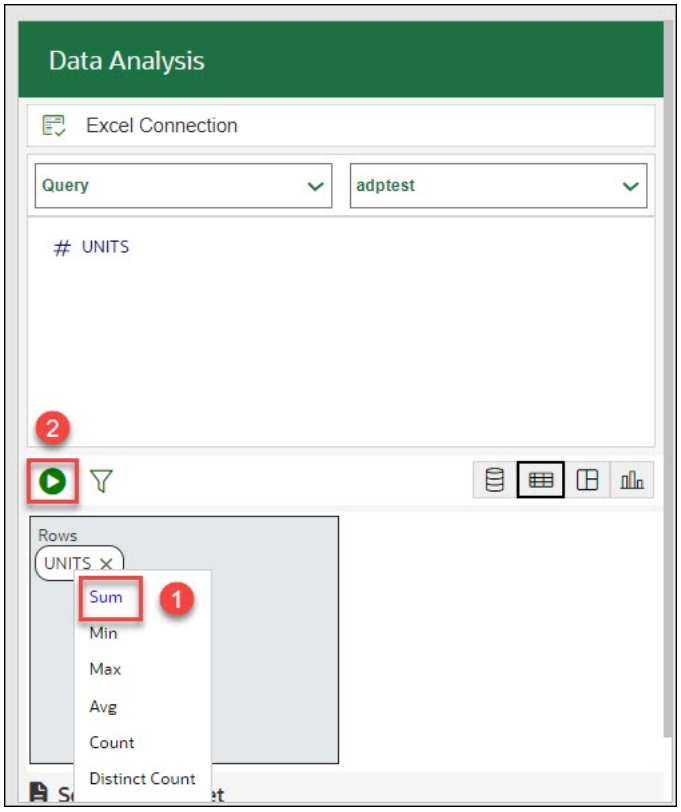
7. Select **Run** to generate the results of the custom query in the worksheet. Click **Pause** to make any changes to the query, such as updating the columns of the table without updating the worksheet.



Perform aggregate functions using the Excel add-in

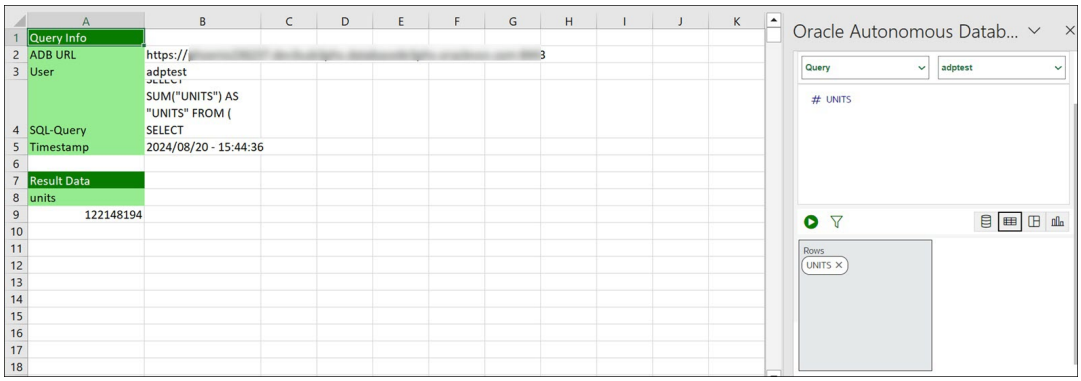
You can also perform aggregate functions such as `SUM`, `MIN`, `MAX`, `AVG`, `COUNT`, and `DISTINCT COUNT`. In this example, we're primarily going to focus on using the Data Analysis feature to gain insights from our sales data.

1. Select **Data Analysis**. Selecting Data Analysis opens a Data Analysis panel. On the Data Analysis panel, select **Query** from the drop-down and the **schema** you want to use from the drop-down.
   Drag and drop the sales value to the query editor and click **Table** to view the sales in tabular format.
2. To calculate the sum of the sales value, click the sales value and select **Sum** from the list of available aggregate functions.



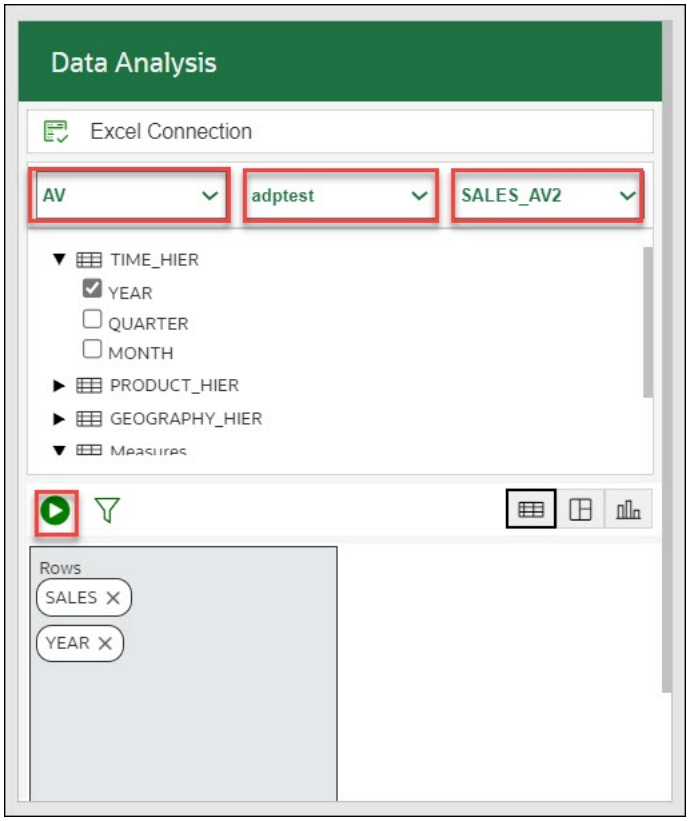
Click **Run** to generate the sum of sales amount in the tabular format.

You will view the result generated in the excel worksheet.



To query an Analytic View and explore the **Data Analysis** menu in the Excel Sheets:

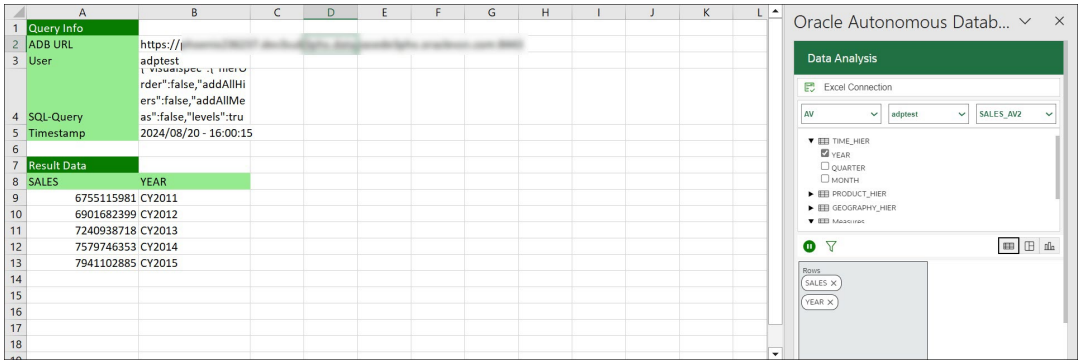
1. Select the ribbon item **Autonomous Database > Data Analysis** on the Excel Sheet. This opens a **Data Analysis** panel.
2. Select **AV** from the drop-down, select a schema you can access from the schema drop-down, and the Analytic View from the list of available Analytic Views.



3. You can select any of the three modes to visualize the results of the AV query you generate:
   - **Table**: You can view the AV query results in tabular form. By selecting this view, a column drop zone appears, enabling you to drag and drop selected columns from the Table browser. Moving the selected columns in the drop zone allows you to view only those columns in the Result data generated in the worksheet. Select the cross mark beside the Column name to remove it from the drop zone.
   - **Pivot**: You can view the SQL query results in pivot format. By selecting this format, an X and Y drop zone appears where you can drag and drop the selected columns from the Tables browser to the drop area.
   - **Chart**: You can view the results of the AV query in chart format. By selecting this format, an X and Y drop zone appears where you can drag and drop the chosen hierarchies and measures from the AV browser to the drop area.

   > **Note:**
   >
   > You are allowed to drop measures in the Y-axis.

4. Click **Run** to view the results in the worksheet . You can view the total Sales generated along with it's year of generation.



Natural Language in Excel Sheet

You can use Natural Language Query to query the Oracle Autonomous Database using the **Natural Language** menu in the **Oracle Autonomous Database for Excel** .

Prerequisites

- Before you use the Natural Language menu in the add-on, you must perform prerequisites to use `DBMS_CLOUD_AI` to configure AI profiles.
- Create and Set an AI Profile

After your AI Profile is ready, the Data Studio tool uses these profiles in the Data Studio Settings menu to configure access to a Large Language Model (LLM) and to setup for generating, running, and explaining SQL based on natural language prompts. This also enables chatting with the LLM.

> **Note:**
>
> You can run Natural Language Queries on Tables and not Analytic Views.

Generate SQL Queries from Natural Language on Tables

Using natural language to interact with your database data is now achievable with Oracle Autonomous Database add-on for Google Sheets.

This means you can use natural language, for example, plain English, to query the database. You can provide a natural language prompt instead of SQL code to interact with your data. When you select the **Generate SQL from natural query** icon, the add-on converts natural language to SQL.

To run a natural query using the add-on:

1. On the Excel Sheet, select the menu **Oracle Autonomous Database**.
2. Select **Natural Language**. Selecting **Natural Language** opens a **Natural Language** wizard.
3. On **Natural Language** wizard, select the **Table** on which you perform the query from the drop-down and the **AI profile** you want to use from the drop-down.

> **Note:**
>
> This is the same AI profile you configure in the Data Studio Settings.

4. Let's say you want details of the product with highest sales. Enter the following natural language query in the query display area:

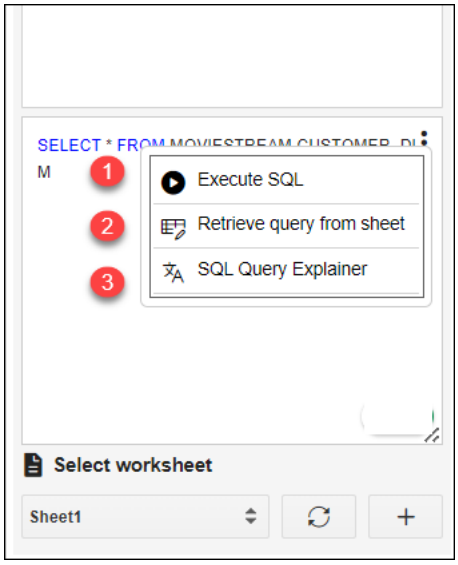
```
show movie with highest sales
```

5. Select **Generate SQL from natural query** to produce the equivalent SQL query in the bottom query display area.
You will view the following code in the bottom SQL code area.

```
SELECT
    *
FROM
    "MOVIESTREAM"."MOVIESALES_WEEKEND_USA" MS
WHERE
    MS."SALES_AMT" = (
        SELECT
            MAX("SALES_AMT")
        FROM
            "MOVIESTREAM"."MOVIESALES_WEEKEND_USA"
    )
```

6. Selecting the **Generate SQL from natural query** icon also displays the result of the natural language query in the Google Sheet.
You can click the **+** sign beside the **Select worksheet** drop-down to display the results in a new worksheet.
7. Click **Clear query content** to clear the content displayed on the natural language query area and the SQL code query area.

Actions you can perform from the SQL code area:



- **Execute SQL**: Runs SQL code and displays the query result in Google Sheet.
- **Retrieve query from Sheet**:
  You can retrieve SQL query displayed in the Google Sheet to the SQL code area in the Natural Query.

  This feature is to retrieve the latest data from the database for a future date. For example, if you want to receive the latest data from the table the next day, the connection may be inactive, but you can still connect to the instance, **retrieve query from Sheet** and **select Execute SQL** to display the latest query results in the Google Sheet.

- **SQL Query Explainer**: Explains SQL code in Natural language form.
  Select this icon to translate SQL queries to natural language that is understood by you. The natural language query is displayed in the natural language query display area of the wizard.

Ask Questions with Chats

Chats option in the Oracle Autonomous Database for Google Sheets refer to an interactive conversation between you and the add-on where the add-on uses natural language to query or interact with the autonomous database.

The data we use in this example is of a company called `Oracle MovieStream` to analyze movie sales data.

The Chats displays recommendations for the default table you select.

Let's find out the top 5 streamed movies in the Moviestream company by weekend gross:

1. Enter the text in the **Start typing your question..** text field.

   ```
   What are the top 5 movies by Weekend Gross?
   ```

2. Click **Enter**.
3. The Chat displays top 5 movies by Weekend Gross.

   It also displays equivalent SQL code of the result.

   Click **SQL** to expand the SQL.

Click **Copy to clipboard** to copy the SQL.

Click **Run SQL** to run the generated SQL query and display the query result in Google Sheet.

The **Chats** option remembers the context of previous chat history.

For example, if you enter `now show me top 10` in the text field.

It displays the top 10 movies by Weekend Gross. It remembers which metric top 10 to fetch without us having to type the whole thing.

Select **New Chat** for deleting the present conversation.

Select **Home** to go back to the main home page.

> **Note:**
>
> LLMs are remarkable at inferring intent from the human language and they are getting better all the time; however, they are not perfect! It is very important to verify the results.

FAQs for Troubleshooting errors with Excel Add-in

If you experience issues with Oracle Autonomous Database for Excel, refer to frequently asked questions in this section to identify and resolve issues.

1. Why is the **My Add-ins** icon from the **Insert** ribbon in the MS Excel workbook greyed out?

   Even before installing the Excel add-in, sometimes the **My Add-ins** icon from the **Insert** ribbon in the MS Excel workbook appears greyed out.

   a. From the **File** menu in the Excel ribbon, go to **Account** and select **Manage Settings** from the Account page.

   b. Ensure that you select the **Turn on optional connected experiences**.

   c. From the **File** menu in the Excel ribbon, go to **Options** and select the **Trust Center** option from Excel Options.

   d. Click **Trust Center Settings** and ensure that deselect **Disable all Application Add-ins** (if selected) from the Add-ins tab in the Trust Center dialog box.

   e. Select the **Trusted Add-in Catalogs** menu from the Trust Center dialog box and ensure that you deselect the **Don’t Allow any web add-ins to start** checkbox.

2. Why doesn’t the sign-in page of the Excel Add-in load or appear?

   Sometimes you might encounter issues with the Excel Add-in even after loading it correctly. For example, an add-in fails to load or is inaccessible. Check the compatibility version of Excel and the operating system you use.

   If the compatibility is correct and the sign-in page to the Excel Add-in still does not show up, or does not load properly, we recommend applying all pending Windows, Office, and browser updates.

   a. Select **Settings**, **Update** & **Security**, and then **Windows Update** from the Windows Start menu.

   b. If updates are available on the Windows Update page, review the updates and click **Install Now**.

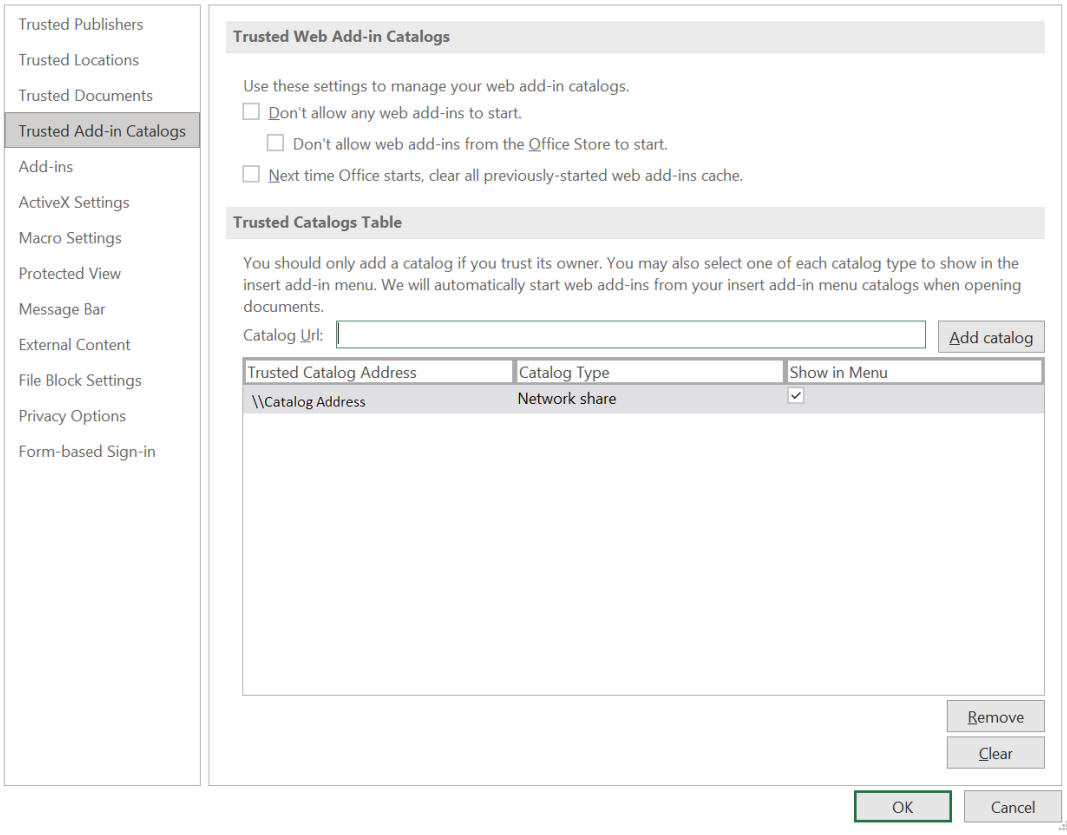
> **Note:**
>
> The details of applying Windows updates can vary from version to version and if required, check with your system administrator for assistance.

3. Why doesn't the add-in work correctly after re-installing?

   Configure the Excel trusted Add-in catalog to set the add-in correctly after re-installation.

   To configure the Excel add-in, check or remove the add-in if it is pointing at the wrong location in the Trusted catalog address. This address should be the same as the shared manifest folder's location (share path).

   Click Excel's **File> More> Options>Trust Center >Trust Center Settings> Trusted Add-in Catalogs**



   Checking is only required the first time you use the installer, or if the shared manifest folder is changed. The change occurs during uninstalling and re-pointing to a new ADB.

   To remove the catalog from the trusted table and add a new catalog pointing to a different address:

   - Select the Catalog you want to remove from the trusted catalog table and click **Remove**.
   - Enter the correct share path of the shared manifest folder in the **Catalog URL field** and click **Add Catalog** to add the shared folder to the trusted catalog.

   Restart Excel to make the new shared folder active to access the add-in.

4. Why doesn't the add-in work correctly even after configuring the Excel trusted Add-in catalog?

   Let's say you configure the Excel trusted add-in catalog after re-installing the add-in but, it does not load correctly. Sometimes the database server changes are not reflected in Excel even after you set the share path of the shared manifest folder as a trusted add-in catalog. Clear the Office cache to resolve this issue.

   Refer to this page: https://docs.microsoft.com/en-us/office/dev/add-ins/testing/clear-cache#clear-the-office-cache-on-windows to clear the Office cache on Windows and Mac.

   Clearing the Office cache unloads the Excel add-in. Install the add-in and check the configuration of the Excel trusted add-in catalog. This should solve the issue of the Excel add-in needing to be correctly loaded.

5. What should you do if you cannot view My Add-ins icon from the Excel ribbon?

   After installing the add-in once, if you cannot view the **My Add-ins** icon from the Insert ribbon and instead you view the Add-ins icon 

.

   Follow these instructions:

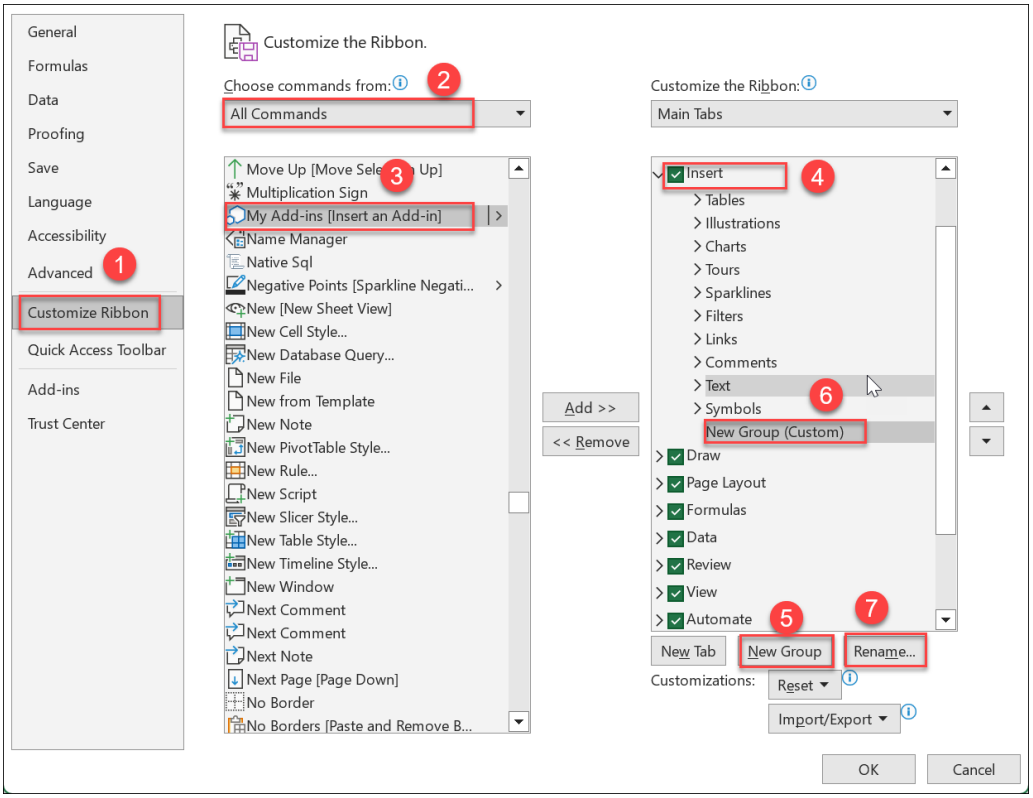
   a. From the **File** menu in the Excel ribbon, go to **Options** and select the **Customize Ribbon** option from Excel Options.

   b. Select **All Commands** drop-down from the **Choose commands from** drop-down.

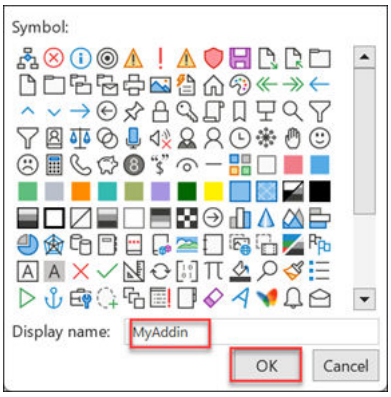
   c. Click **My Add-ins** and click **Insert** option from the Main Tabs list.

   d. Click **New Group** to add a new menu item under the **Insert** menu. You will view a **New Group (Custom)** menu option added to the Insert menu.

   e. Click **Rename** to rename the newly created menu.

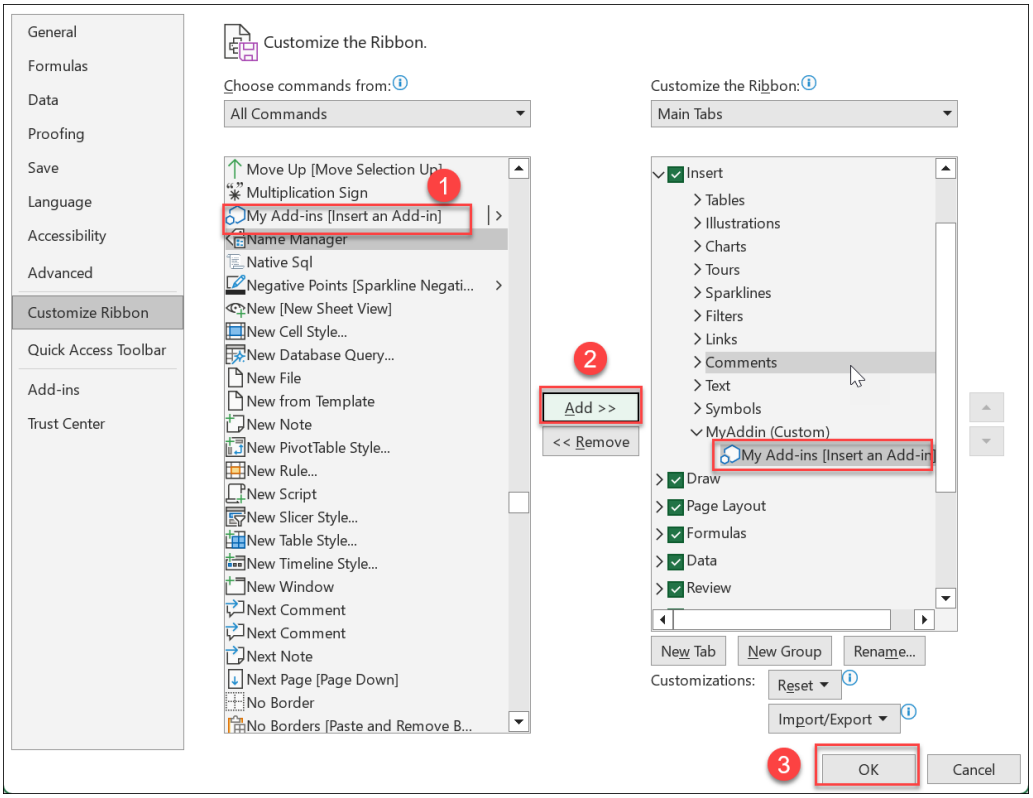


f. Specify the name of the add-in. For example, **My Addin**. Click **OK**.

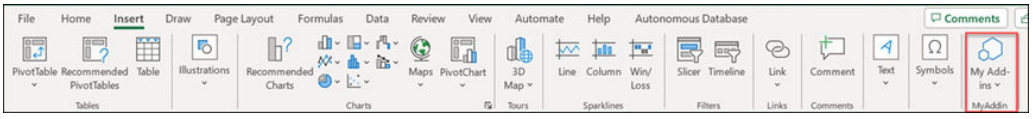


g. Click **My Add-ins[Insert and Add-in]** from All Commands list and select **Add**. My Add-ins menu is added to the newly created menu “MyAddin”.

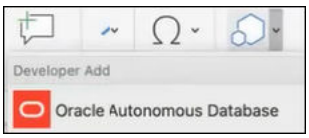
h. Click **OK** to save the changes.



i. Clicking **OK** takes you to the main Excel sheet page where you can view “**MyAdd-ins**” menu under the Insert menu.



j. Click **My Add-ins**. You can now view the Oracle Autonomous Database for Excel loaded.



6. What happens when you cannot view the latest added menu items in the **Oracle Autonomous Database for Excel** add-in?

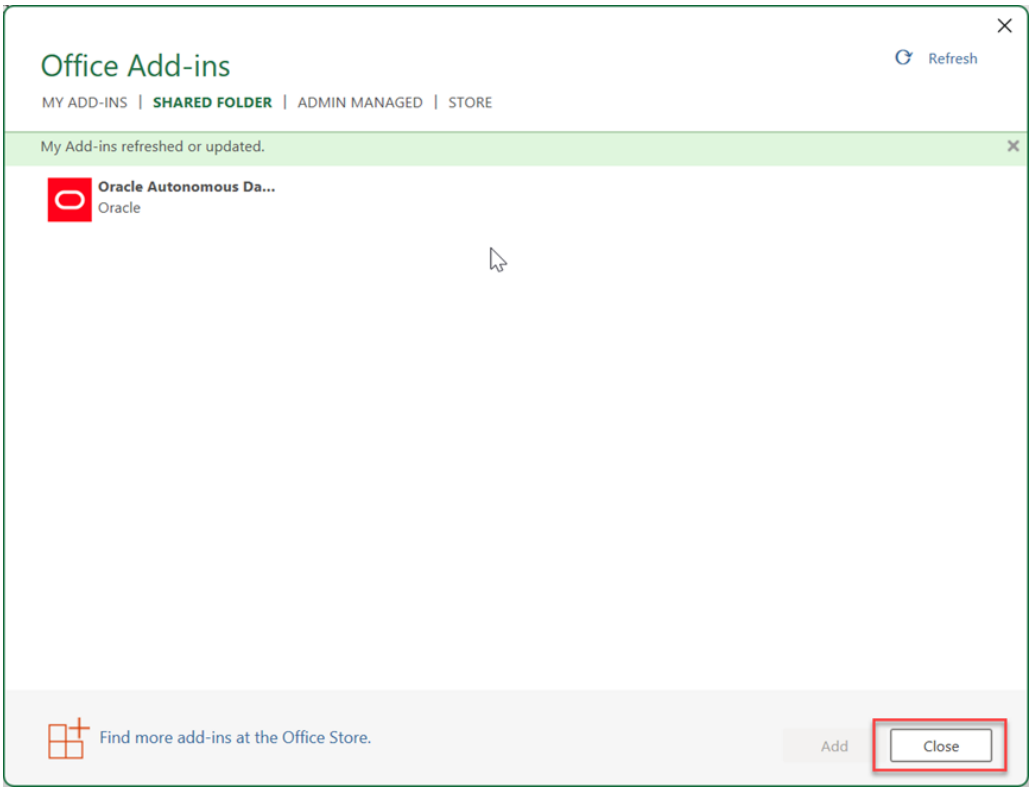
   Sometimes when you cannot view the changes updated in the latest version of the plug-in, you must:

   a. Select **My Add-ins** from the Insert menu on the Excel ribbon.

   b. Click the **Shared Folder** tab of the Office Add-ins dialog box. You will see the add-in under the Shared Folder list.

   c. Click **Refresh**.
   After the add-ins are refreshed, you will receive a notification on the dialog that says "*My Add-ins refreshed or updated*". The refresh button loads the manifest file again for the latest changes to appear.

   d. Click **Close** to close the dialog box.



Oracle Autonomous Database for Excel Privacy Policy Details

This topic covers details for writing policies to control access to Autonomous Database resources.

Data use by **Oracle Autonomous Database for Excel** complies with Oracle's Privacy Policy.

Oracle Autonomous Database for Google Sheets

The Oracle Autonomous Database add-on  enables you to query tables using SQL or Analytic Views using a wizard directly from Google Sheets for analysis.

The data retrieved from the Autonomous Database is available locally in Google Sheets for further analysis. The results are stored in the local copy and cannot be written back to the Autonomous Database. You can run direct SQL queries or query Analytic Views and view their results in the worksheet. The add-on allows you to filter the query results, and perform table joins and calculations.

> **Note:**
>
> The Oracle Autonomous Database add-on for Google Sheets must comply with Privacy Policy. For information on details of privacy policy, see Oracle Autonomous Database for Google Sheets Privacy Policy Details.

How does the add-on for Google Sheets work?

To query an Analytic View or Tables from the Autonomous Database, you must select an Analytic View or Table to work with. While retrieving data from the Analytic View, you can configure the query according to your requirements. You can select specific hierarchies and create custom calculations on the wizard. The add-on configures your query and returns the result to the Google Sheets. You can save the results of your queries locally in the Google Sheet. The add-on can also query the schema directly to which you have access. Using the Web UI, you can also view reports and analyses you create in the Data Analysis menu in the Data Studio tool.

To use the add-on, you must enable Web Access on the Autonomous Database account. You must have the CONNECT, DWROLE, RESOURCE and ADPUSER roles grant in the SQL worksheet to access the Google Sheets add-on.

> **Note:**
>
> - The **Download Microsoft Excel/Google Sheets add-in** is available to you under the **Downloads** menu of your Database Actions instance only if you have the `ADPUSER` role.
> - The Oracle Autonomous Database add-on for Google Sheets is not supported in Safari web browser.

- Install and setup the add-on for Google Sheets
  Before you install the Oracle Autonomous Database add-on for Google Sheets, download the *oracleGoogleAddin* zip file from your Database Actions instance.

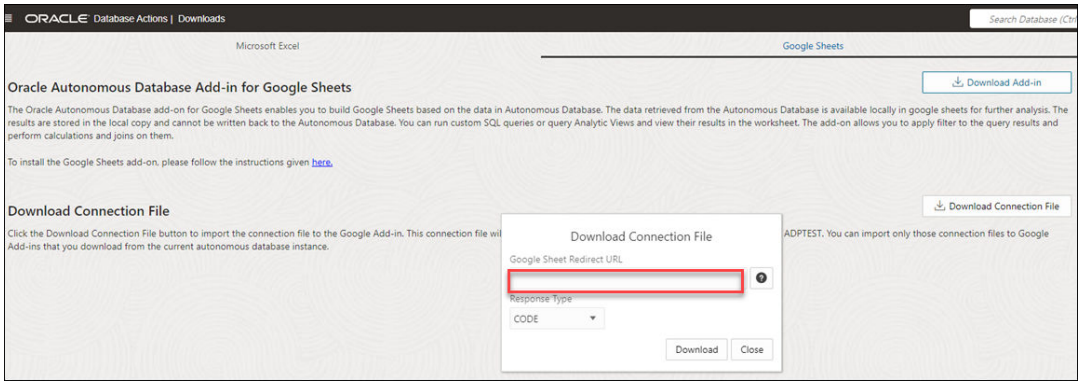
- Download Connection File
  To connect to the Autonomous Database, you can download a connection file from the Database Actions instance and import it to the Google Sheet add-on you have setup.
- Connecting to Autonomous Database
  The Oracle Autonomous Database add-on for Google Sheets enables you to connect to multiple Autonomous Databases with a single add-on using the Connections feature. The add-on connects to Google Sheets by providing authentication to Google. Multiple users or databases can connect simultaneously to the add-on. However, only one connection can remain active.
- Generate Client ID and Client Secret using UI
  In this section you use the Web UI to obtain the `client_id` and `client_secret` .
- Authorize Google Sheets to use Autonomous Database
  After your identity is determined using OAuth authentication, Google Sheets needs permission to access the Autonomous Database.
- Natural Language in Google Sheets
  You can use Natural Language Query to query the Oracle Autonomous Database using the **Natural Language** menu in the **Oracle Autonomous Database for Google Sheets** "add-on".
- Data Analysis in Google Sheets
  Selecting **Data Analysis** opens an Oracle Autonomous Database wizard in the Google sheet.
- Run Direct SQL Queries
  The Oracle Autonomous Database add-on for Google Sheets lets you run SQL queries to work with your data in a Google Sheet. With the add-on, you can type your SQL code in the SQL editor area and click **Run** to run the command.
- Reporting and Analysis in Google Sheets
  You can view Reports and Analytic Views or visualize data for analysis purposes.
- Clear Sheet
  Once the add-on runs the query and retrieves the data into the worksheet, you can view the Timestamp, User, AV-query and SQL-query of the Analytic View in the automatically generated query results.
- Delete all sheets
  Use this option to delete all the sheets existing in the spreadsheet.
- About Oracle Autonomous Database menu
  Use this option to view details about the add-in
- Share or Publish
  Once you generate the query results in the Google Sheet, you can share it with other users. With sharing, creates a copy of the worksheet and sends it with the design tools hidden and worksheet protection turned on.
- Oracle Autonomous Database for Google Sheets Privacy Policy Details
  Effective Date: `10/4/2024`
- Oracle Autonomous Database for Google Sheets Support
  Welcome to the support page for the **Oracle Autonomous Database for Google Sheets**. This resource is designed to assist you with any issues or questions you may have while using the add-on.

Download Connection File

To connect to the Autonomous Database, you can download a connection file from the Database Actions instance and import it to the Google Sheet add-on you have setup.

Follow the steps shown below to download the connection file.

1. Navigate to the launchpad of your Database Actions instance, and select the **DOWNLOAD MICROSOFT EXCEL/ GOOGLE SHEETS ADD-IN** Card. Click the **Download Connection File** button in the Google Sheets tab of the **Downloads** page to import the connection file to the Google Add-in.
2. This connection file will allow you to connect to the Autonomous Database with the logged-in user. You can import only those connection files to Google Add-ins that you download from the current Autonomous Database instance.



3. Selecting the Download Connection File button opens a Download Connection File wizard. Specify the following field values in the wizard:
   - **Google Sheet Redirect URL**: This is the **Web application deployment URL** you copied from step number nine of Deploy the Google Script as a Web app section.
   - Choose a **Response Type**:
     - **Explicit Connection**
       You use the OAuth Client ID and OAuth Client Secret values to authenticate and authorize Google Sheets to use the Autonomous Database. Use this when you use **CODE** as the Response Type while downloading the connection file from the Database Actions page. This is the more secure method and is preferred to use if the Autonomous database has public access.
     - **Implicit Connection**
       You will need an OAuth Client ID to implicitly access the Autonomous Database. Use this when you use Token as the **Response Type** while downloading the connection file from the Database Actions page. Use this when the autonomous database is in a private subnet or within a customer firewall.

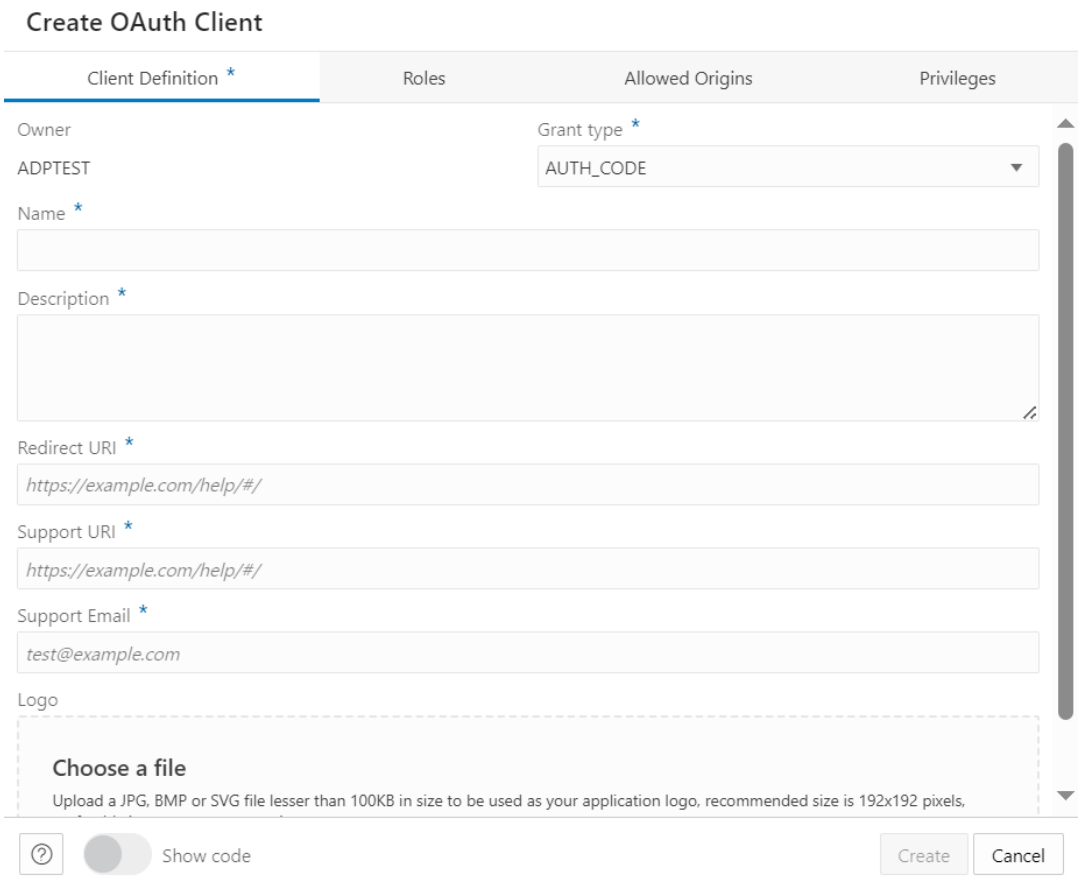
Generate Client ID and Client Secret using UI

In this section you use the Web UI to obtain the `client_id` and `client_secret` .

You generate the client keys by accessing the Autonomous Database instance URL appending with oauth/clients.

For example, if your instance is “ https://<hostname>-<databasename>.adb.<region>.oraclecloudapps.com/ords/<schema Name>/_sdw/", you need to sign in to the link " https://<hostname>-<databasename>.adb.<region>.oraclecloudapps.com/ords/<schema Name>/oauth/clients/". Be sure to include the trailing slash.

1. Sign in to Database Actions with the "https://machinename.oraclecloudapps.com/ords/SchemaName/oauth/clients/" link. You can view an OAuth Clients page in the link "https://localhost:port/ords/schemaName/_sdw/?nav=rest-workshop&rest-workshop=oauth-clients".
2. Click the **+Create OAuth Client** button to create a new client.

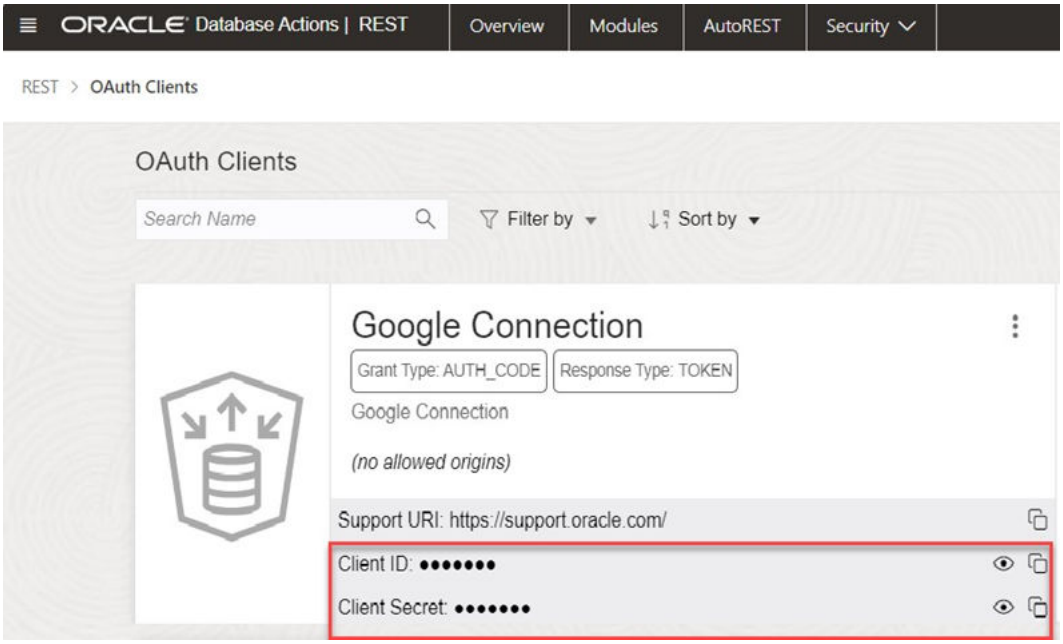


3. From the Grant type drop-down, select the type of client connection you want. You can select the following options:
   - AUTH_CODE: Select this option for implicit connection. Use this response type when the autonomous database is in a private subnet or within a customer firewall.
   - IMPLICIT: Select this option for explicit connection. This is the more secure method and is preferred to use if the Autonomous database has public access.
4. Enter the following fields. The fields with an asterisk (*) are mandatory:
   - **Name**: Name of the client.
   - **Description:** Description of the purpose of the client.
   - **Redirect URI**: web application deployment URL you copied from step 10 of Deploy the Google Script as a Web app

- **Support URI**: Enter the URI where end users can contact the client for support. Example: *https://script.google.com/*
- **Support Email**: Enter the email where end users can contact the client for support.
- **Logo**: Optionally, select an image from your local system to insert a logo for your new client.

Navigate to the **Roles** tab to select the roles of the client. This is not a mandatory field.

5. Progress to the **Allowed Origins** tab. Specify and add the list of URL prefixes in the text field. This is not a mandatory field.
6. Progress to the **Privileges** tab to add any privilege. You are not required to have any privileges to create an OAuth Client.
7. Click **Create** to create the new OAuth Client. This registers the OAuth Client which you can view on the OAuth Clients page.

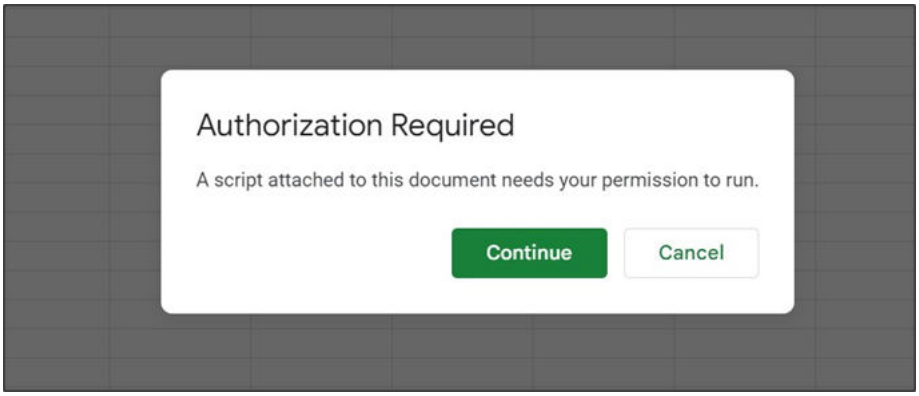


8. Click the show icon to view the `Client ID` and the `Client Secret` fields.

How do I connect manually?

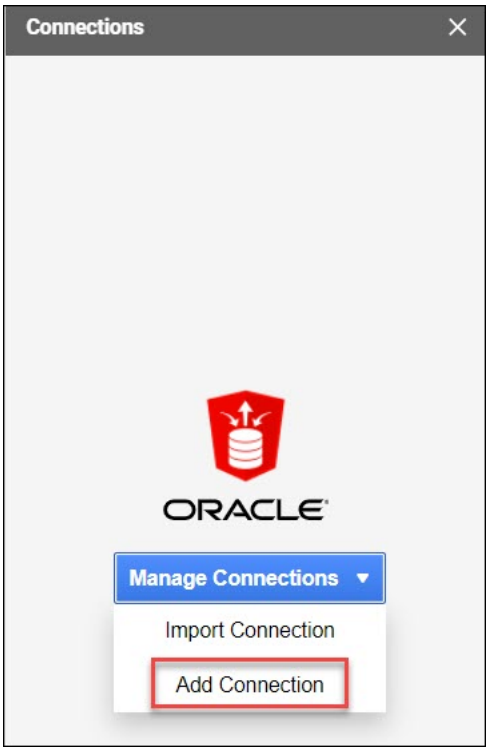
The following sections demonstrate how to connect using implicit and explicit connections. Google Sheets needs permission to access the Autonomous Database. You must first complete the authorization to connect to the autonomous database. The add-on requires one-time authentication for the setup.

1. On the Google Sheet, click **Oracle Autonomous Database** and select **Connections**. Selecting Connections requires one-time Google authentication.
   - Clicking **Connections** opens a pop-up window that asks your permission to run the authorization. Click **Continue**.

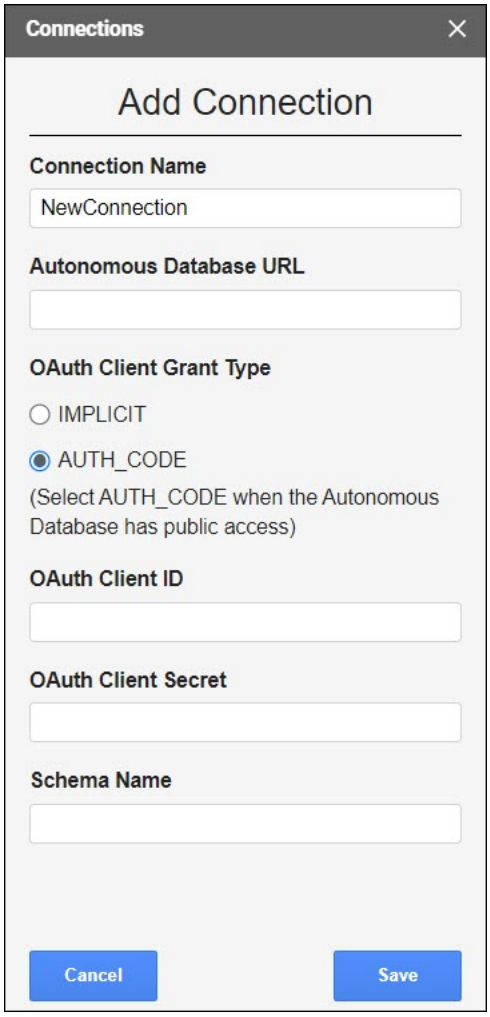


- You will now view a window that informs you that the application requests access to sensitive information in your Google account.
- Click **Advanced** and select the **Go to Untitled project (unsafe)** link. Selecting the link opens new window, ensuring you trust the application. Click **Allow** to continue. You have now completed the setup.

2. On the Connections wizard, click **Add Connection** from the **Manage Connections** drop-down menu to add a connection.



3. Selecting Add Connection opens an **Add Connection** wizard in the Connections wizard's connection list panel.



4. Specify the following field values in the wizard:

   Connection Name: Enter the connection's name—for example, *TestConnection*.

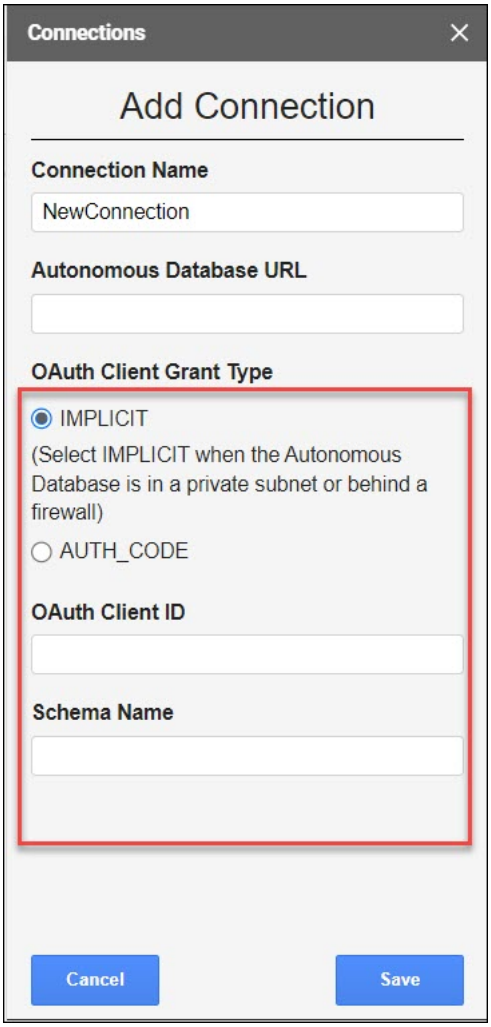
   Autonomous Database URL: Enter the URL of the Autonomous Database you wish to connect to. For example, "*https://<hostname>-<databasename>.adb.<region>.oraclecloudapps.com/*"

   In the OAuth Client Grant Type field, select one of the two options based on the type of connections you want. Refer to the Generate Client ID and Client Secret using the UI section.

   This option varies with implicit and explicit connections.

   Implicit: Select this option for implicit connection. Use this response type when the autonomous database is in a private subnet or within a customer firewall.
   AUTH_CODE: Select this option for explicit connection. This is the more secure method and is preferred to use.

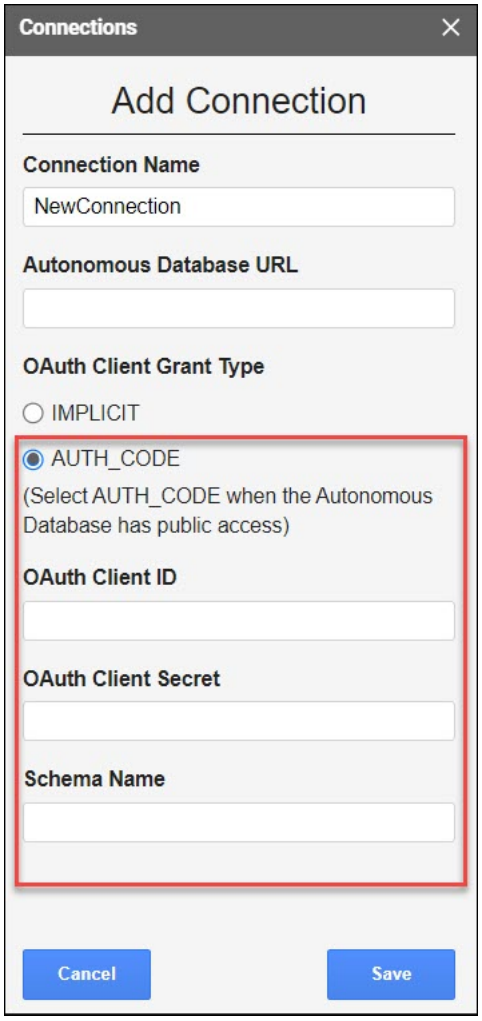
   When you select the Implicit option, you can view the following fields:



OAuth Client ID: `client_id` you generate using the *Create New Client* wizard in the UI. Refer to the Generate Client ID and Client Secret using the UI section.

Schema Name: Specify the name of the schema.

When you select AUTH_CODE, you can view the following fields:



OAuth Client ID: `client_id` you generate using the Create New Client wizard in the UI. Refer to the *Generate Client ID and Client Secret using the UI* section.

OAuth Client Secret: `client_secret` you generate using the Create New Client wizard in the UI. Refer to the *Generate Client ID and Client Secret using the UI* section.

Schema Name: Specify the name of the schema.

Click **Save**.

After you click **Save**, you can view the new connection in the connection list panel. The connection list displays the connection's name, the schema's name, and the OAuth type you grant. However, it is still in a disconnected state.

5. Click the three vertical dots beside the connection name and perform the following operations:

Connect: Select **Connect** to the Autonomous Database and change the connection status to active. Selecting **Connect** opens the sign-in page of the Autonomous database. After you log in, you will view a page that shows that database access has been granted to you. Close the window and return to Google Sheets. You will now see that the connection is active.

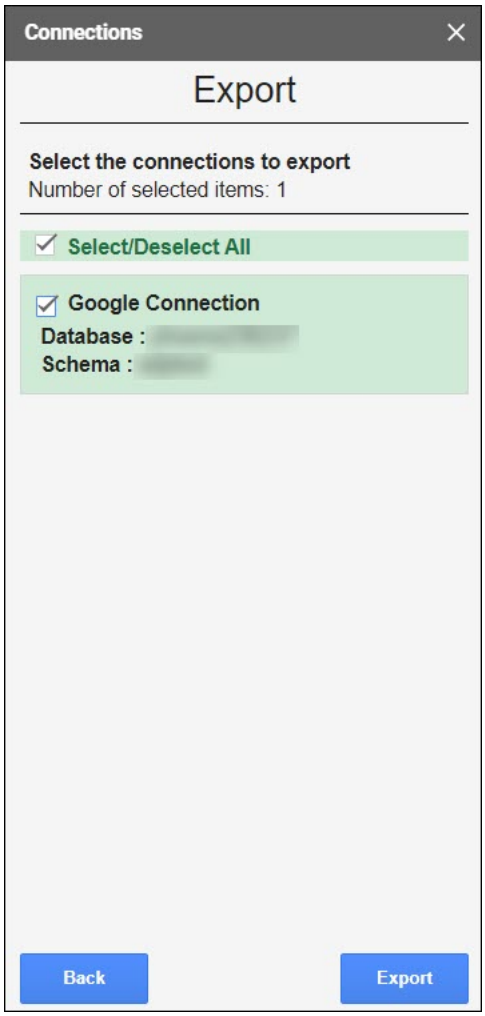
Edit: Select **Edit** to update any value of the connection. Click **Save** to update the edited values.

Duplicate: Select **Duplicate** to create a duplicate connection.

Remove: Select **Remove** to remove the connection from the connection list.

Exporting Connections

1. Click **Export Connection** from the **Manage Connections** drop-down menu to export the selected connection.
2. Select the connection you want to export, and click **Export**.



3. Click **Export**.
4. The exported connection downloads in your local system. The connection file is saved as *spreadsheet_addin_connections.json*.

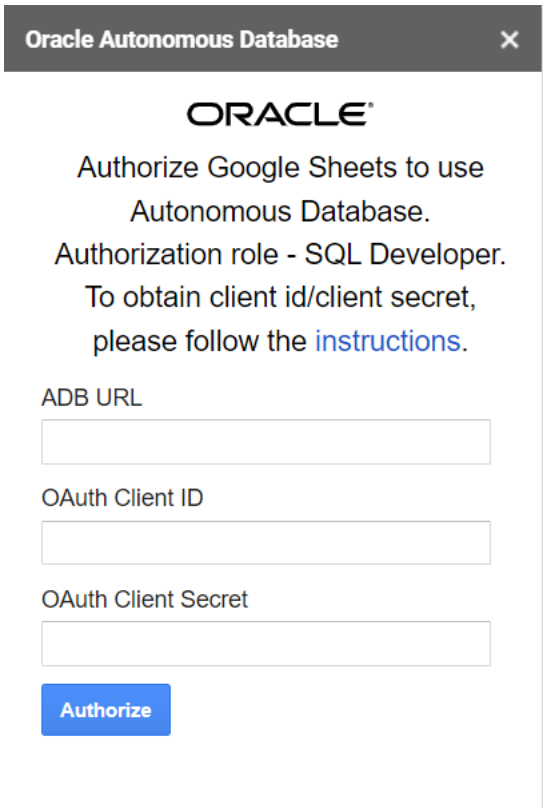
Authorize Google Sheets to use Autonomous Database

After your identity is determined using OAuth authentication, Google Sheets needs permission to access the Autonomous Database.

The `client_id` and `client_secret` values you generate during OAuth authentication are used for authorization.

1. Click on the **Oracle Autonomous Database** menu in the Google Sheet you are working on and select **Register**. This requires one-time Google authentication.

2. Clicking Register opens a pop-up window that asks your permission to run the authorization. Click **Continue**. Selecting Continue will redirect you to the Google Accounts page, where you must select your Gmail account.
3. You will now view a window that informs you that the application requests access to sensitive information in your Google account. Click Advanced and select the Go to Untitled project (unsafe) link.
4. Selecting the link opens a new window, ensuring you trust the application. Click **Allow** to continue.
5. You have now completed the setup. Select **Register** from the Oracle Autonomous Database menu in the Google sheet.
   This opens an Oracle Autonomous Database wizard in the Google sheet. Specify the following fields:
   - **ADB URL**: Enter the ADB URL. For example, "*https://<hostname>-<databasename>.adb.<region>.oraclecloudapps.com/ords/<Schema Name>*".
   - **OAuth Client ID**: client_id you generate during authentication.
   - **OAuth Client Secret**: client_secret you generate during authentication. Refer to the Create Connections with the Google spreadsheet section for more details.



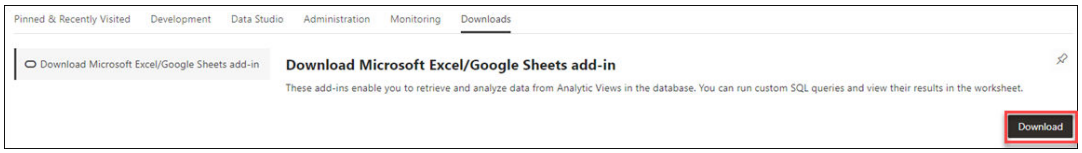
6. Select **Authorize**.

After successfully authorizing the credentials, you can view **Connections**, **Direct SQL**, **Data Analysis**, **Analyses and Reports Clear Sheet**, **Delete All Sheets**, **About Autonomous Database**, and **Sign Out** menu items under **Oracle Autonomous Database**.

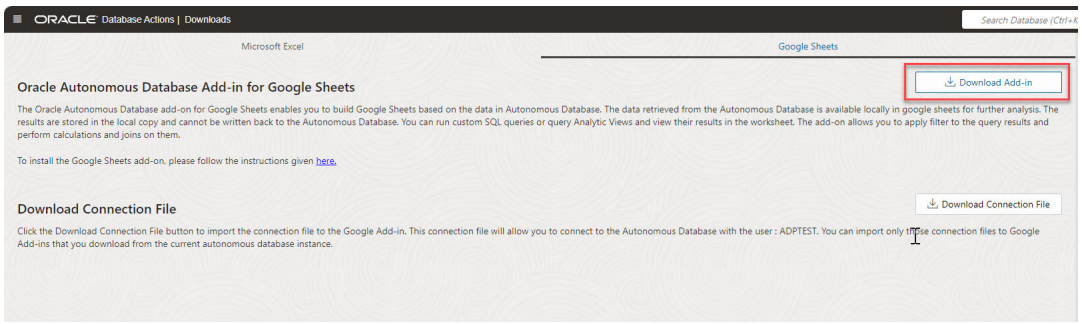
Install and setup the add-on for Google Sheets

Before you install the Oracle Autonomous Database add-on for Google Sheets, download the *oracleGoogleAddin* zip file from your Database Actions instance.

- Open the Database Actions Launchpad.
- Under the Downloads tab, select the **DOWNLOAD MICROSOFT EXCEL/GOOGLE SHEETS ADD-IN** pane.



- This opens a Download screen with Microsoft Excel and Google Sheets tabs. Click the Google Sheets tab and select **Download Add-in**.



You can now view a zip file in the your system's Downloads folder. Extract the contents of the zip file onto your system.

To set up the Oracle Autonomous Database add-on for Google Sheets, import the files in the *oracleGoogleAddin* folder to Google Apps Script.

> **Note:**
>
> Importing the files is a one-time activity, and typically, this is done by an administrator.

After you import or upload the files to Google Apps Script follow these steps to complete the setup of the Oracle Autonomous Database add-on for Google Sheets:

- Deploy the Google Script as a web app.
- Generate Client ID and Client Secret keys using UI.
- Upload oracleGoogleAddin files to Google Apps Script using Clasp
  To upload all the files present in the oracleGoogleAddin folder, you must use the Command Line Apps Script Project (clasp).
- Deploy the Google script as a web app
  After all the files from the oracleGoogleAddin folder are imported or uploaded to the Google Apps Script files, you must deploy the Google script as a web app.

Upload oracleGoogleAddin files to Google Apps Script using Clasp

To upload all the files present in the oracleGoogleAddin folder, you must use the Command Line Apps Script Project (clasp).

Clasp is an open-source tool to develop and manage the Google Apps Script projects from your terminal.

> **Note:**
>
> Clasp is written in Node.js. and distributed via the Node Package Manager (NPM) tool. It is required to install Node.js version 4.7.4 or later to use clasp.

1. Enter *sheet.new* in the web browser's address bar to open Google Sheets. Make sure you are logged in with your Google account.
2. Select **Apps Script** from the **Extensions** menu. You can view the Apps Script editor window.



3. Select the Code.gs file in the Apps Script editor window, which already exists by default. Click on the vertical dots beside the Code.gs file. Select **Delete** to delete the existing Code.gs file.
4. After you install Node.js, enter the following npm command in the command prompt to install clasp. You must enter this command in the location where you have downloaded and extracted the oracleGoogleAddin folder.

```
C:\Users\username\Desktop\oracleGoogleAddin>npm install @google/clasp -g
```

To run the command as an administrator for UNIX- and Linux-based systems, enter the following command:

```
sudo npm install @google/clasp -g
```

After you install Clasp, the command is available from any directory on your computer.

5. Enter the following command to log in and authorize managing your Google account's Apps Script projects.

```
clasp login
```

Once this command is run, the system launches the default browser and asks you to sign into your Google account where your Google Apps Script project will be stored. Select **Allow** for clasp to access your Google Account.

> **Note:**
>
> If you have not enabled the Apps Script API in Google Apps Script, the above command will not be successful. Enable the API by visiting the https://script.google.com/home/usersettings site and allow site and enabling the Google Apps Script API by selecting the **On** button.

6. In your existing Google Apps Script project, click the Project Settings in your left pane. Click **Copy to Clipboard** to copy the Script ID under IDs.
7. Go back to the command prompt and enter the following command with the Script ID you copied in the previous step as displayed in the image below:

   ```
   clasp clone <Script ID>
   ```

8. Push all the files from your folder to the Google Apps Script files by specifying the following command:

   ```
   clasp push
   ```

   This command uploads all of the script project's files from your computer to Google Apps Script files.
9. Go to the newly created Google Sheet, click the Extensions menu, and select **Apps Script**. Under Files, you can view all the files in the *oracleGoogleAddin* folder.
10. After you import or upload the files to Google Apps Script, follow these steps to complete the set up of the Oracle Autonomous Database add-on for Google Sheets:
    - Deploy the Google Script as a web app
    - Generate Client ID and Client Secret keys using UI

Deploy the Google script as a web app

After all the files from the oracleGoogleAddin folder are imported or uploaded to the Google Apps Script files, you must deploy the Google script as a web app.

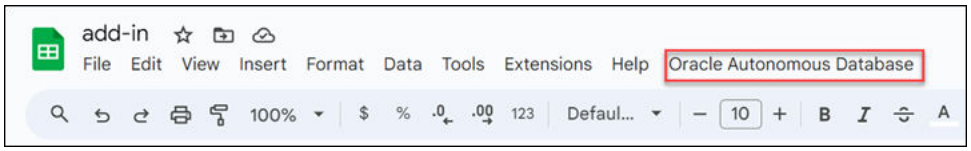
To deploy the Google script as a web app:

1. Click on the **Extensions** menu in the Google Sheet you are working on and select **Apps Script**. This opens the window.
2. Click **Deploy** button on the top right and select New deployment. A New deployment window opens.
3. Next to **Select type**, click the **settings** icon and select **Web app**.
4. Under Configuration, specify a Description of the deployment in the **Description** field. For example, *Web app deployment*.
5. Under **Web app** , select the Google account you used to log in from the **Execute as** drop-down. Optionally, you can choose anyone who has access to this deployment.
6. Select **Deploy**.

> **Note:**
>
> - If you receive a window that asks you to Authorize access, select it. This will redirect you to the Google Accounts page where you must to select your Gmail account.
> - Click **Advanced** and select the **Go to Untitled project (unsafe)** link.
> - Selecting the link opens a new window, ensuring that you trust the application. Click **Allow** to continue.
>
> If you Authorize access at this stage, you need not follow steps 2-4 in the *Authorize Google Sheets to use Autonomous Database*.

7. Click **Done** to close the New Deployment window.
8. Click the **Deploy** button on the top right and select **Manage Deployments**.
9. On the Manage Deployments page, you can view a Web app URL. Use the **Copy to Clipboard** to copy the Web app URL. For example, here is a sample of the web app URL "*https://script.google.com/macros/s/abcdefghijk/exec*". This is the **Web application deployment URL**.
10. Save this URL, which you will use later in the **Google Sheet Redirect URL** field when downloading a connection file from Database Actions or manually creating a connection from the Google Sheet to the Autonomous Database.

    For details on selecting **Response Type**, see Download Connection File.
11. You can close the Apps Script browser tab and navigate to the Google Sheets browser tab. You are now ready to create a connection to the Autonomous Database.
12. Ensure you save the worksheet after uploading all the files to Apps Script. Click the **Refresh** button once you have uploaded all the files. You can now view a new **Oracle Autonomous Database** menu in the Google Sheets.



> **Note:**
>
> Generate **OAuth Client ID** and **OAuth Client Secret** fields by using the UI.

Connecting to Autonomous Database

The Oracle Autonomous Database add-on for Google Sheets enables you to connect to multiple Autonomous Databases with a single add-on using the Connections feature. The add-on connects to Google Sheets by providing authentication to Google. Multiple users or databases can connect simultaneously to the add-on. However, only one connection can remain active.

The connection icon enables you to manage one or more database connections to databases in your Google Sheets. You can connect, edit, duplicate, and remove a connection.

You use the OAuth authentication and credentials to access the Oracle Autonomous Database for Google Sheets. The Add-on connects with the database using implicit and explicit types of connections.

Explicit Connection

You use the OAuth Client ID and OAuth Client Secret values to authenticate and authorize Google Sheets to use the Autonomous Database. Use this when you use **CODE** as the **Response Type** while downloading the connection file from Database Actions page.

Implicit Connection

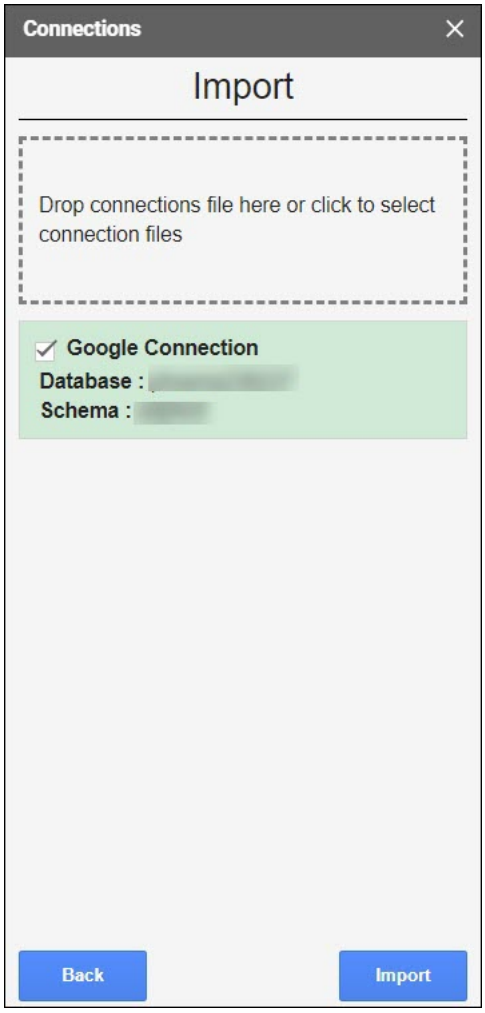
You will need an OAuth Client ID to implicitly access the Autonomous Database. Use this when you use **Token** as the **Response Type** while downloading the connection file from the Database Actions page.

Connection Management in Google Sheets

Import the connections file

Import an existing connection by clicking the **Import Connection** menu in the **Manage Connections** drop-down field.

1. Click the drop connection area and drag and drop the connection file saved in your local system to import the connection. You can import the **Connection File** you downloaded from the Download Connection File section.
2. Click and drop the connection file from your system to the drop area of the wizard. After the connection file loads, select the check box beside the connection file you want to import from the file.



3. Click **Import**. After you import the connection, you can view the connection in the list of connections.
4. Select the **three vertical dots** beside the connection name and click **Connect** to connect to the database.

Manual connection to the database

If you do not have a connections file to connect to the autonomous database and have access to the Schema via SQL Developer web, follow the steps outlined here to connect to the database.

Natural Language in Google Sheets

You can use Natural Language Query to query the Oracle Autonomous Database using the **Natural Language** menu in the **Oracle Autonomous Database for Google Sheets** "add-on".

Prerequisites

- Before you use the Natural Language menu in the add-on, you must perform prerequisites to use `DBMS_CLOUD_AI` to configure AI profiles.
- Create and Set an AI Profile

After your AI Profile is ready, the Data Studio tool uses these profiles in the Data Studio Settings menu to configure access to a Large Language Model (LLM) and to setup for generating, running, and explaining SQL based on natural language prompts. This also enables chatting with the LLM.

> **Note:**
>
> You can run Natural Language Queries on Tables and not Analytic Views.

Generate SQL Queries from Natural Language on Tables

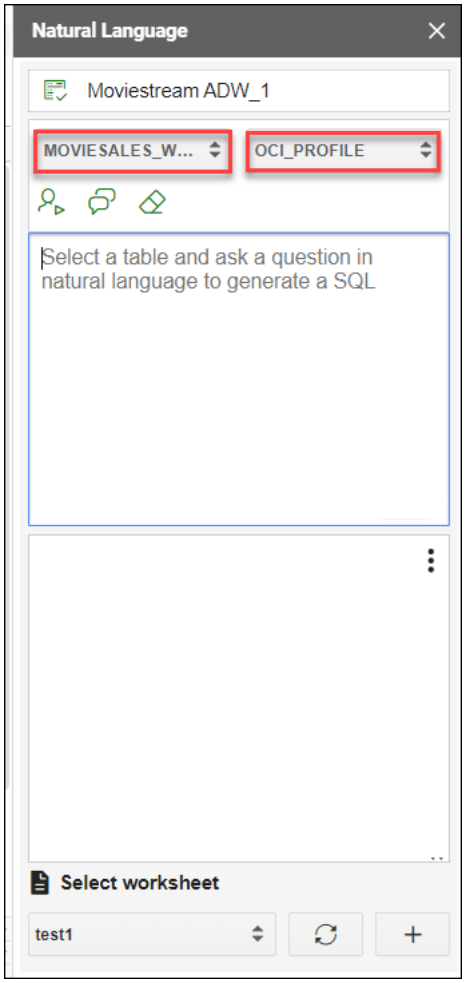
Using natural language to interact with your database data is now achievable with Oracle Autonomous Database add-on for Google Sheets.

This means you can use natural language, for example, plain English, to query the database. You can provide a natural language prompt instead of SQL code to interact with your data. When you select the **Generate SQL from natural query** icon, the add-on converts natural language to SQL.

To run a natural query using the add-on:

1. On the Google Sheet, select the menu item **Oracle Autonomous Database**.
2. Select **Natural Language**. Selecting **Natural Language** opens a **Natural Language** wizard.
3. On **Natural Language** wizard, select the **Table** on which you perform the query from the drop-down and the **AI profile** you want to use from the drop-down.

> **Note:**
>
> This is the same AI profile you configure in the Data Studio Settings.



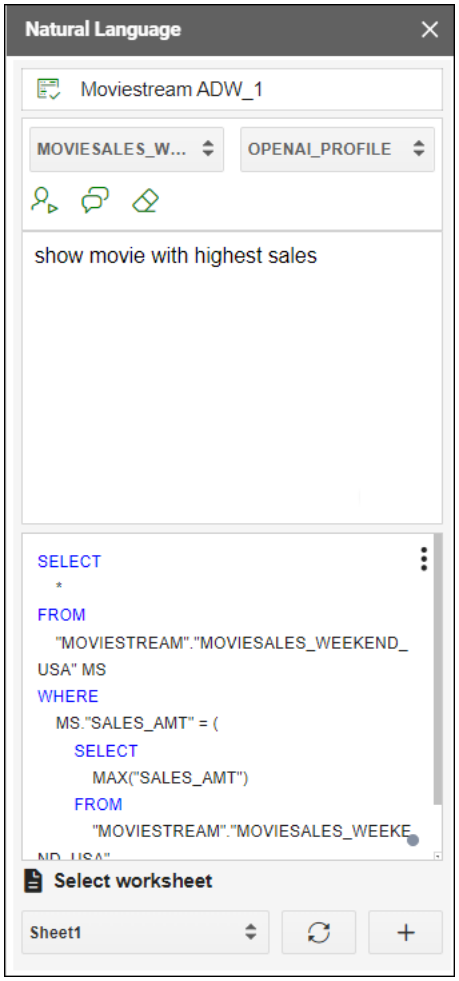
In this example, `MOVIESALES_WEEKEND_USA` is the name of the **Table** and `OCI_PROFILE` is the **AI Profile**.

4. Let's say you want details of the product with highest sales. Enter the following natural language query in the query display area:

```
show movie with highest sales
```

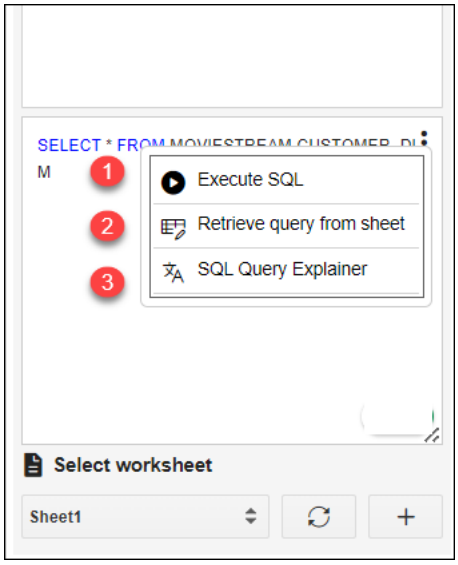
5. Select **Generate SQL from natural query** to produce the equivalent SQL query in the bottom query display area.
You will view the following code in the bottom SQL code area.

```
SELECT
    *
FROM
    "MOVIESTREAM"."MOVIESALES_WEEKEND_USA" MS
WHERE
    MS."SALES_AMT" = (
        SELECT
            MAX("SALES_AMT")
        FROM
            "MOVIESTREAM"."MOVIESALES_WEEKEND_USA"
    )
```



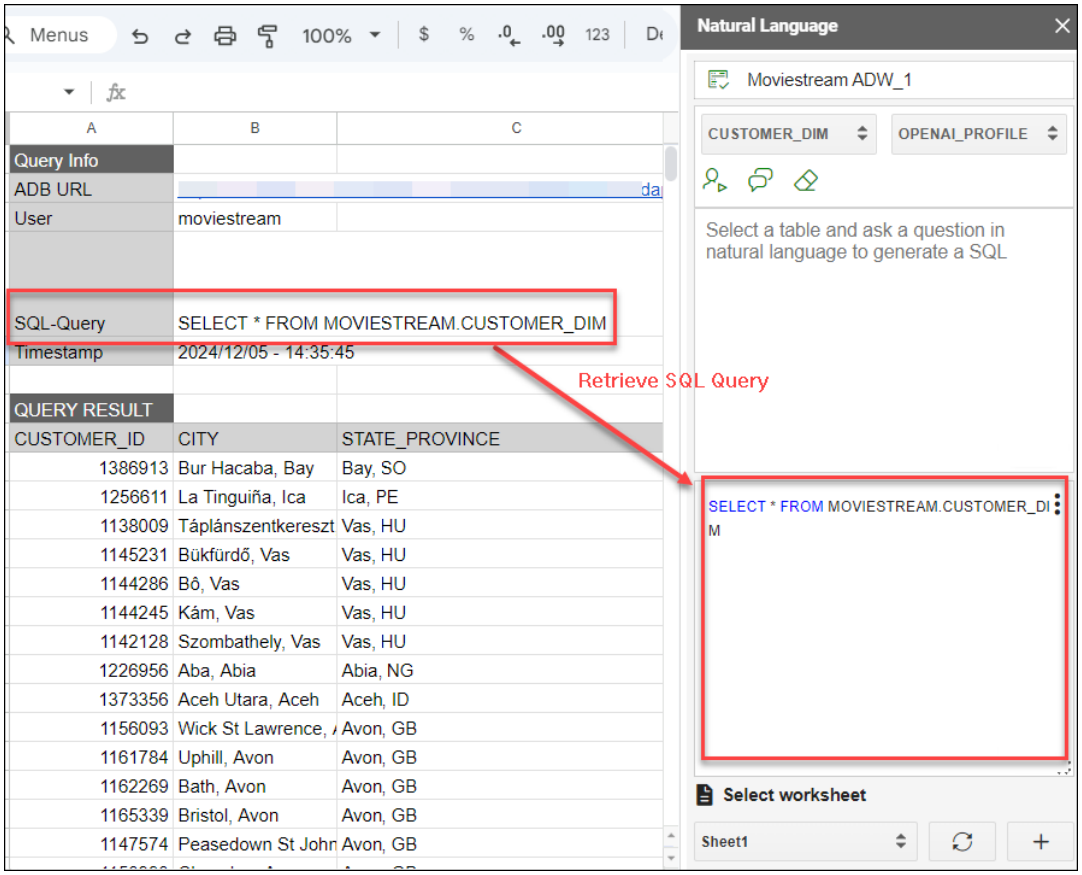
6. Selecting the **Generate SQL from natural query** icon also displays the result of the natural language query in the Google Sheet.
   You can click the **+** sign beside the **Select worksheet** drop-down to display the results in a new worksheet.
7. Click **Clear query content** to clear the content displayed on the natural language query area and the SQL code query area.

Actions you can perform from the SQL code area:



- **Execute SQL**: Runs SQL code and displays the query result in Google Sheet.
- **Retrieve query from Sheet**:
  You can retrieve SQL query displayed in the Google Sheet to the SQL code area in the Natural Query.

  This feature is to retrieve the latest data from the database for a future date. For example, if you want to receive the latest data from the table the next day, the connection may be inactive, but you can still connect to the instance, **retrieve query from Sheet** and **select Execute SQL** to display the latest query results in the Google Sheet.



- **SQL Query Explainer**: Explains SQL code in Natural language form.
  Select this icon to translate SQL queries to natural language that is understood by you. The natural language query is displayed in the natural language query display area of the wizard.

Ask Questions with Chats

Chats option in the Oracle Autonomous Database for Google Sheets refer to an interactive conversation between you and the add-on where the add-on uses natural language to query or interact with the autonomous database.

The data we use in this example is of a company called `Oracle MovieStream` to analyze movie sales data.

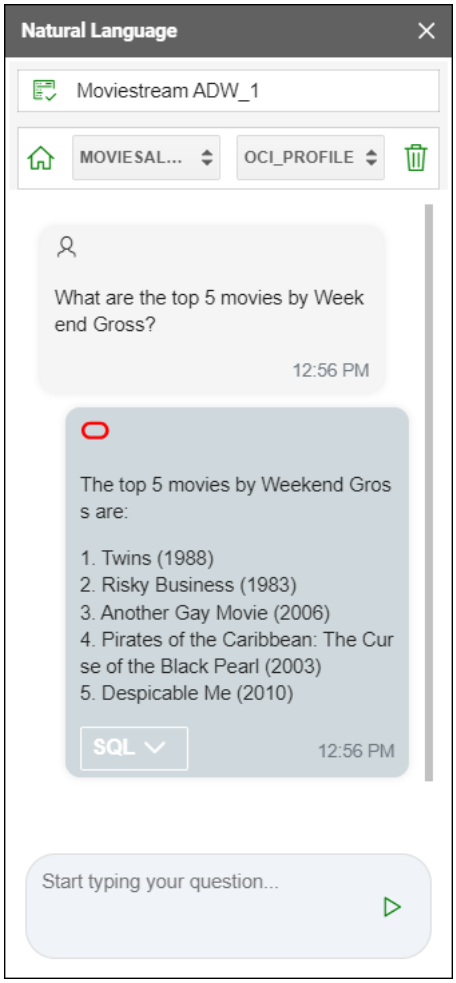
The Chats displays recommendations for the default table you select.

Let's find out the top 5 streamed movies in the Moviestream company by weekend gross:

1. Enter the text in the **Start typing your question..** text field.

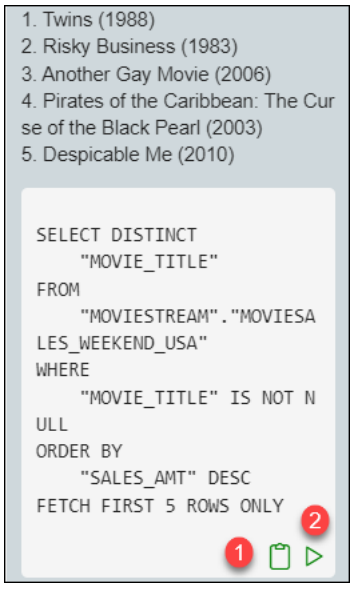
   ```
   What are the top 5 movies by Weekend Gross?
   ```

2. Click **Enter**.
3. The Chat displays top 5 movies by Weekend Gross.



It also displays equivalent SQL code of the result.

Click **SQL** to expand the SQL.



Click **Copy to clipboard** to copy the SQL.

Click **Run SQL** to run the generated SQL query and display the query result in Google Sheet.

The **Chats** option remembers the context of previous chat history.

For example, if you enter `now show me top 10` in the text field.

It displays the top 10 movies by Weekend Gross. It remembers which metric top 10 to fetch without us having to type the whole thing.

Select **New Chat** for deleting the present conversation.

Select **Home** to go back to the main home page.

> **Note:**
>
> LLMs are remarkable at inferring intent from the human language and they are getting better all the time; however, they are not perfect! It is very important to verify the results.

Data Analysis in Google Sheets

Selecting **Data Analysis** opens an Oracle Autonomous Database wizard in the Google sheet.

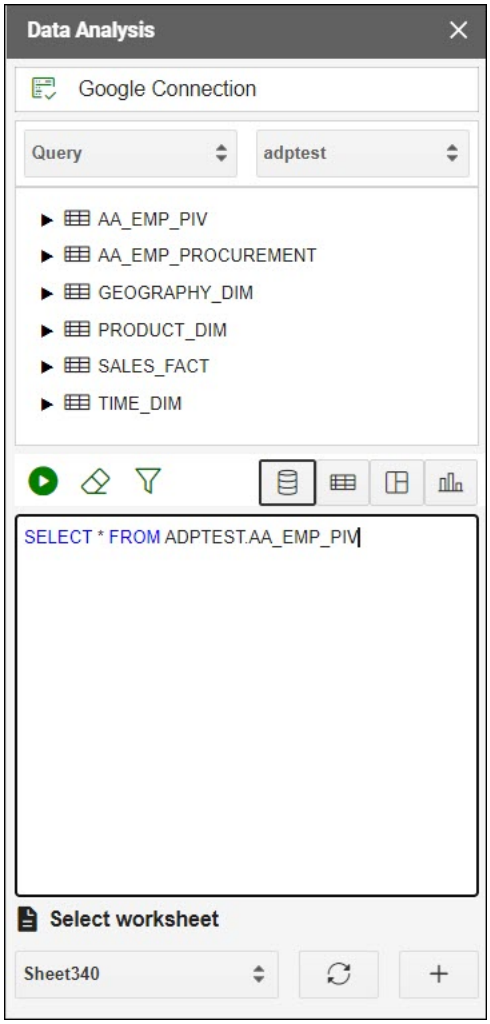
The add-on enables you to receive a copy of data from the Autonomous Database to the Google sheet. You can query an existing Analytic View and run SQL Query using the Oracle Autonomous Database Wizard.

You can retrieve the Analytic View and manipulate the query according to your requirements to visualize the result data in the worksheet. You can search for the Analytic View and select measures, hierarchies, and levels from the query. You can also add filters and calculated measures to the query and view the result in the sheet. By default, the data is retrieved in tabular format.

You can run custom queries. The add-on enables you to apply a filter to the query results. The add-on lets you to view query results that can be customized with selected columns using a faceted filter.

To run a custom query using the add-on:

1. On the Google Sheet, select the menu item **Oracle Autonomous Database**.
2. Select **Data Analysis**. Selecting Data Analysis opens a Data Analysis wizard. On the Data Analysis wizard, select **Query** from the drop-down and the **schema** you want to use from the drop-down.



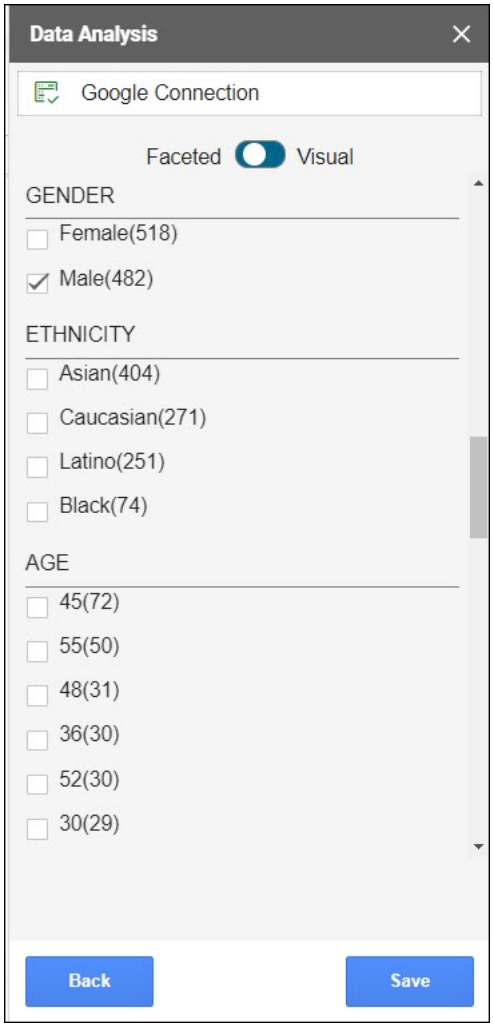
3. You will view the default query in the query editor area. You can select any of the four modes to visualize the results of the SQL query report you generate:



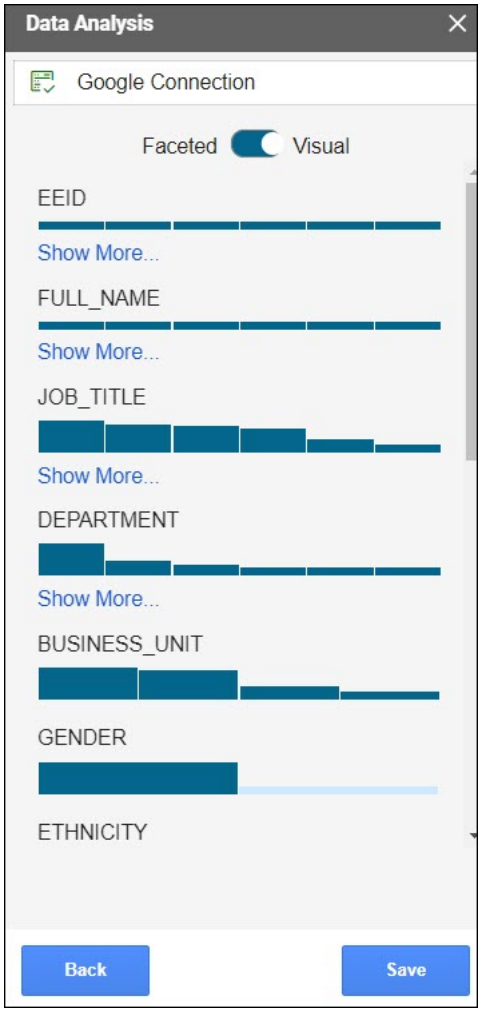
- **Base Query**: This type of view is by default. The query written in the SQL editor is the Base Query.
- **Table**: You can view the SQL results in tabular form. By selecting this view, a column drop zone appears, enabling you to drag and drop selected columns from the Table browser. Moving the selected columns in the drop zone allows you to view only those columns in the Result data generated in the worksheet. Select the cross mark beside the Column name to remove it from the drop zone.
- **Pivot**: You can view the SQL query results in pivot format. By selecting this format, an X and Y drop zone appears where you can drag and drop the selected columns from the Tables browser to the drop area.
- **Chart**: You can view *Area Chart*, *Bar Chart*, *Line Chart*, or *Pie Chart* when you select this option.. The mappings displayed when you select one of the options are as follows:

- **Orientation**: Choose between horizontal and vertical orientation types from the drop-down list.
- **X axis label and Y axis label**: Optionally enter labels for X axis and Y axis.

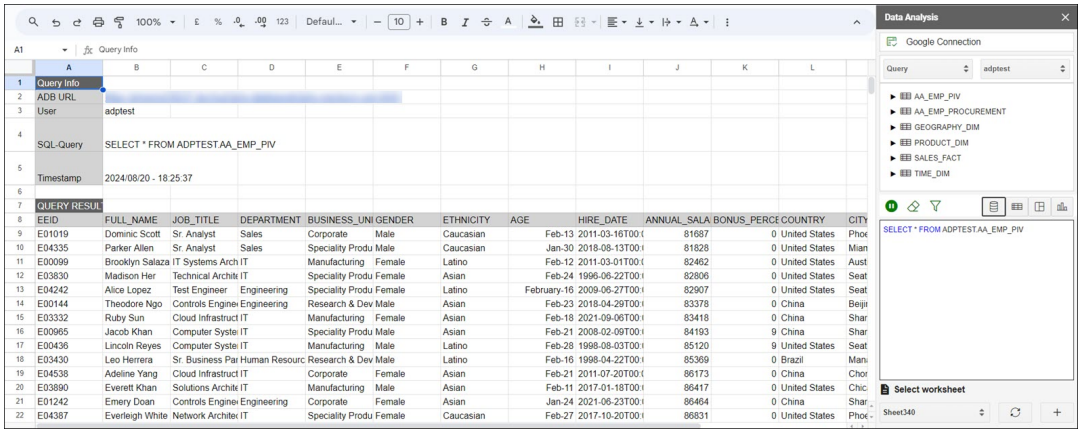
4. Click the **funnel icon** (Faceted Filter) to add filters to the result. The wizard generates a filter for each value in the column retrieved from the query result. You can filter different columns on the faceted filter panel and view the results in the worksheet to view only the data you wish to view. For example, to view the customer reports by Gender, click the **faceted filter** and select **Male** under Gender.



Use the Visual facet: Use the visual indicator to graphically represent the faceted filter.



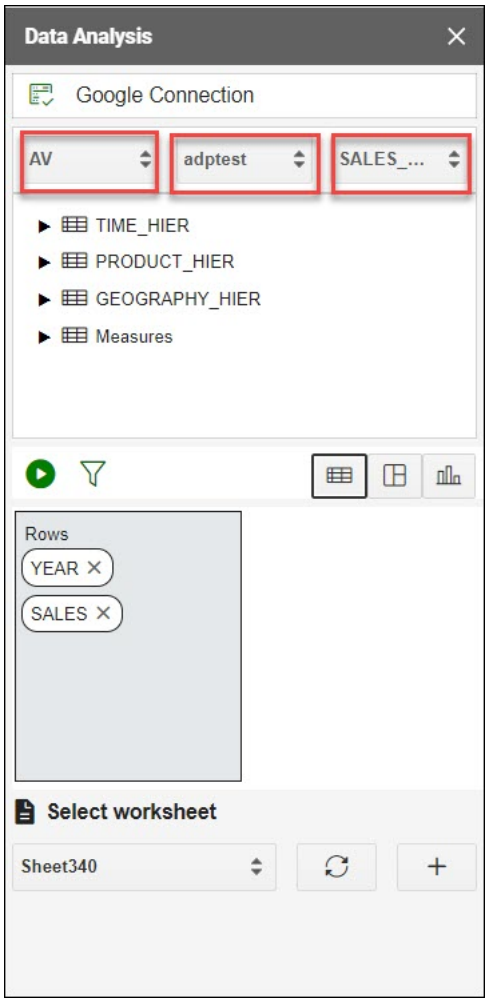
5. Select **Run** to generate the results of the custom query in the worksheet.



To query an Analytic View and explore the **Data Analysis** menu in the Google Sheets:

1. Select the menu item **Oracle Autonomous Database > Data Analysis** on the Google Sheet. This opens a **Data Analysis** wizard in the Google task pane.

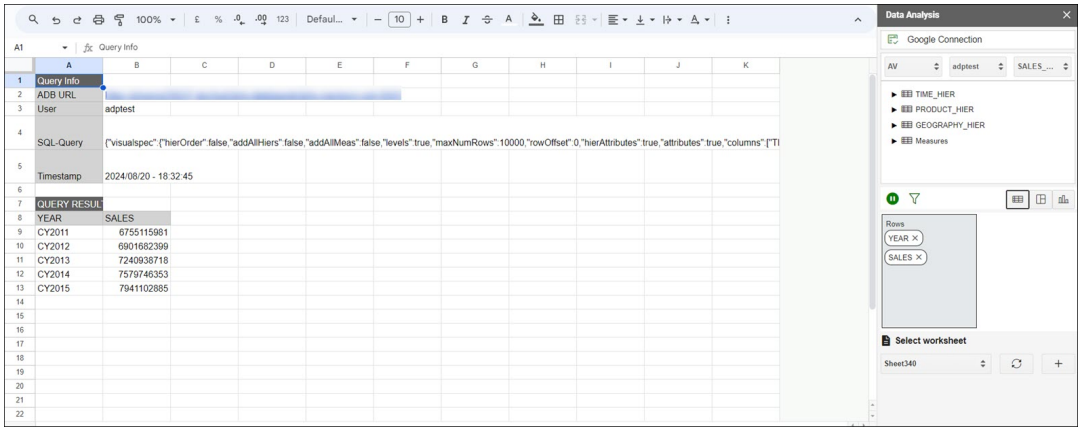
2. Select **AV** from the **AV** or **Query** drop-down, select a schema you can access from the schema drop-down, and the AV from the available Analytic Views.



3. You can select any of the three modes to visualize the results of the AV query you generate:
   - **Table**: You can view the AV query results in tabular form. By selecting this view, a column drop zone appears, enabling you to drag and drop selected columns from the Table browser. Moving the selected columns in the drop zone allows you to view only those columns in the Result data generated in the worksheet. Select the cross mark beside the Column name to remove it from the drop zone.
   - **Pivot**: You can view the SQL query results in pivot format. By selecting this format, an X and Y drop zone appears where you can drag and drop the selected columns from the Tables browser to the drop area.
   - **Chart**: You can view the results of the AV query in chart format. By selecting this format, an X and Y drop zone appears where you can drag and drop the chosen hierarchies and measures from the AV browser to the drop area.

> **Note:**
>
> You are allowed to drop measures in the Y-axis.

4. Click the **funnel** icon to view the Faceted and Visual Filter list. The wizard generates a filter for each value in the column retrieved from the query result. You can filter different columns on the faceted filter panel and view the results in the worksheet to view only the data you wish to view.
5. Click **Save**.
6. Click **Run** to view the results in the worksheet you select.



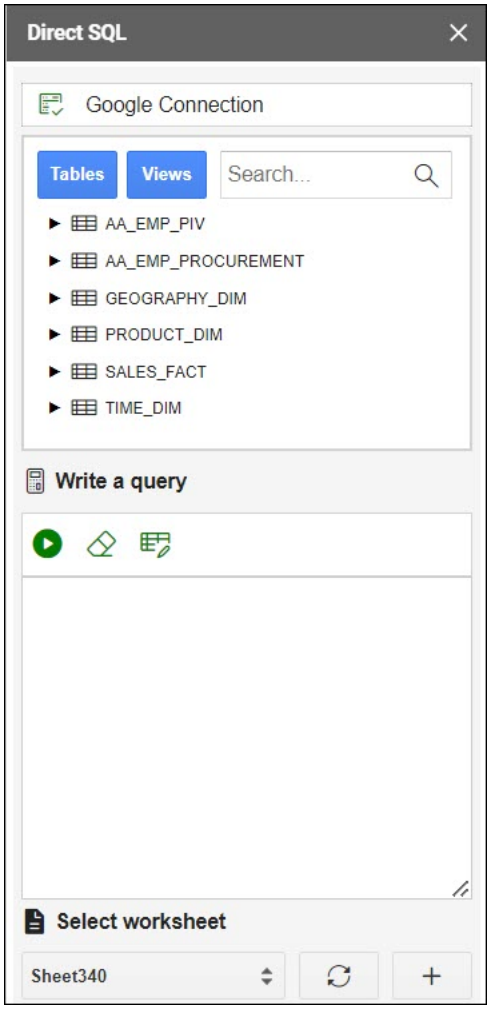
Run Direct SQL Queries

The Oracle Autonomous Database add-on for Google Sheets lets you run SQL queries to work with your data in a Google Sheet. With the add-on, you can type your SQL code in the SQL editor area and click **Run** to run the command.

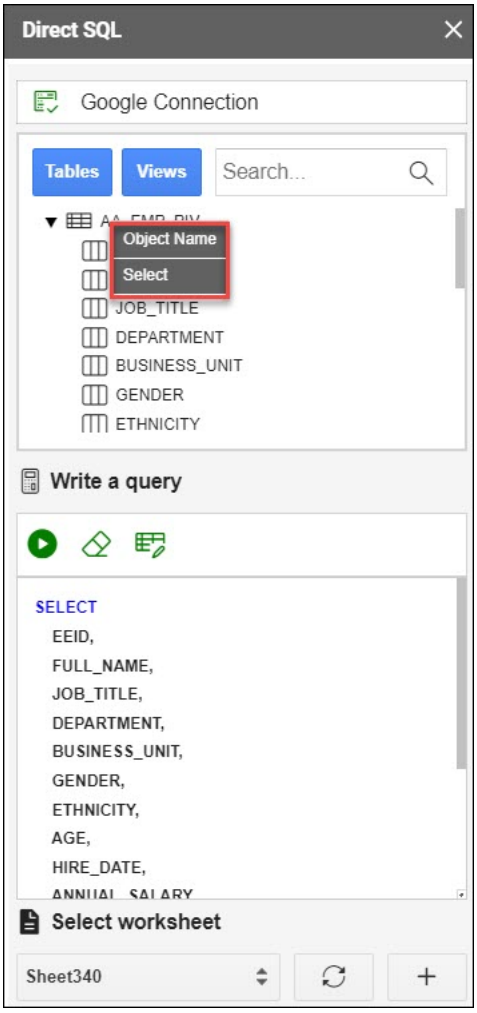
The add-on loads the result in the Google Sheet. The time taken to load the results depends on the number of records and the complexity of the query.

To run a query using the add-on, open Google Sheets and a blank workbook.

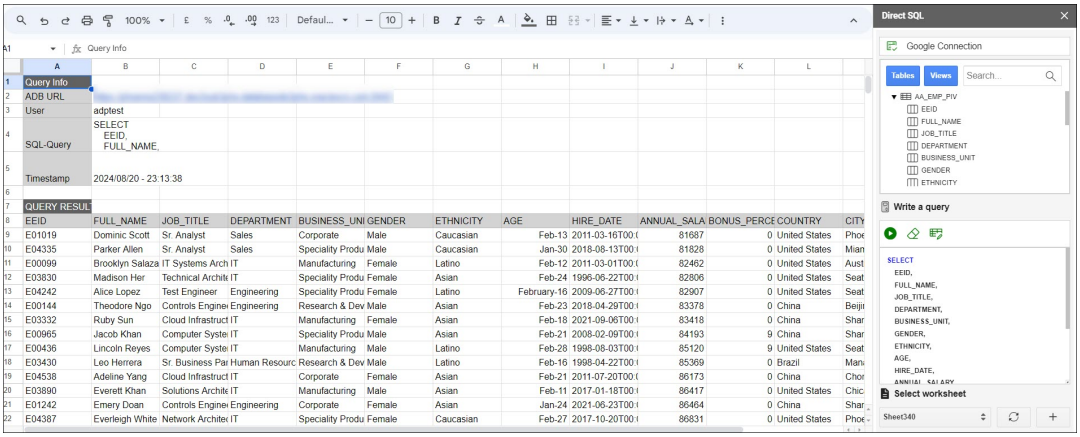
1. In the Google Sheet, select the menu item **Oracle Autonomous Database**.
2. Select **Direct SQL** to type and run the SQL command.
3. The Oracle Autonomous Database wizard opens Tables and Views icons and a search field beside it.



4. Select **Table** to view all the tables in the database. Perform the same operations for Views.
5. You can right-click on the table whose data you want to query and choose **Select** to view all the table's columns. The column names will be displayed in the **Write a Query** section. You can click on the table and view individual columns as well.



6. Click **Run** to run the query and display the results in the worksheet. You can click the **+** sign beside the **Select worksheet** drop-down to display the results in a new worksheet.
7. The worksheet also displays the timestamp, the user who creates and runs the query, the ADB URL, and the SQL Query.



Reporting and Analysis in Google Sheets

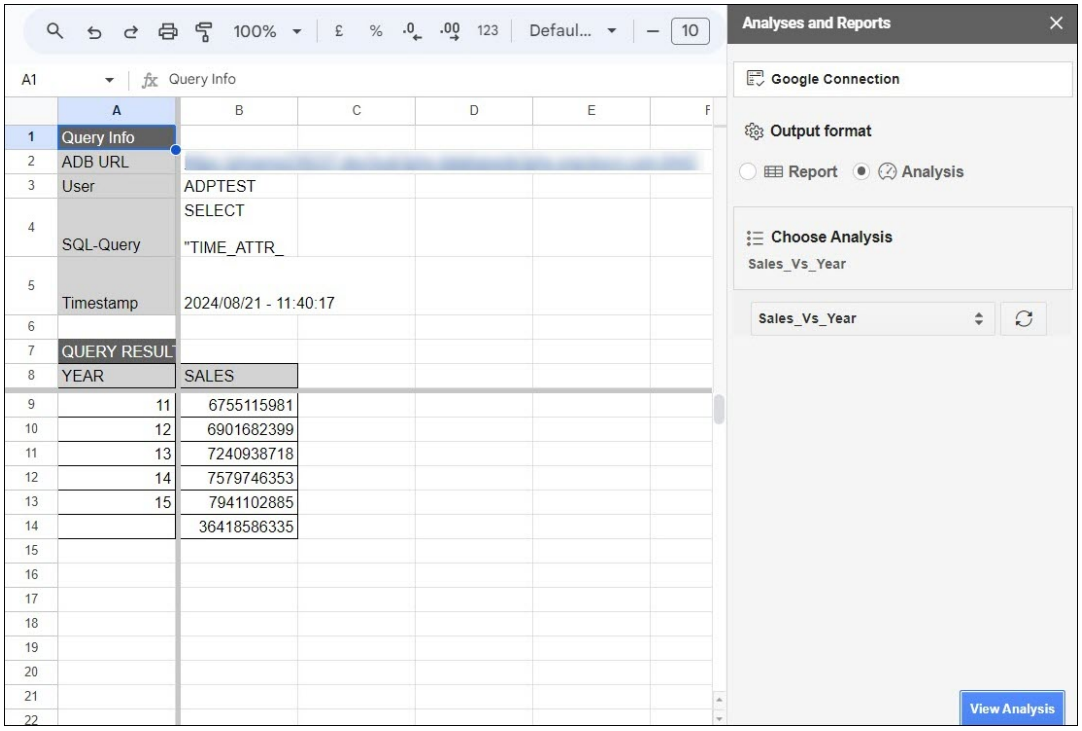
You can view Reports and Analytic Views or visualize data for analysis purposes.

The reports and charts can be viewed in various charts: Bar Charts, Area charts, Line Charts, and Pie Charts. Reports provide analytical insights that you create from the Analytic Views. An Analysis can contain multiple reports. The Analyses and Reports icon enables you to retrieve Analyses and Reports from the Autonomous Database.

View Analysis

To view Analysis and explore the Analyses and Reports menu:

1. Select **Analysis** under Output format.
2. Use the Select an Analysis drop-down to choose the Analysis you want to view.
3. Click **View Analysis** to view the analysis in the Google Sheet.

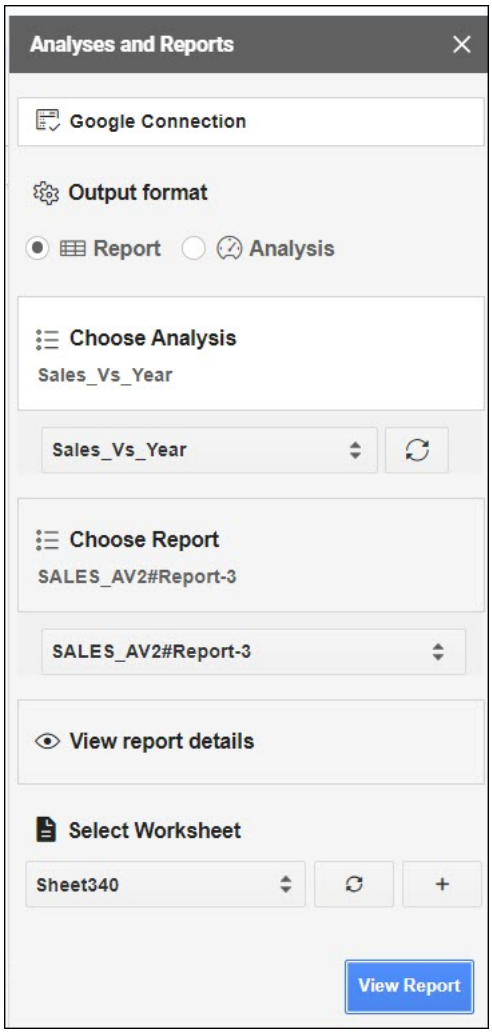


View Report

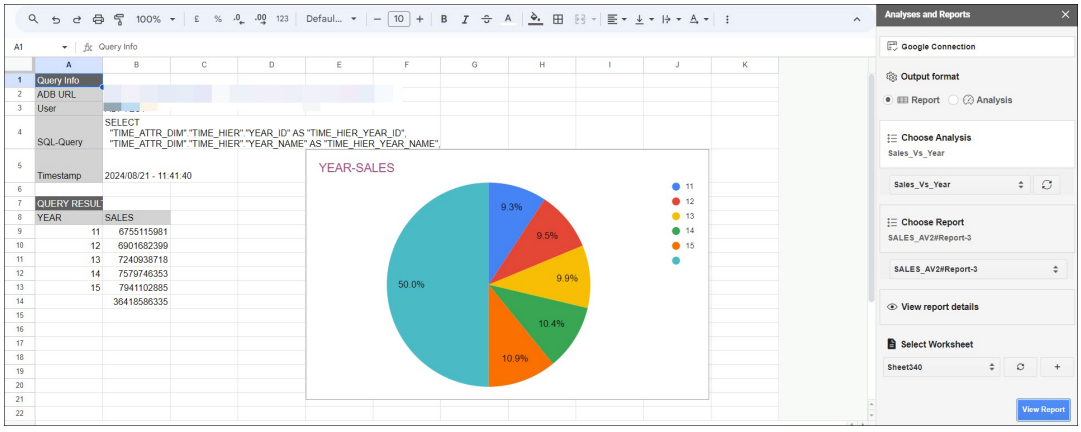
To view Reports :

1. Select the **Analyses and Reports** menu from the **Oracle Autonomous Database** menu. This opens the Analyses and Reports wizard.
2. Select **Report** under Output Format.
3. Use the **Select an Analysis** drop-down under **Choose Analysis** to choose the Analysis you want to view.
4. After you select the Analysis, to view the report present in the Analysis, click the **Select a report** drop-down and select the report you wish to view.

5. Click **View Report Detail** to view more information about the report: Analytic View Name, Type of visualization, and rows, columns, and values you select while creating the report.



6. Select the worksheet from the drop-down where you would want to view the report.
7. Click **View Report** to view the report in the selected sheet from the previous step. You can now view the report in the worksheet you select.



Clear Sheet

Once the add-on runs the query and retrieves the data into the worksheet, you can view the Timestamp, User, AV-query and SQL-query of the Analytic View in the automatically generated query results.

Once the add-on runs the query and retrieves the data into the worksheet, you can view the Timestamp, User, AV query, and SQL query of the Analytic View in the automatically generated query results.

The worksheet displays the result of the query in one go. Consider, for example, if you want to modify the query and generate the query result in the same sheet. You must clear the existing data in the sheet.

To clear query results in the Google sheet, click the menu item **Oracle Autonomous Database** and select **Clear Sheet**.

This option erases all data types in the selected sheet, including images and formatting.

Delete all sheets

Use this option to delete all the sheets existing in the spreadsheet.

Select **Delete All Sheets** from the **Oracle Autonomous Database** menu to delete all sheets from the spreadsheet.

About Oracle Autonomous Database menu

Use this option to view details about the add-in

The **About Oracle Autonomous Database** menu from **Oracle Autonomous Database** displays if the add-on is connected to server, the ORDS version, the Add-in version, the ORDS Schema version, the database major and minor version, the ADB URL and the Schema.

Share or Publish

Once you generate the query results in the Google Sheet, you can share it with other users. With sharing, creates a copy of the worksheet and sends it with the design tools hidden and worksheet protection turned on.

The recommended steps to take before you publish are:

1. Review and inspect to remove personal or sensitive information.
2. Save the source version of the worksheet. Consider adding a file name suffix of –src for the source worksheet. Then, remove the suffix in the distributed copy.
   Once you are ready to distribute to the users, click **Share**.
   - In the Share window that appears, add the user email IDs with whom you want to share the Sheets and to whom you want to provide permissions for accessing the Sheets.
   - You can select the permission of the users from the drop-down. Select **Editors** if you want the user to share the worksheet. **Viewers** and **commenters** can see the option to download, print and copy but not share the sheets.
   - Select **Notify people** check-box to notify the users of the share.

- Under General access, select **Restricted** from the drop-down to share it with people who have access to the link. You could also share it with people who do not have access by selecting **Anyone with the link** from the drop-down.

Oracle Autonomous Database for Google Sheets Privacy Policy Details

Effective Date: `10/4/2024`

Oracle Autonomous Database add-on for Google Sheets facilitates seamless data integration between Google Sheets and Oracle Autonomous Database. This Privacy Policy topic explains how we access, use, share, protect, retain, and delete your Google user data when you use our add-on.

Data use by **Oracle Autonomous Database add-on for Google Sheets** complies with Oracle's Privacy Policy.

1. **Google User Data We Access**
   - **Google Sheets Data**: The add-on accesses the content of your Google Sheets to import data from your Oracle Autonomous Database.
   - **User Profile Data**: User Profile Data, which includes personal information linked to a Google account, is not accessed or used by our add-on, as authentication is done with Oracle credentials. The data our application handles is limited to Google Sheets content.
2. **How We Use Google User Data**
   - **Data Integration**: Your Google Sheets data is used to enable functionalities like importing data from Oracle Autonomous Database into Google Sheets.
   - **Authentication**: Authentication is performed using your Oracle credentials. We do not access or use any Google authentication data or Basic Profile Information.
3. **Data Sharing and Disclosure**
   - **No Third-Party Sharing**: We do not share, sell, or disclose your Google user data to any third parties.
   - **Legal Obligations**: We may disclose your information if required by law or in response to valid legal processes.
4. **Compliance with Google API Services User Data Policy**
   - **Adherence to Limited Use Requirements**: To the extent the **Oracle Autonomous Database add-on for Google Sheets** uses and/or transfers information received solely from Google APIs to any other app, such use and transfer will adhere to Google API Services User Data Policy, including the Limited Use requirements.
5. **Data Protection Measures**
   - **Secure Transmission**: All data transferred between Google Sheets and Oracle Autonomous Database is encrypted using industry-standard protocols as detailed in our Security and Authentication for Autonomous Database.
   - **Access Controls**: Access to your data within the Add-on is restricted to authorized processes, ensuring that only necessary operations are performed.
   - **Oracle's Security Standards**: We adhere to Oracle's comprehensive security practices, including regular security assessments and compliance with global security standards.
6. **Data Retention and Deletion**

- **No Data Storage**: We do not store your Google Sheets data on our servers. Data is accessed in real-time solely for the purpose of facilitating the integration between Google Sheets and Oracle Autonomous Database.
- **User Control**: You can revoke the Add-on's access to your Google Sheets data at any time through your Google account settings.
- **Deletion Requests**: Since we do not retain your data, deletion requests are generally unnecessary. However, if you have concerns, please contact us at *adb_googlesheets_dev_ww_grp@oracle.com*.

7. **User Choices and Controls**
   - **Manage Permissions**: You can manage or revoke the add-on's permissions through your Google account settings at any time.
   - **Uninstallation**: To discontinue using the add-on, you can uninstall it from Google Sheets and revoke its access permissions.
8. **Changes to This Privacy Policy**
   We may update this Privacy Policy periodically. Changes will be reflected by the `Effective Date` at the top of this document. We encourage you to review this policy regularly to stay informed about how we protect your information.
9. **Contact Information**

   If you have any questions or concerns regarding this Privacy Policy, please contact us at:
   - **Email**: *adb_googlesheets_dev_ww_grp@oracle.com*
   - **Address**:
     Oracle Corporation

     500 Oracle Parkway

     Redwood Shores, CA 94065

     USA

Oracle Autonomous Database for Google Sheets Support

Welcome to the support page for the **Oracle Autonomous Database for Google Sheets**. This resource is designed to assist you with any issues or questions you may have while using the add-on.

Support for the Add-on

If you encounter any problems specifically related to the Oracle Autonomous Database Add-on, please reach out to our development team at `adb_googlesheets_dev_ww_grp@oracle.com`

Our team is dedicated to providing prompt assistance to ensure you have a seamless experience.

Support for Oracle Database Issues

For issues related to Oracle Database that are not specific to the Google Sheets add-on:

- **Contact Your Oracle Administrator**: Your organization's Oracle administrator can help with database access, permissions, and other database-related concerns.
- **Oracle Support Portal**: File a support ticket for further assistance at My Oracle Support.

The Catalog Tool

The Catalog Tool allows you browse, search and discover data in, and connected to, the Oracle Autonomous Database.

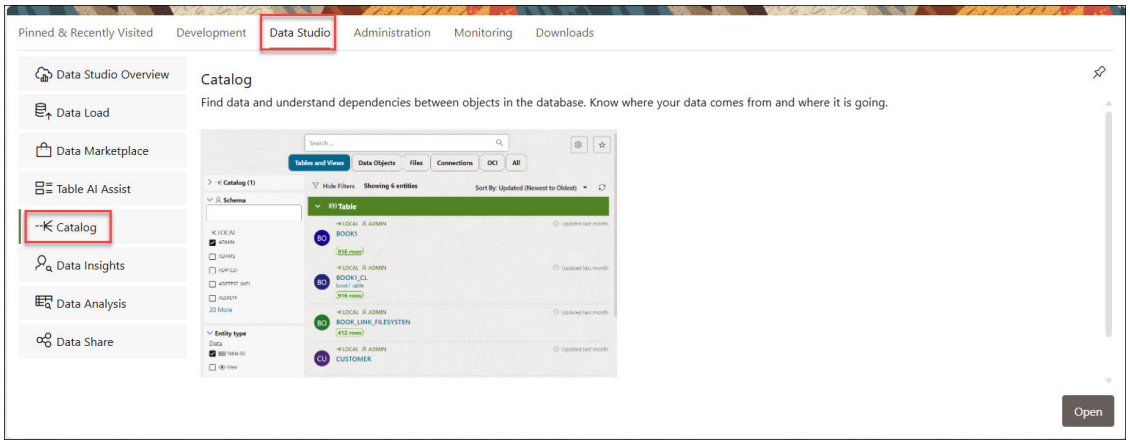
The Autonomous Database Catalog is a multi-catalog tool that provides a way to search for data and other objects in your currently connected Autonomous Database, and also in a wide range of other connected systems.

The Catalog tool enables you to search, find, load, or link connected data assets from anywhere in the cloud and beyond. It also enables you to mount new catalogs over other systems and search for data and other items in the connected Autonomous Database.

The Catalog tool can be used to mount new catalogs over:

- Other Autonomous Databases in your tenancy.
- Any other database that can be connected using a DB Link, for example an on-premises database that your Autonomous Database can connect to.
- Shared data, for example data shared from DataBricks using Delta Sharing.
- Existing external data catalogs such as AWS Glue, or the OCI Data Catalog.

To reach the Catalog tool, select the **Catalog** menu in the **Data Studio** tab of the Launchpad.



or click the Selector icon and select **Catalog** from the Data Studio menu in the navigation pane.

> **Note:**
>
> If you do not see the Catalog card then your database user is missing the required DWROLE role.
>
> Run the following command to grant `DWROLE` to the user:
>
> ```
> GRANT DWROLE TO ANALYST;
> ```

This chapter will take you through the features, capabilities, and technical architecture of the Catalog tool.

Enable the multi catalog feature in the tool

To enable the multi catalog feature in this tool, you need to have the `ORACLE_CATALOGS` role.

> **Note:**
>
> Until the user has `ORACLE_CATALOGS` role granted, they are limited to mounting only the local, connected Autonomous Database catalog.

In case the `ORACLE_CATALOGS` role is missing, run the following command as an **admin** to grant the `ORACLE_CATALOGS` role to the `user`:

```
CREATE ROLE ORACLE_CATALOGS;
    GRANT ORACLE_CATALOGS TO USER;
```

The following topics describe the Catalog tool and how to use it.

- About the Catalog Page
- Browse and Search Catalogs
- Sorting and Filtering Entities
- Filter Entities
- Display Entity Search
- About the Catalog Page
  Use the Catalog Tool to browse, search and discover data either in your Oracle Autonomous Database, or connected to it and get information about the entities in and available to your Oracle Autonomous Database. You can see the data in an entity, the sources of that data, the objects that are derived from the entity, and the impact on derived objects from changes in the sources.
- Browse and Search Catalogs
  The Catalog Tool aims to allow you to find data quickly and easily, whether or not the data is currently resident in the connected Autonomous Database.
- Sorting and Filtering Entities
  On the Catalog page, you can sort and filter the displayed search results.
- Filter Entities
  By selecting different quick filters, the faceted results of the entity types also vary.
- Display Entity Search
  If you do not select a quick filter, the **Tables and Views** filter is selected by default.
- Catalog Settings
  The Catalog Settings can manage the general settings, add, and view query scopes, create and view saved searches and change the look and feel of the Catalog page.
- Manage Catalogs
  You can now enable multiple catalogs using the Catalog Tool.

- Viewing Entity Details
  When you have found an entity you are interested in, you can click the entity to view its details, or click one of the other available actions from the menu that appears on the right hand side when you hover over that row.
- Registering Cloud Links to Access Data
  Cloud Links enable you to remotely access read only data on an Autonomous Database instance.
- Export Data to Cloud
  Use the **Export Data to Cloud** menu in the table actions of the Browse Catalog page to export data as text from an Autonomous Database to a cloud Object Store. The text format export options are CSV, JSON, Parquet, or XML.
- Gathering Statistics for Tables
  You can generate statistics that measure the data distribution and storage characteristics of tables.
- Editing Tables
  You can create and edit objects using Edit Table wizard available from the **Edit** menu in **Actions** (three vertical dots) besides the table entity.

About the Catalog Page

Use the Catalog Tool to browse, search and discover data either in your Oracle Autonomous Database, or connected to it and get information about the entities in and available to your Oracle Autonomous Database. You can see the data in an entity, the sources of that data, the objects that are derived from the entity, and the impact on derived objects from changes in the sources.

The Catalog page lists and displays details about:

- Database objects such as tables, columns, views, cloud objects, analytic views, packages, procedures, shares, share recipients, catalogs, catalog storage links, live table feed, share providers.
- Other types of object connected to the database, such as Cloud Storage Links, Cloud Storage Buckets, Cloud Objects (files on connected Cloud Storage), Compartments, Data Catalog Links, Catalog Storage Links, Autonomous Database.
- Database data dictionary objects such as tables, columns, views, database links, analytic views, packages, and procedures that have been created by the database tools or a database application such as SQL Developer.

The Catalog tool now provides a unified search and discovery of the following data assets that are connected to the Autonomous Database:

- Data in connected databases – for example, other Autonomous Databases in the same tenancy.
- Data in connected cloud storage systems – for example, in OCI, AWS S3, Azure, or Google Cloud Platform.
- Data in registered data catalogs – for example, in OCI Data Catalog, or AWS Glue.
- Data in registered data shares – for example, shared from Data Bricks.

Catalog Tool Terminology

The following are some major terms related to the Catalog Tool:

- **Catalog**: A catalog consists of one or many schemas, each with one or many tables. A catalog simply represents the set of data in that database as it is organized into schemas

and tables. The way you connect to the catalog determines the level of access to those schemas and tables. For other types of catalog, for example, catalogs that are mounted over external data catalogs or shares, the objects in the external system are still organized into schemas and tables but may not exist physically in such a form. For example, a catalog that is mounted over an AWS Glue data catalog may contain tables that actually exist as multiple files on AWS S3 cloud storage.

- **Mounting a Catalog**: This is the process of adding a new catalog connection to an external system so that the metadata in that system can be represented in the Autonomous Database catalog in a standard form.

A **Local** catalog is all the objects in or directly connected to your Autonomous Database as described in the above points. In addition to the local catalog, Catalog Explorer can add in additional catalogs so that you can discover data from other sources.

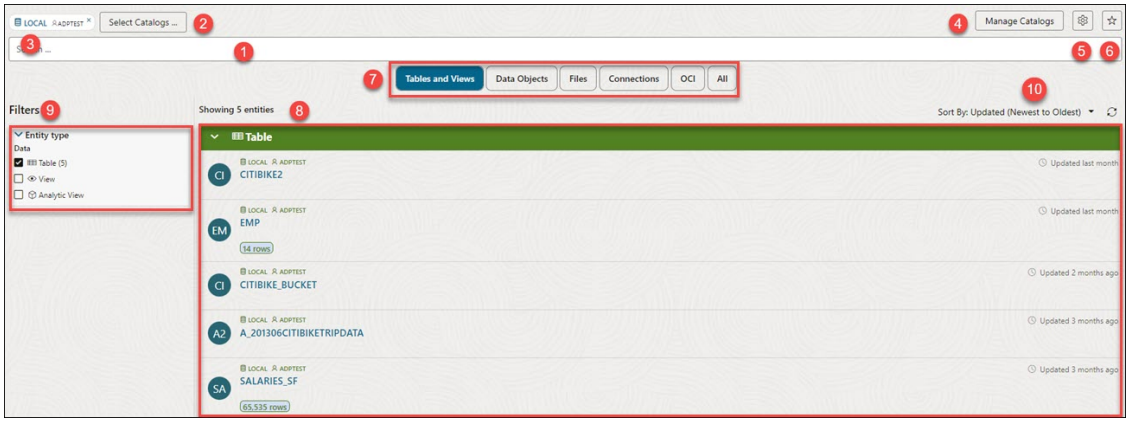
Additional catalogs can be mounted over:

- Mount Autonomous Database as Catalog: You can add a catalog to another autonomous database in your tenancy, which will allow you to search and discover all of the objects in that database based on the permissions for connection.
- Mount Database Link: The tool can mount an existing Database Link. Where a database link already exists, for example to an on-premises Oracle database, this can be mounted as a catalog. As with adding a catalog over an Autonomous Database, this option enables all the objects that are available to connection for search and discovery.
- Mount Data Share: The tool can mount registered live share or delta share as a catalog.
- Mount Amazon Glue: A catalog can be mounted over an AWS Glue Data Catalog database region.
- Mount OCI Data Catalog: The tool can mount data registered in OCI Data Catalog.

Once a new catalog has been mounted, you can add the catalog in the Catalog tool, and search and discover data from that catalog, or across many catalogs.

When you first open the Browse catalog page, it contains the **Search catalog** field, the **Catalog Settings**  button, **Manage Catalogs** icon, a list of recently viewed objects, and several suggested searches. You can enter a search string, click on one of the recent objects to see its details, or click on **Show search suggestions** icon to view suggestions on the right side of the page.

As soon as you select Catalog menu from Data Studio menu bar, the page is displayed as below.



When you hover your mouse over a particular catalog entity, it will show a floating toolbar of common actions. The list is context-sensitive depending on the capabilities available for the particular entity. For example, tables in the current database offer more possibilities than tables in remote catalogs.

The Catalog page contains:

1. **Search** field

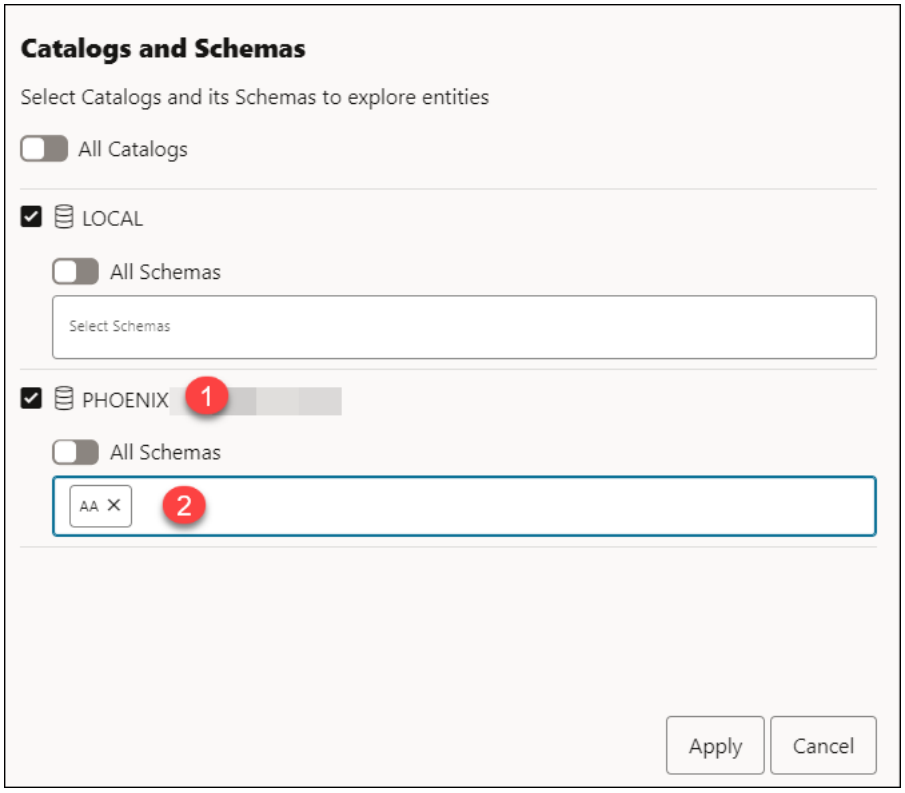
   You can search for the entities in the search field. For more details see Browse and Search Catalogs.

2. **Select Catalogs...** field
   Clicking this icon opens the **Catalogs and Schemas** dialog box.

   You can view a list of all mounted catalogs and schemas in this dialog box.

   Select specific schemas from a database of your choice to refine your search.

   - To show tables of only the **AA** schema from **PHOENIX** database, in addition to the default **local** schema:
     - On the **Catalogs and Schemas** dialog box, click the **PHOENIX** database and select **AA** from the **Select Schemas** search field.



3. **Default Database name and Schema name**
   This field displays the default database name along with the schema name.
4. **Manage Catalogs**
   Refer to Manage Catalogs for detailed information.
5. **Catalog Settings**
   Refer to Catalog Settings for detailed information.
6. **Saved Search Suggestions**

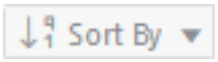
   Click to open or close the **Saved Search Suggestions** panel.
7. **Filters** option
   Click any of the filters icon to filter and refine the search results.

   Refer to Browse and Search Catalogs for detailed information.
8. **Display area**

   The area beneath the **Search catalog** field displays the entities returned by a search and that match the filter criteria set in the **Filters** panel. You can sort the entities by clicking the **Sort By** button and then setting sort values.

   Hover besides the name of the entity to view details about the entity.
9. **Faceted Filters** panel

   Select one or more filter values to limit the entities shown on the page. Only those entities that match the filter values are shown. That is, the items returned by a search are filtered by these filter settings. The filters also change depending on the quick filters you select. See Filter Entities.
10. **Toolbar**
    The toolbar appears after you run an initial search. It contains these buttons:

- **Sort By** 

  To select sorting values, click the **Sort By** button to open the list of options. Then click the **Ascending** or **Descending** icon next to **Name** or **Updated**. For example, if you select the **Ascending** icon next to **Name** the entities will be sorted in alphabetical order by entity name. If you select **Descending** icon next to **Updated**, the entities will be sorted in reverse of the latest updated entity..

  Click **Reload** to refresh the choices in the list.

Browse and Search Catalogs

The Catalog Tool aims to allow you to find data quickly and easily, whether or not the data is currently resident in the connected Autonomous Database.

Using the Catalog tool, you can discover data in connected systems, preview it, and choose whether or not to load, link or feed the data into the local database.

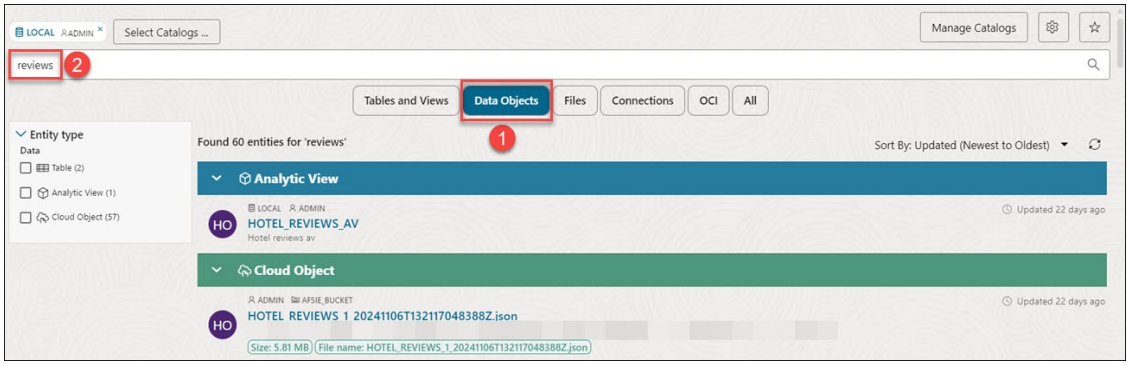
You can view the search bar on the Catalog's home page. To enter a search term, click on the search field located at the top. You can filter the search entry by selecting a quick filter displayed below the search field.

You must first select the type of assets you want to search from the filter buttons at the top of the screen, then search for objects by name.

Selecting the quick filters helps you display filtered searches in the display area. You can combine multiple filters. They are:

- **Tables and Views**
  This filter searches only for the database tables (both internal and external tables) and views (this includes both standard database views and analytic views)
- **Data Objects**
  This filter searches for data objects in the database and on connected cloud storage. This is a superset of the **Tables and Views** and **Files** scopes.
- **Files**
  This filter searches only for files in connected cloud storage.
- **OCI**
  This filter searches for other Oracle Cloud Infrastructure objects, such as other Autonomous Databases, buckets on OCI cloud storage, or a registered OCI Data Catalog. This field is available if you have an OCI credential. You can create an OCI credential from the Data Studio Settings icon.
- **Connections**
  This filter searches for connections that are registered in Data Studio, such as connections to external Data Catalogs or databases.
- **All**
  This filter searches for all objects in the catalog. Since there are many thousands of such objects, this scope may be slower to use.

For example, to find all tables, views or files that may contain data related to reviews, click **Data Objects** and search for *reviews*:



This displays all data objects that contain review data.

Sorting and Filtering Entities

On the Catalog page, you can sort and filter the displayed search results.

Sort Entities

To set a sort order for the entities on the page,

1. Click **Sort by** on the toolbar.
2. Find the value you want to use to sort, for example **Entity Name**.
3. Click the **Name (A - Z)** (Ascending) icon or the **Name (Z - A)** (Descending) icon next to that value.
4. You can set only one parameter. You can also set **Updated: (Oldest to Newest)** or **Updated: (Newest to Oldest)** based on the time you update the entity.

Filter Entities

Restrict which entities that are returned by a search are displayed on the page by setting filters in the **Filters** suggestions panel on the top of the Catalog page. Select one or more filter values. Only those entities that are returned by a search and that match the filter values are shown. Selecting all or none of the options shows all entities returned by the search.

By default, system tables and private tables are not displayed. To display them, click **Catalog Settings**  in the toolbar and then click the **Show system tables** slider or the **Show private tables** slider, or both.

Filter Entities

By selecting different quick filters, the faceted results of the entity types also vary.

The faceted search displays various options in entity types based on different categories. You can select multiple entity types in various categories for more specific search results. It displays the count value in bracket.

The following table lists the entity types and filter condition subcategories with their description:

Table 3-6 Entity Types and Filter Condition

Field	Description
Filter for	

Table 3-6 (Cont.) Entity Types and Filter Condition

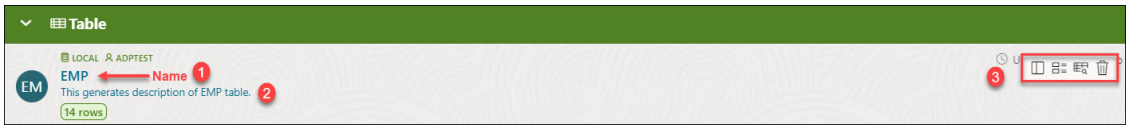
Field	Description
Entity Type	Specifies the name of the entity type under which you select entity filters
Filter condition subcategory	Data
Data Entity type values	• Table • View • Analytic View • Column • File • Cloud Object
Share Entity type values	• Share • Share Recipient • Cloud Link Namespace • Share Schema • Share Table
Connection Entity type values	• Catalog • Cloud Storage Link • Data Catalog Link • Live Table Feed • Share Provider • Cloud Link • Data Catalog Asset • Database Link • Share Link
Dictionary Entity Type Values	• Attribute • Attribute Dimensions • Directory • Credential • Function • Level • Mining Model • Procedure • Synonym • Schema
OCI	• Autonomous Database • Bucket • Compartment • Data Dictionary • Database Connection • Data Catalog • Vault • Vault Secret

Display Entity Search

If you do not select a quick filter, the **Tables and Views** filter is selected by default.

You can also view the Table after you enter the search term, a search is performed across **Tables and Views**.

Expand the Table header to view the Table name.



1. You can view the schema name. The icon color uniquely reflects schema or catalog. All entities within the same schema or catalog have the same entity color. If multiple catalogs are being searched, the catalog name appears next to the schema name.
2. **Auto-generate description**: You can also Auto generate a description by clicking the **Add description** link.
   On the **Add Description** wizard, click **Auto-Generate**. You need to set up an AI profile using the Data Studio Settings icon to configure the profile. Click **OK** on the **Add description** wizard. You can now view the auto generated description. Your profile must be AI enabled to generate auto description.
3. **Viewing Entity Details**
   You can view the entity details of the selected entity by just clicking it. You can view the menu of the available actions by hovering over the entity. The icons vary from entity to entity.

   For **Table and Views** type of entity, you can view the set of available actions of the entity:

   - **View Details**: Select **View Details** to view details about the table.
   - **Gather Statistics**: Select **Gather Statistics** to display new statistics after you modify the table's structure. Refer to the Gathering Statistics for Tables for more detailed information.
   - **Create Share**: Select **Create Share** to share from the associated table.
   - **Register to Cloud Link**: Select **Register to Cloud Link** to register the table for remote access for a selected audience you define. See Registering Cloud Links to Access Data section for more information.
   - **Create Analytic View:** Select **Create Analytic View** to create an Analytic View from the selected associated table. See Creating Analytic Views
   - **Query**: Selecting **Query** takes you to the Data Analysis page with the selected table.
   - **Export Data to Cloud Storage Location**: Select Export Data to Cloud to export data to a cloud object store.
   - **Edit**: Select **Edit** to edit the properties of the table. Refer to Editing Tables for more information.
   - **Drop**: Select **Drop** to delete the table.

Catalog Settings

The Catalog Settings can manage the general settings, add, and view query scopes, create and view saved searches and change the look and feel of the Catalog page.

You can alter the Catalog tool's user interface and page layout through themes and theme styles in the Catalog Settings.

Select **Show Settings** on the Catalog page to set the behavior of the Catalog. When you set these options, they take place immediately and are also saved as the default behavior for the

page. Clicking **Show Settings** opens the General tab where you can view the following options.

- **Show system tables**

  Select this option to include system tables in the search results.

- **Show private tables**

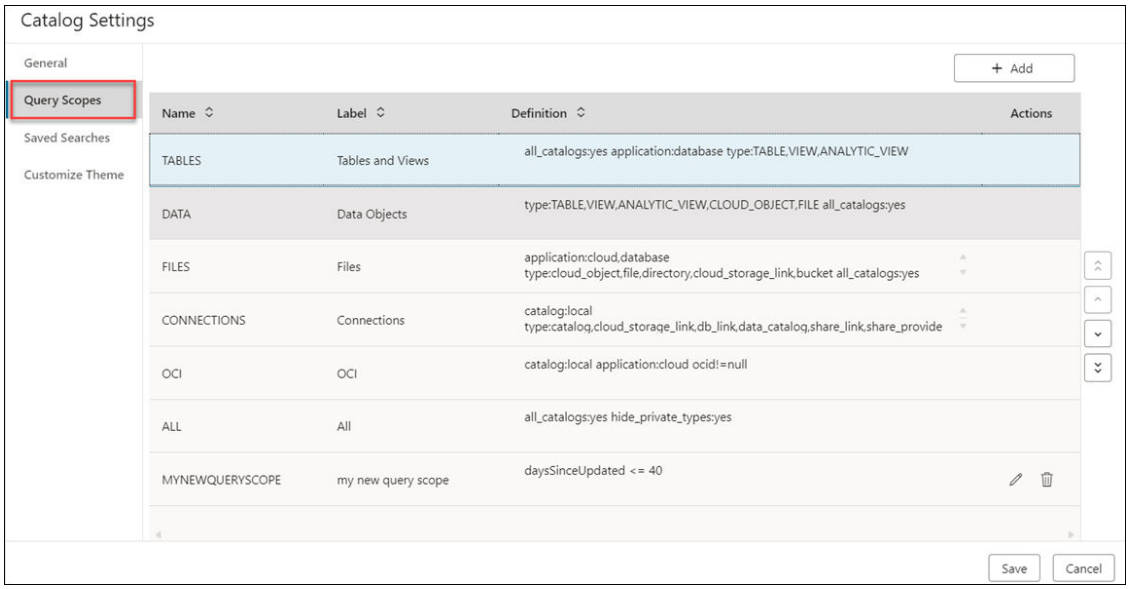
  Select this option to include private tables in the search results.

- **Page size**

  Select the number of entities to display on the page.

Gradually progress to use the **Query Scopes** tab to view, create, and delete query scopes. You can search for catalogs and save this search using query scopes.

Selecting Query Scope icon enables you to create new query scopes based on your preferences.



You can view, create, and edit the previously saved query scopes in this wizard. The query scopes are categorized based on by whom its created.

- Select **Custom** to view, create, edit, or delete the query scopes created by you.
- Select **Predefined** to view the query scopes which are already defined by the Database Actions. This option does not allow you to create, edit, or delete the query scopes.
- Select **All** to view all the query scopes. This includes Custom and Predefined query scopes.

Click **+Add** to create a Query Scope. Click **Cancel** to cancel its creation.

Specify the following fields in the Query Scope tab.

- **Name**: Enter the name of the Query Scope. This is a mandatory field.
- **Label**: This is a mandatory field. Enter a descriptive name here. You will use this field to refer to a query scope.

- **Definition**: Enter the Oracle Autonomous Database Data Definition Language (DDL) that creates the search entity. This is the same search criteria you enter in the Search Catalog field.

Once you have added the new query scope, it is visible in the list of query scopes in the Catalog Settings wizard.

Progress to the **Saved Searches** tab.

You can save your time from redefining the same search again in the future. You can diagnose problems faster since you are just few clicks away from accessing a saved search. Here is how you can create a saved search.

- Select this tab to save the five previous search queries that you entered into the **Search catalog** field. When the **Search catalog** field is empty and you click in the field, the last five search queries are listed in the drop-down list. (Any predefined searches selected from the **Suggestions** panel won't appear in this list, unless you edited them to make a new search.)

  The **Saved Searches** tab saves the current search query as per the criteria you specify or create a new saved search. Click **+Add** to create a new saved search.

- On the **Saved Searches** tab, click **+Add** to create a new search.
  - On the **Create New Search** dialog, specify the **Title**, **Scope**, **Definition** and **Description** (optional field) in their respective fields. For example, you want to search for a specific Entity type and a specific owner.
  - Select **Save**.

  After the creation of the saved search, it appears on the list of Saved Searches.

  You can change the columns displayed in the search results by clicking the pencil icon in the Actions column. Click the **delete** icon in the Actions column to delete the search you save. The saved searches you create are available for selection in the Saved Search panel in the right of the Catalog page.

- Progress to the **Customize Theme** tab to expand each entity type to select the properties and preview it.
  - On the **Customize Theme** tab, you can select the entity type whose appearance you want to define by editing the options in this tab.
    You can view the following menu options as you move from left to right.

Options	Description
Entity Type	This option describes the type of entity you want to edit. You can expand the entity type to select properties and to preview what the change looks like.
Is actions toolbar?	This option when enabled displays the actions menu items as floating toolbar on hover, else displays the actions menu.
Header background color	This option enables you to set the background color of the title of the entity header. This is the same as the text color for the properties of the entity listed.
Text color	This option enables you to set the text colour of the entity title.

Options	Description
Text highlighter	This option enables you to set the highlight colour of the box surrounding the properties of the entity.

Click **Save** in the Catalog Settings dialog to save all your preferences you set for the Catalog page.

Manage Catalogs

You can now enable multiple catalogs using the Catalog Tool.

You can mount additional catalogs from other Autonomous Databases in the same tenancy, so that you can search for data across multiple databases.

Mount Autonomous Database as Catalog

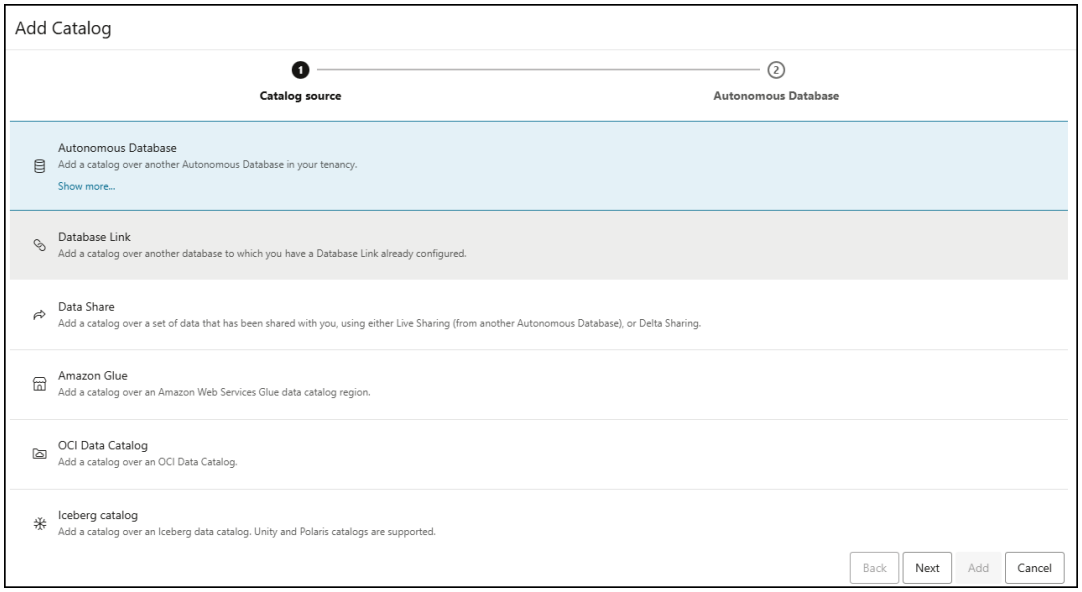
You will use an existing Autonomous Database to create a Database link and provide catalog name to mount it.

> **Note:**
>
> To mount an Autonomous Database as a Catalog, you must have the `CREATE DATABASE LINK` and `CREATE ANY DIRECTORY` system privileges. The ADMIN user is granted `CREATE DATABASE LINK` and `CREATE ANY DIRECTORY` system privileges. The ADMIN user can grant `CREATE DATABASE LINK` and `CREATE ANY DIRECTORY` system privileges to other users.

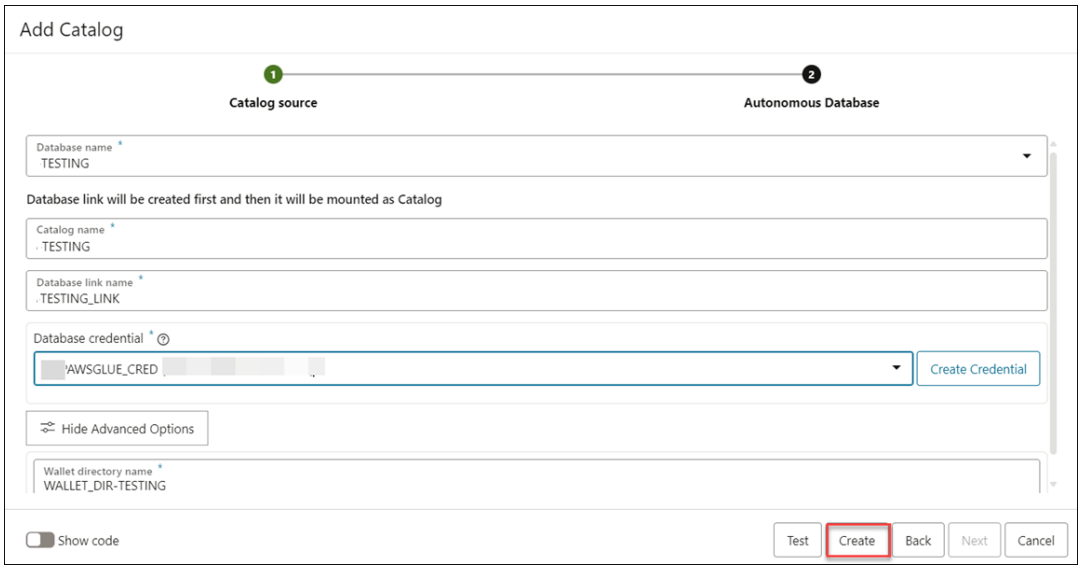
To add an autonomous database as a catalog, you need to perform the following operations:

- Store your object storage credentials in your Autonomous Database
- Create a database link or use an existing database link
- Mount the database link as a catalog

1. On the Catalog page, select **Manage Catalog**. Click **+Add**.
2. On the **Add Catalog** page, select **Autonomous Database** as Catalog Source.



3. Select the **Database name** from the drop-down. For example, *TESTING*.
4. You can view the following fields after you select the database name. Specify the following field values:



- **Catalog Name**: Enter the name of the catalog you want to mount the database link as. For example, *TESTING*.
- **Database link name**: Enter the name of the database link the tool creates using the existing database. For example, *TESTING_LINK*.
- **Database credential**: Select the credential from the drop-down which is required to create the database link. For example, *AWSGLUE_CRED*.
  Click **Create Credential** in case there is no existing credentials.
- Under **Show Advanced Options**, select the **Wallet Directory name** from the drop-down. You need to have permission to access the directory. For example, *WALLET_DIR-TESTING*.

> **Note:**
>
> The wallet file, along with the Database user ID and password provide access to data in the target Oracle Autonomous Database.

- Click **Test** to test the creation of the Database Link.

Click **Create** to create the autonomous database as Catalog.

After successful creation, you can view the Autonomous database listed under Manage Catalog.

Mount an existing Database Link as Catalog

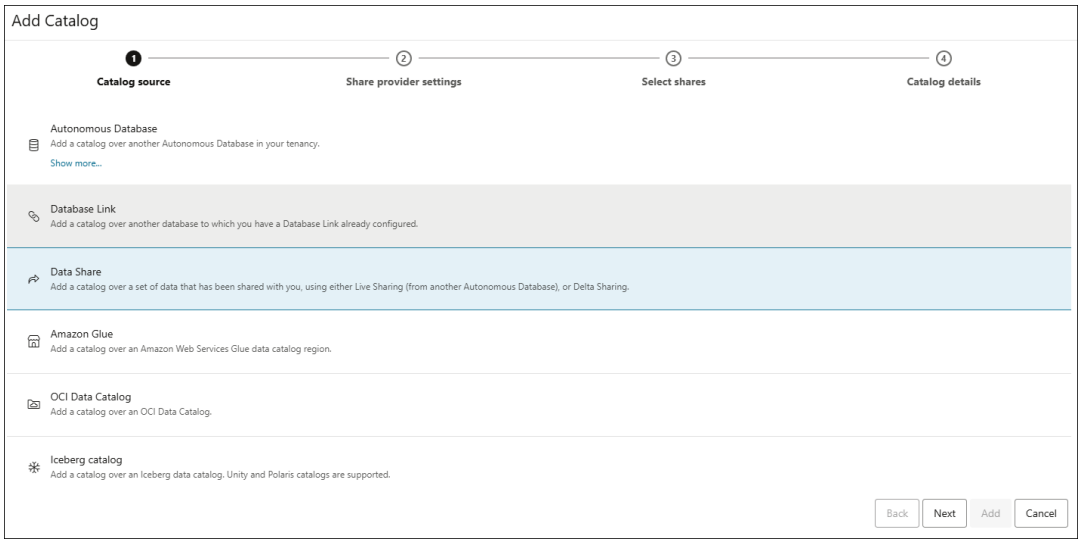
You can instantiate a catalog schema as a local database user. If a database link has already been set up, you can mount a catalog over the linked database in order to search and discover data sets in that database.

Mounting a catalog requires the following information:

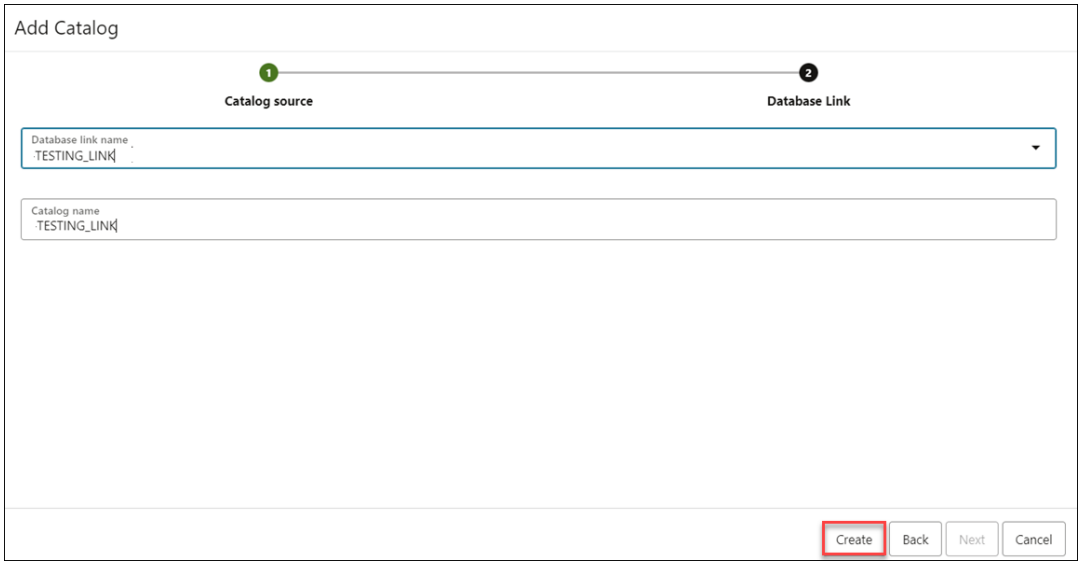
- Name of the database link
- Cloud credentials required to access the data

You can add a database link to one of the databases and mount its metadata as a new catalog.

1. On the Catalog page, select **Manage Catalog**.
2. Click **+Add**.
3. On the **Add Catalog** page, select **Database Link** from Catalog Source.



4. Select the **Database Link** tab.



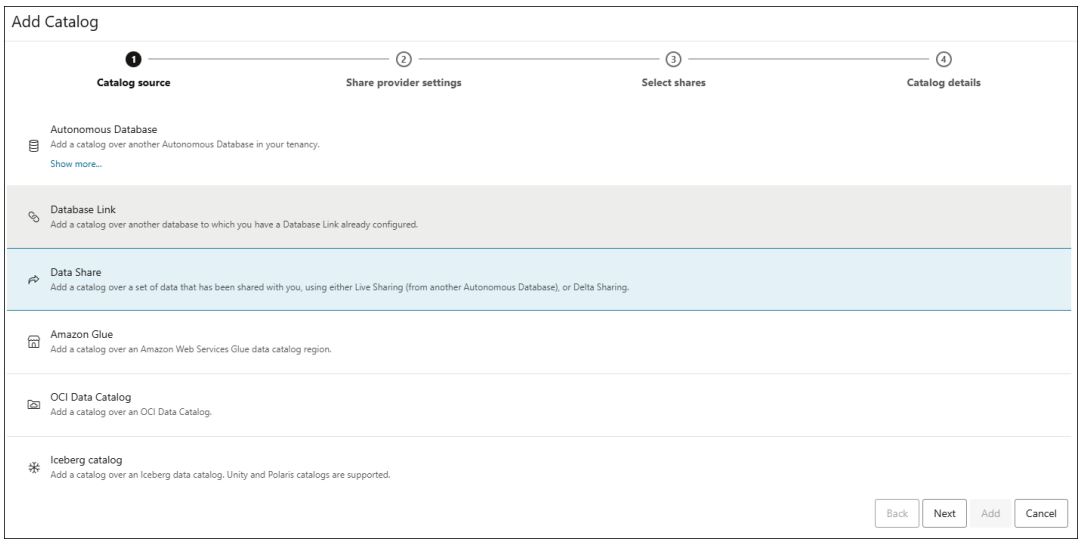
5. Select the **Database Link name** from the drop-down. For example, *AA-TESTING_LINK*.
6. Enter the name of the catalog in the **Catalog name** field. For example, *AA-TESTING_LINK*.
7. Click **Create**.
   You can view the newly mounted catalog in the **Manage Catalogs** dialog box.

   After you have successfully mounted a new catalog over a database connected using a Database link, you can select the catalog in Catalog Explorer and search for data in it.

Mount either Live Share or Delta Share Provider as Catalog

You can mount catalogs from registered data shares. This enables the consumers to access tables defined by other providers in other databases. The external tables need to be defined just once.

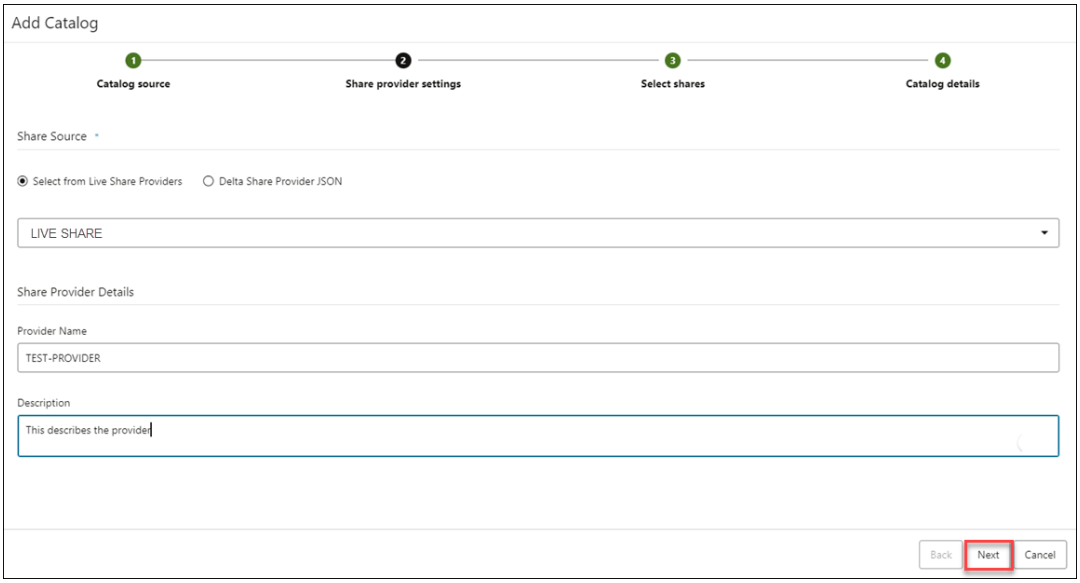
1. On the Catalog page, select **Manage Catalog**. Click **+Add**.
2. On the **Add Catalog** page, select **Data Share** from Catalog Source.



3. Click **Next** to progress to the **Share provider settings** tab.
   Under Share Source, select any of the available options:
   - **Select from Live Share Providers**
   - **Delta Share Provider JSON**

   In this example, you will select share source from Live Share Providers.

   Under **Share Provider Details**, specify the following details:
   - **Provider Name**: Enter the name of the share provider. For example, *TEST-PROVIDER*.
   - **Description**: Enter any description. This field is optional.



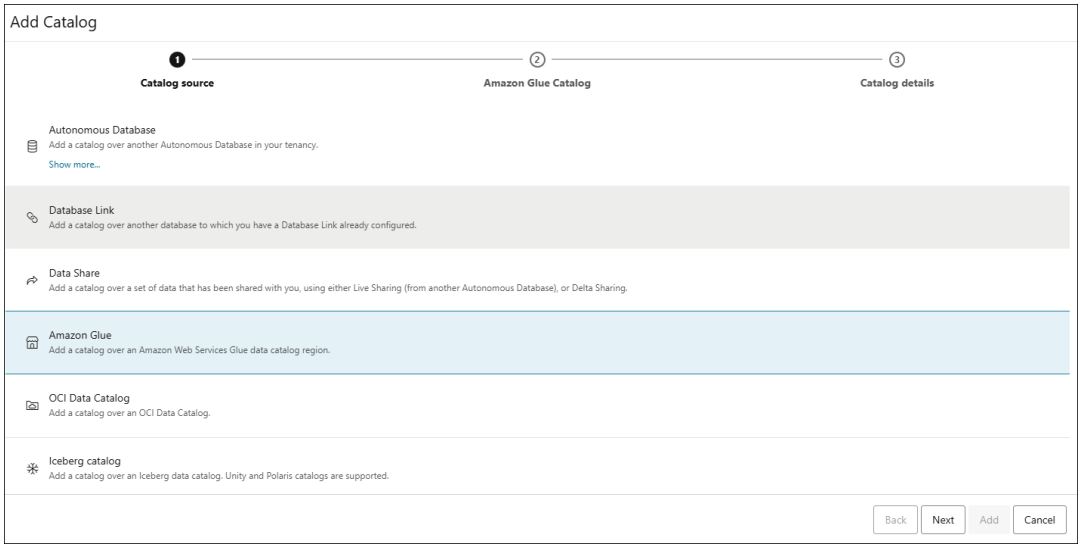
4. Click **Next** to progress to the **Select shares** tab.
5. On the Share provider settings tab, drag and drop the share you want to mount as a catalog from the Available Shares to the Selected Shares.
6. Click **Next** to progress to the Catalog details tab. On the Catalog details tab, specify the following details:
   - **Catalog name**: Enter a catalog name of your choice.
   - **Catalog description**: Enter a description.
7. Select **Create** to mount a Live Share or Delta Share as a catalog.

After you have successfully mounted a new catalog over a Share, you can select the catalog in Catalog Explorer and search for data in it.

Mount Amazon Glue as Catalog

You can mount catalogs from registered AWS Glue data connections.

1. On the Catalog page, select **Manage Catalog**.
2. Click **+Add**.
3. On the **Add Catalog** page, select **Amazon Glue** from Catalog Source.



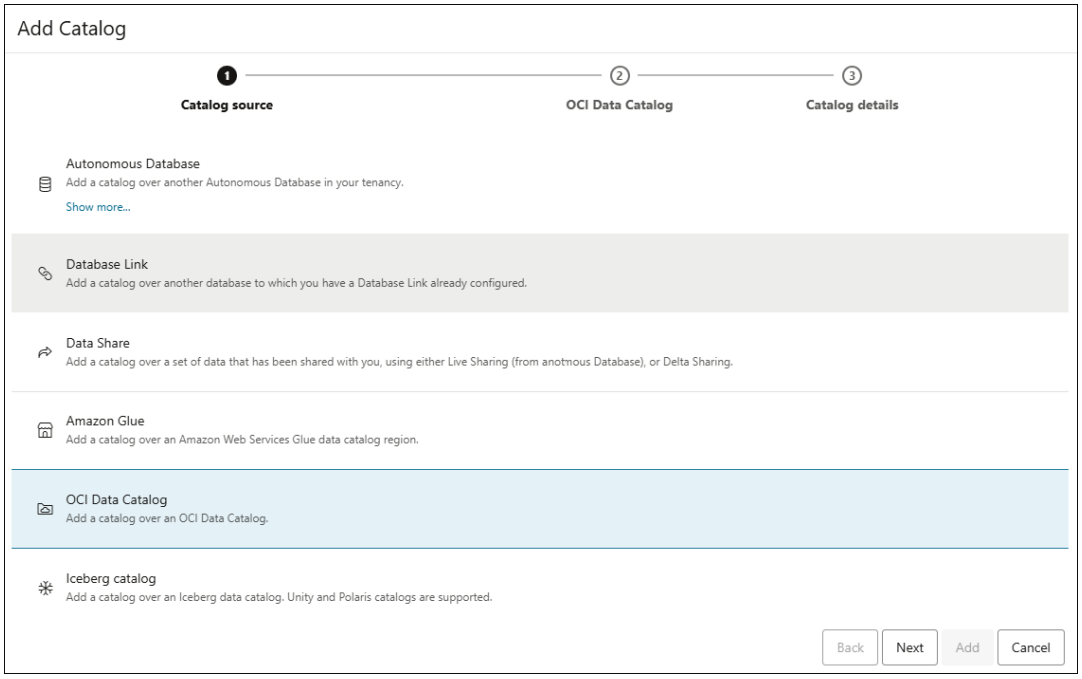
4. Click **Next** to proceed to the **Amazon Glue Catalog** tab.
5. On the **Amazon Glue Catalog** tab, specify the following details:
   - **Credential for Data Catalog Connection**: Select a credential from the drop-down.
   - **Region**: Select a region from the drop-down.
6. Click **Next** to progress to the **Catalog details** tab. Specify the following field values on the **Catalog details** tab:
   - **Catalog Name**: Enter the name that you want to use in Data Studio to refer to the catalog.
   - **Catalog Description**: Enter a description of the catalog.
7. Click **Create** to mount the AWS Glue Data Catalog region as a catalog.

After you have successfully mounted a new catalog over an AWS Glue Data Catalog region, you can select the catalog in Catalog Explorer and search for data in it.

Mount OCI Data Catalog as Catalog

You can mount catalogs from a registered OCI Data Catalog.

1. On the Catalog page, select **Manage Catalog**. Click **+Add**.
2. On the **Add Catalog** page, select **OCI Data Catalog** from Catalog Source.

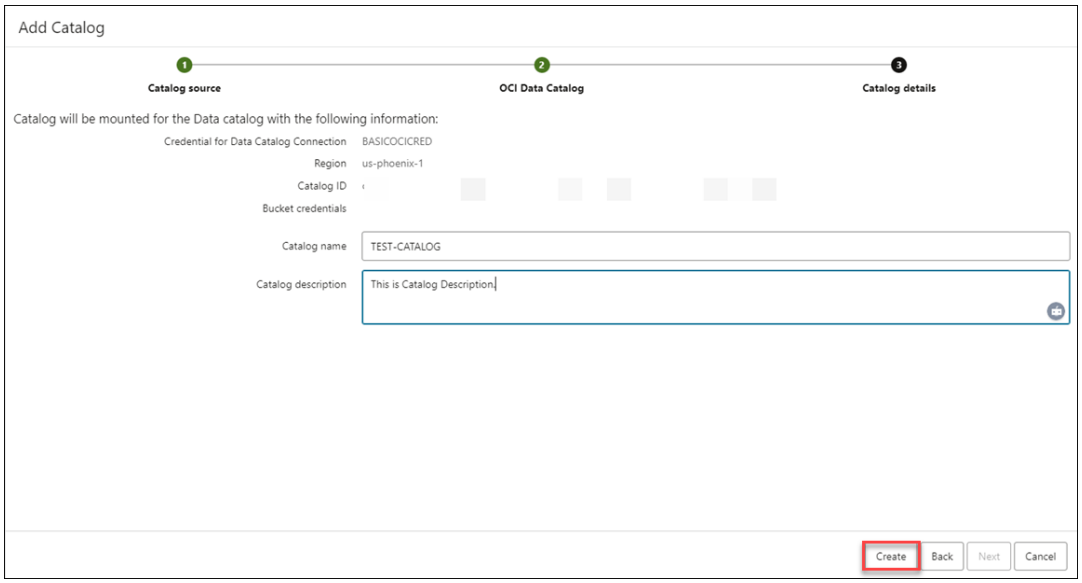


3. Click **Next** to proceed to the **OCI Data Catalog** tab.
4. On the OCI Data Catalog tab, specify the following field values:
   - **Credential for Data Catalog Connection**: Select a credential from the drop-down. This field is mandatory. For example, *BASICOCICRED*.
   - **Region**: Select a region from the drop-down. This field is mandatory. For example, *us-phoenix-1*.
   - **Data Catalog ID**: Select a Data Catalog ID from the drop-down. A Data Catalog ID is the ID of the account to which the data catalog belongs. For example, *catalog1*.

     Select **Use separate credential for Object Storage** to use a different credential other than the selected one.



5. Click **Next** to progress to the **Catalog details** tab.
6. On the **Catalog details** tab, specify the following field values:
   - **Catalog Name**: Enter the name that you want to use in Data Studio to refer to the catalog. For example, *TEST_CATALOG*.
   - **Catalog Description**: Enter a description of the catalog.



7. Click **Create** to mount the OCI Data Catalog as catalog.

After you have successfully mounted a new catalog over an OCI Data Catalog, you can select the catalog in Catalog Explorer and search for data in it.

Mount a Catalog over an Iceberg Data Catalog

You can add a catalog over an Iceberg Data Catalog. Integrating a catalog with Iceberg data involves connecting with supported catalogs such as Unity and Polaris. To integrate Iceberg Data Catalogs via Data Studio, you need to ensure your Iceberg Catalogs are properly configured and accessible.

To configure and access your Iceberg Catalog, you must have the following:

1. **Iceberg Endpoint**: Set up Polaris or Unity Catalog using the Iceberg REST API. This involves specifying the catalog type, URI, and warehouse location. You will need the REST endpoint for your external Iceberg catalog. This is the URL that provides access to the Iceberg metadata.
   Example: *https://your-iceberg-catalog.com/api/v1/iceberg*
2. **Token Endpoint**: The */v1/oauth/tokens* endpoint is provided by Iceberg REST catalogs to issue **Bearer Token**.
3. **Client ID & Client Secret**: These are credentials provided by the OAuth server to identify your application (in this case, the Unity and Polaris Catalog integration). You will use these credentials to generate Bearer Token.
4. **Bearer Token**: To authenticate requests to the Iceberg REST endpoint, you will use OAuth with the Client Credentials. The Catalog tool calls the token endpoint to receive **Bearer Token**.

5. **Namespace**: A namespace is equivalent to a schema in a database. This location will serve as the repository for your Iceberg tables within the catalog. Iceberg catalogs group tables into namespaces, which can be nested within other namespaces. The Iceberg specification supports arbitrary levels of nesting.
   For example, a `SALES` table could be defined within a `STORES` namespace, which is defined in a `RETAIL` namespace, which is itself defined within a `CORPORATE` namespace. The Iceberg location of such a table would be `CORPORATE.RETAIL.STORES.SALES`.

The following is a sample configuration to create a REST Catalog integration that uses OAuth to connect to Apache Iceberg tables:

```
CREATE OR REPLACE CATALOG INTEGRATION tabular_catalog_int
  CATALOG_SOURCE = ICEBERG_REST
  TABLE_FORMAT = ICEBERG
  CATALOG_NAMESPACE = 'default'
  REST_CONFIG = (
    CATALOG_URI = 'https://api.tabular.io/ws'
    CATALOG_NAME = '<tabular_warehouse_name>'
  )
  REST_AUTHENTICATION = (
    TYPE = OAUTH
    OAUTH_TOKEN_URI = 'https://api.tabular.io/ws/v1/oauth/tokens'
    OAUTH_CLIENT_ID = '<oauth_client_id>'
    OAUTH_CLIENT_SECRET = '<oauth_secret>'
    OAUTH_ALLOWED_SCOPES = ('catalog')
  )
  ENABLED = TRUE;
```

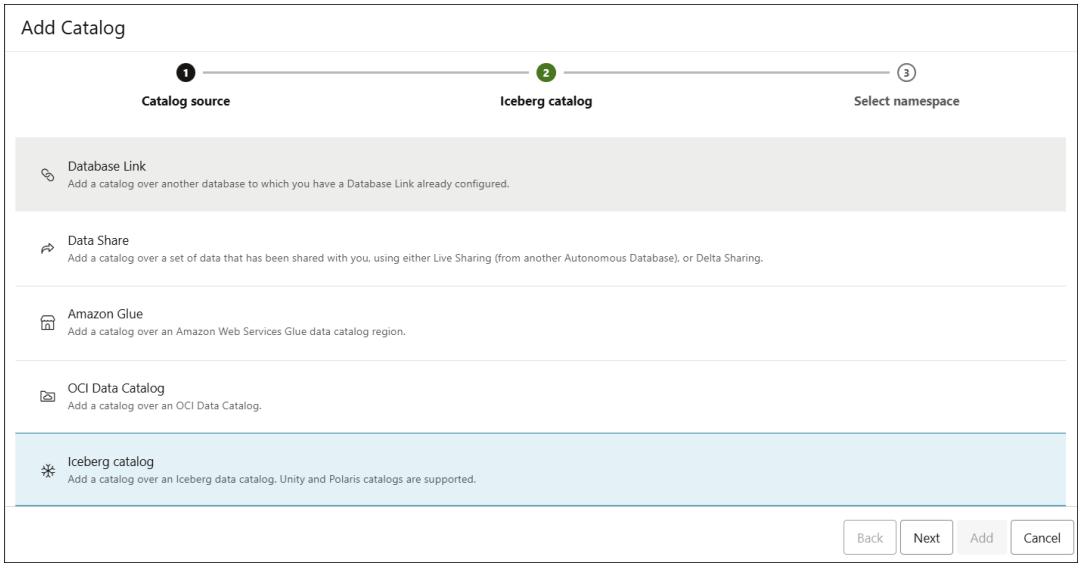
Here, `CATALOG_URI` is the REST end point to access the warehouse.

The `REST_AUTHENTICATION` procedure consists of all the parameters you need to create the `BEARER TOKEN` and use it to create the credential.

See Configure a catalog integration for Apache Iceberg™ REST catalogs for more information on this.

To mount a Polaris Catalog over an Iceberg Data Catalog:

1. On the Catalog page, select **Manage Catalog**. Click **+Add**.
2. On the **Add Catalog** dialog, select **Iceberg catalog** from the **Catalog source** tab.



Click **Next**.

3. Specify the following field values on the Catalog source tab of the Add Catalog dialog:

   - **Catalog name**: Enter the name of the Catalog you want to mount the Iceberg Catalog as.
     For example, *Polaris*.
   - **Iceberg catalog type**: Select *Polaris* from the drop-down field.
   - **Iceberg catalog endpoint**: It specifies your Polaris Catalog account URL. The endpoint for Snowflake's Polaris Catalog is usually in the format of *https://<account_locator>.snowflakecomputing.com/polaris/api/catalog*. For more information about account identifier formats, see the Snowflake documentation.
   - **Iceberg catalog credential**: Select a credential from the list of available credentials in the drop-down list. If you do not have a credential, you can create one. Select **Create Credential**. On the Create Credential dialog, specify the following field values:
     - Credential Name: Enter the name of your credential. For example, *POLARIS_CRED*.
       Under **Get New Bearer Token**,
       * **Access token end point**: Enter the Iceberg REST specification endpoint. This generates access token using OAuth2 defined as `/v1/oauth/tokens`. This endpoint was designed to facilitate authentication by exchanging client credentials (client ID and secret) for an access token.

         > **Note:**
         >
         > The `/v1/oauth/tokens` endpoint is deprecated in Iceberg version 1.6.0 and will be removed in version 2.0. You are encouraged to rely on external OAuth2 providers instead. See OAuth Authentication for more details.

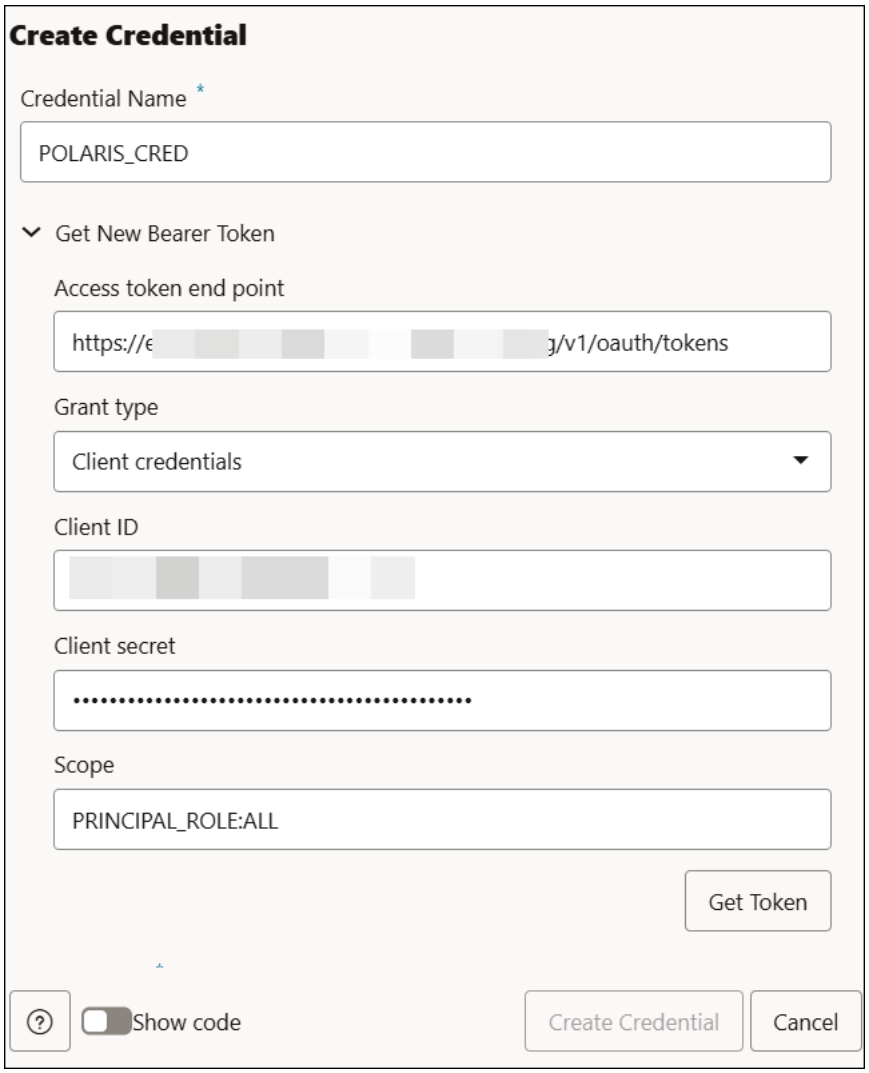
       * **Grant type**: Select *Client credentials*.
       * **Client ID**: Specify the Client ID.
       * **Client Secret**: Specify the Client Secret.

* **Scope**: *PRINCIPAL_ROLE:ALL*.
  Click **Get Token**.

  The **Bearer token** field gets populated.

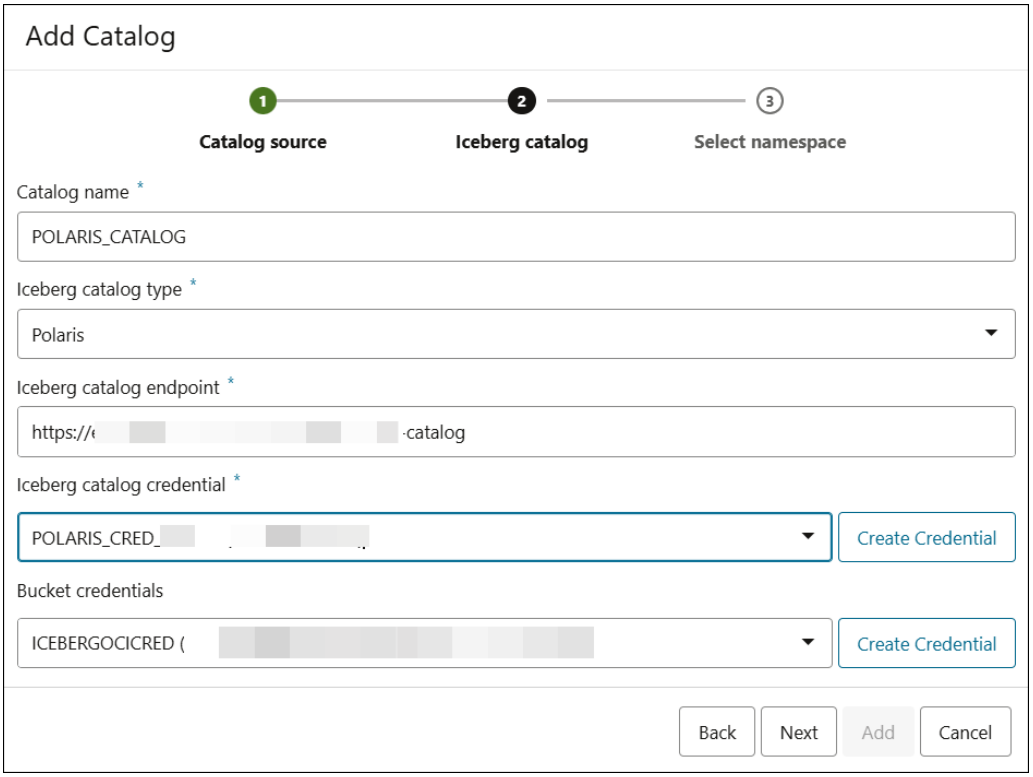
Click **Create Credential**.

POLARIS_CRED is displayed as Iceberg catalog credential.



> **Note:**
>
> The credentials are stored using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL` in the backend.
>
> For the sample example provided above, the procedure would be as follows:

```
BEGIN
  dbms_cloud.create_credential(
    credential_name  =>  'POLARIS_CRED',
    username         =>  'BEARER_TOKEN',
    password         =>  '<BEARER_TOKEN_VALUE This is
generated using token endpoint, client id, client secret>'
  );
END;
```



This credential allows the Oracle Autonomous Database to access the Iceberg table.

- **Bucket credentials**: This credential stores a cloud storage volume to store your Iceberg tables in Parquet format. This field is mandatory for Polaris catalogs.
  Click **Next** to proceed to the **Iceberg catalog** tab.

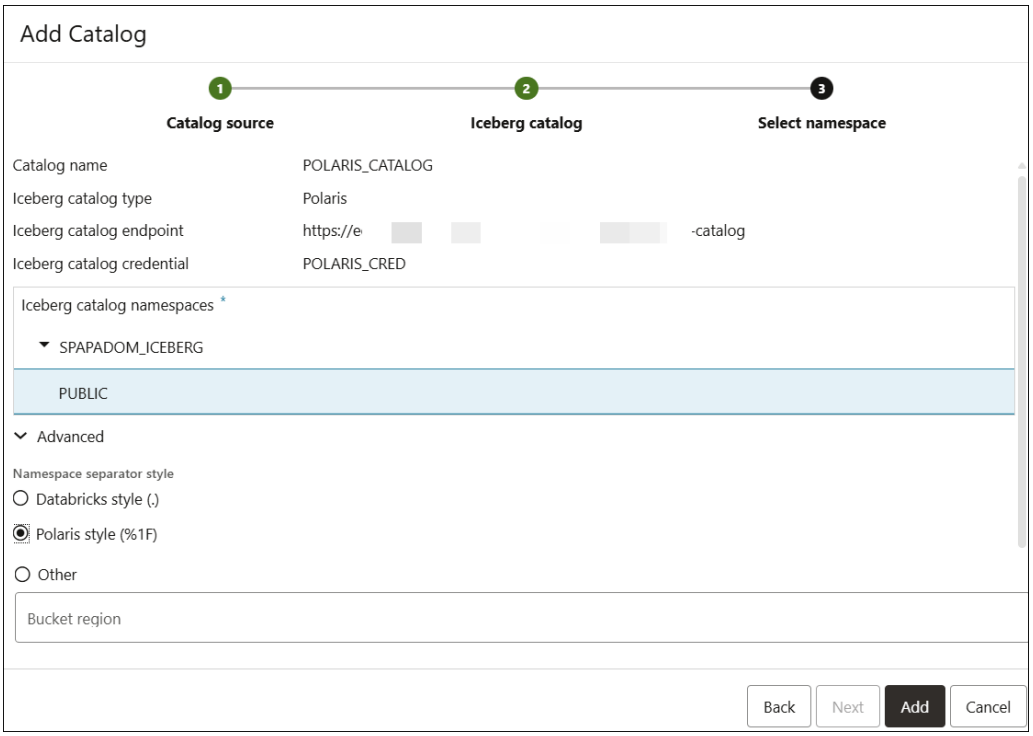
4. On the **Iceberg catalog** tab, you can specify the following field values:
   - **Iceberg catalog namespaces**: This field displays the hierarchical organization of tables. The **Iceberg catalog endpoint** consists of the root namespace. Select the namespace on which you want to have the catalog mounted.

- **Namespace separator style**: Select **Polaris style**. This is the default value. Select any value based on the way your catalog is configured.

  > **Note:**
  >
  > It is preferred that you retain the default value.

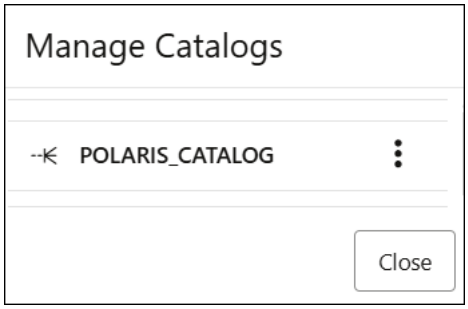
- **Bucket region**: Enter the bucket region. For example, *us-west-2*.
  Iceberg catalog endpoints, serve as interfaces for managing Iceberg tables stored in cloud storage like S3. The bucket region refers to the geographical location where the S3 bucket is hosted.

  This is mandatory field for Polaris catalogs.



Click **Add**.

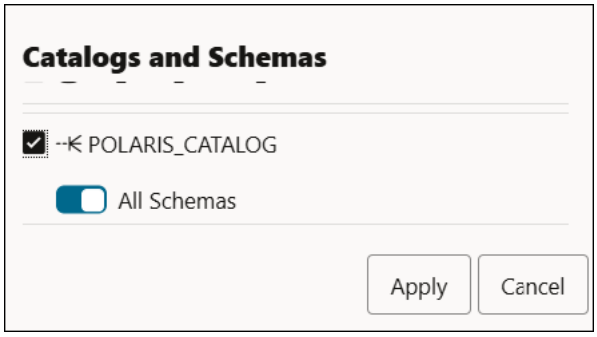
After you have successfully mounted a new catalog over an Iceberg Catalog, you can view the Polaris Catalog on the **Manage Catalogs** page. You can disable and delete the catalog from this page.



View the mounted Polaris Catalog on the Catalog Explorer

Once you have mounted the new catalog, you can load tables from the mounted catalog. To search for the catalog:

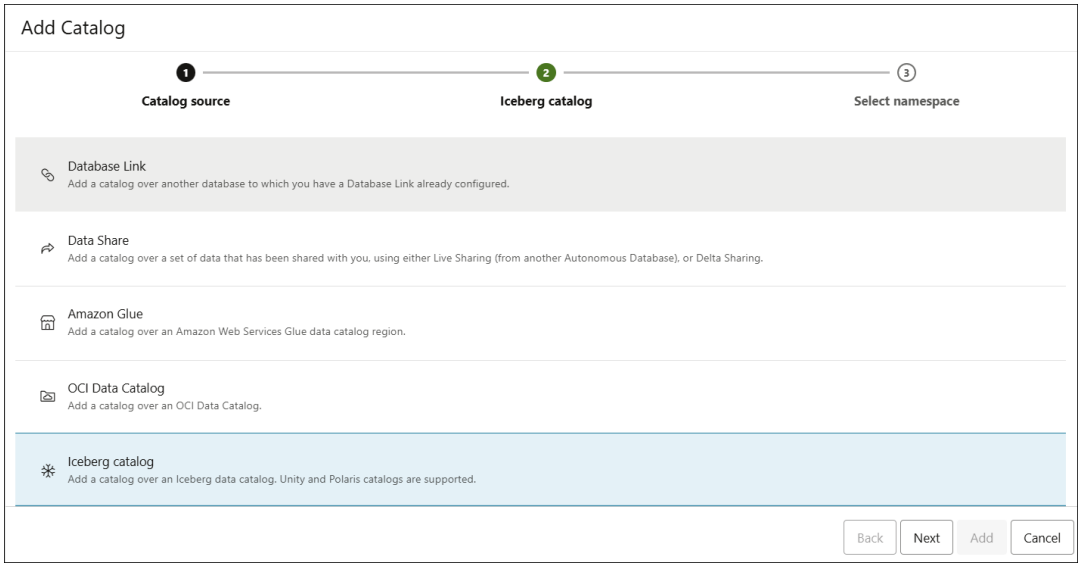
1. On the Catalog tool, click **Select Catalogs...**.
2. Click **POLARIS _CATALOG** filter on **Catalogs and Schemas** wizard to display the **POLARIS_CATALOG** catalogs.
3. Click **Apply** to view the results on the Catalog page.





To mount a Unity Catalog over an Iceberg Data Catalog:

1. On the Catalog page, select **Manage Catalog**. Click **+Add**.
2. On the Add Catalog dialog, select **Iceberg Catalog**.

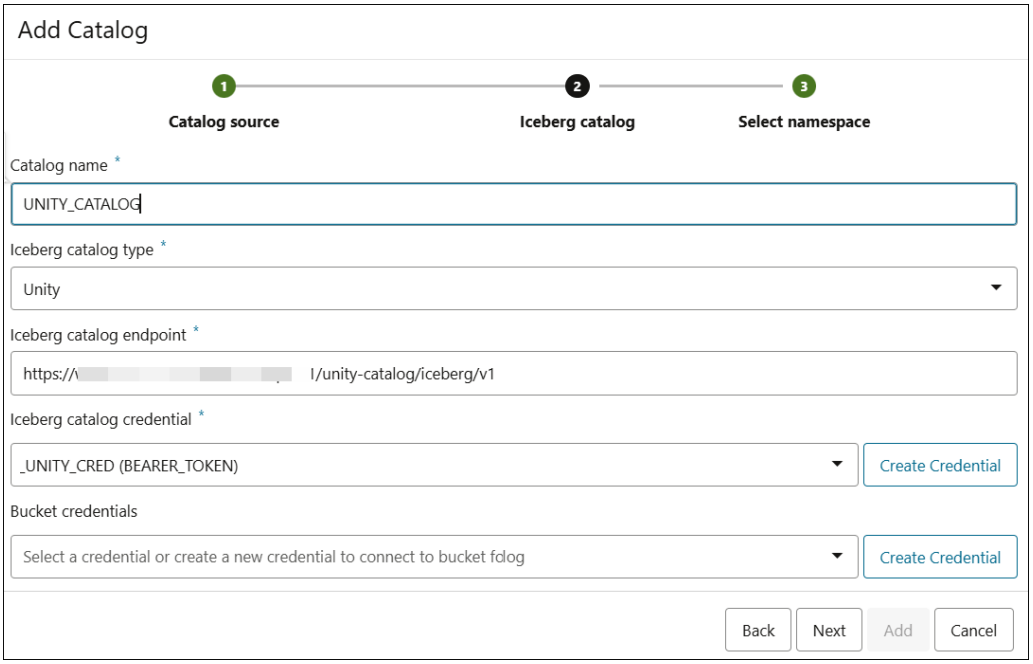


3. Click **Next**.
4. Specify the following field values on the Catalog source tab of the Add Catalog dialog:
   - **Catalog name**: Enter the name of the Catalog you want to mount the Iceberg Catalog as.
     For example, *Unity*.
   - **Iceberg catalog type**: Select *Unity* from the drop-down field.
   - **Iceberg catalog endpoint**:
     For Unity Catalog, you typically use the Databricks REST API. The endpoint for accessing Iceberg tables via Unity Catalog is usually in the format of *https://<databricks-instance-url>/api/2.1/unity-catalog/iceberg*.

     > **Note:**
     >
     > See https://docs.databricks.com/aws/en/external-access/iceberg for more information.

   - **Iceberg catalog credential**: Select a credential from the list of available credentials in the drop-down list. If you do not have a credential, you can create one. Select **Create Credential**.
     On the Create Credential dialog, specify the following values:
     – **Credential Name**: Enter the name of your credential. For example, *UNITY_CRED*.
     – **Bearer Token**: Paste the Bearer Token value.
       Select **Create Credential**.

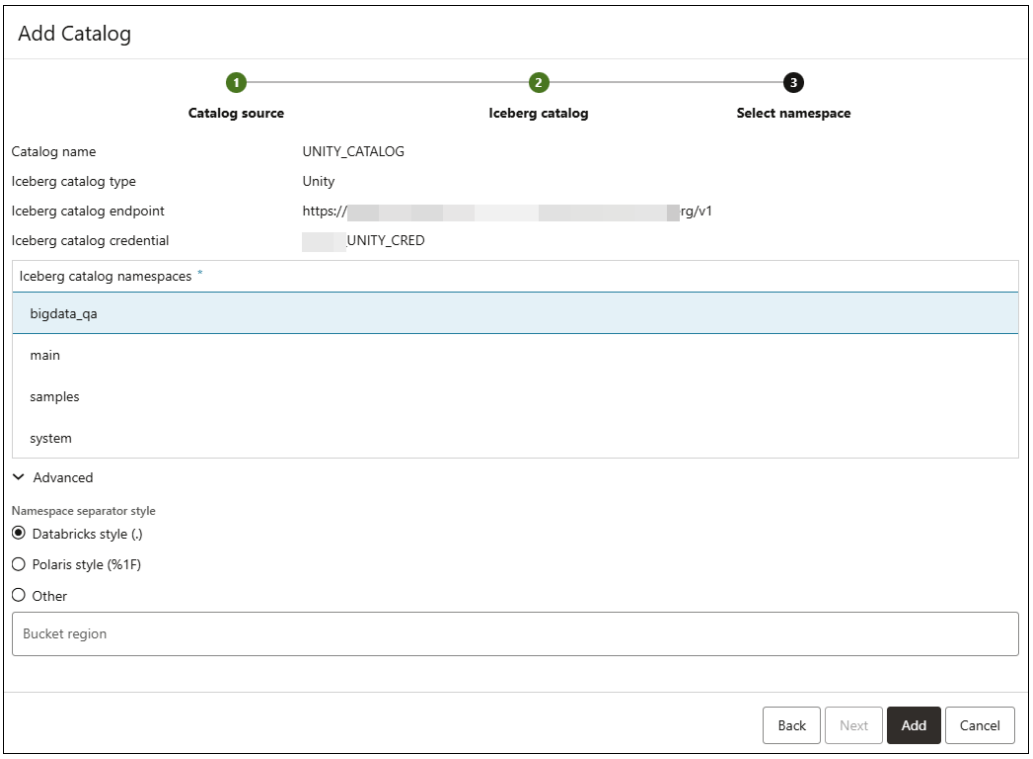
       UNITY_CRED is displayed as **Iceberg catalog credential**.



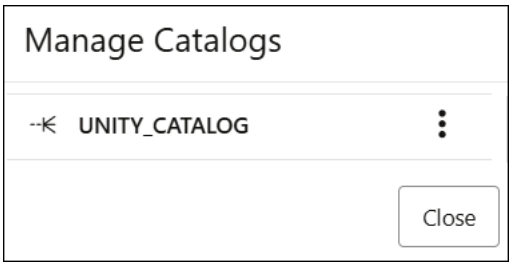
Click **Next** to proceed to **Iceberg catalog** tab.

On the **Iceberg catalog** tab, you can specify the following field values:

- **Iceberg catalog namespaces**: Select any one of the schemas you need access to.
- **Namespace separator style**: The namespace separator style typically refers to how namespaces are structured and separated within the catalog. By default, **Databricks style** is selected.
- **Bucket region**: This field is not mandatory.



Click **Add**.

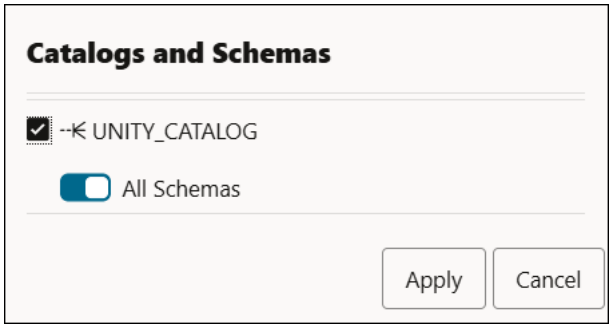


After you have successfully mounted a new catalog over an Iceberg Catalog, you can view the Unity Catalog on the **Manage Catalogs** page. You can disable and delete the catalog from this page.

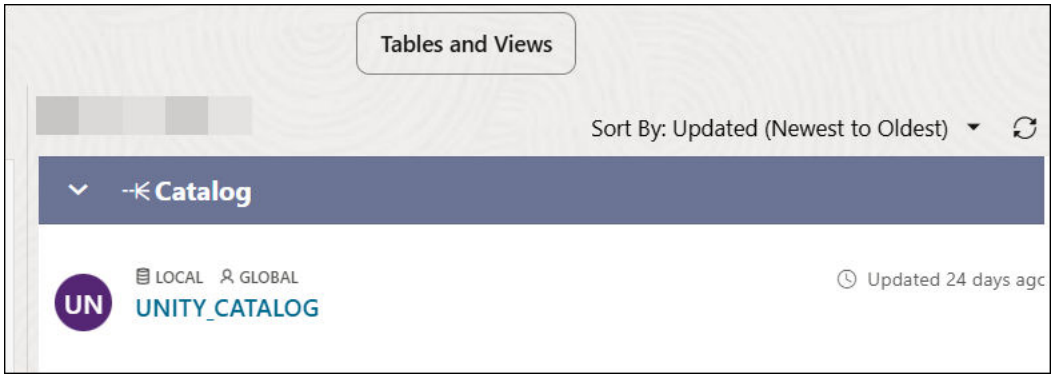
View the mounted Unity Catalog on the Catalog tool

Once you have mounted the new catalog, you can load tables from the mounted catalog. To search for the catalog:

1. On the Catalog tool, click **Select Catalogs…**.
2. Click **UNITY _CATALOG** filter on the **Catalogs and Schemas** wizard to display the **UNITY_CATALOG** catalogs.



3. Click **Apply** to view the results on the Catalog page.



- Inline Data Load of Delta Share from the Catalog Tool
  You can load a Delta Share from the Catalog tool without leaving the Catalog tool for loading.
- Query Remote Catalogs and Databases
  You can access and query external data by mounting catalogs. This feature uses a simple SQL syntax that is similar to how you query data using database links.

Inline Data Load of Delta Share from the Catalog Tool

You can load a Delta Share from the Catalog tool without leaving the Catalog tool for loading.

You need to first download your share profile to receive shares.

1. **Download your share profile**:
   For the data share you want to receive, download your profile sent from the share provider.
   - Open the invitation email you received from the data share provider, when the data share was published.
   - Copy the profile URL from the email.
   - Paste the URL into your browser and click **Get Profile Information**.
   - Your profile has been downloaded to your local machine. The Profile information downloaded message page is displayed.

   You can subscribe to the data share provider using the Catalog tool.

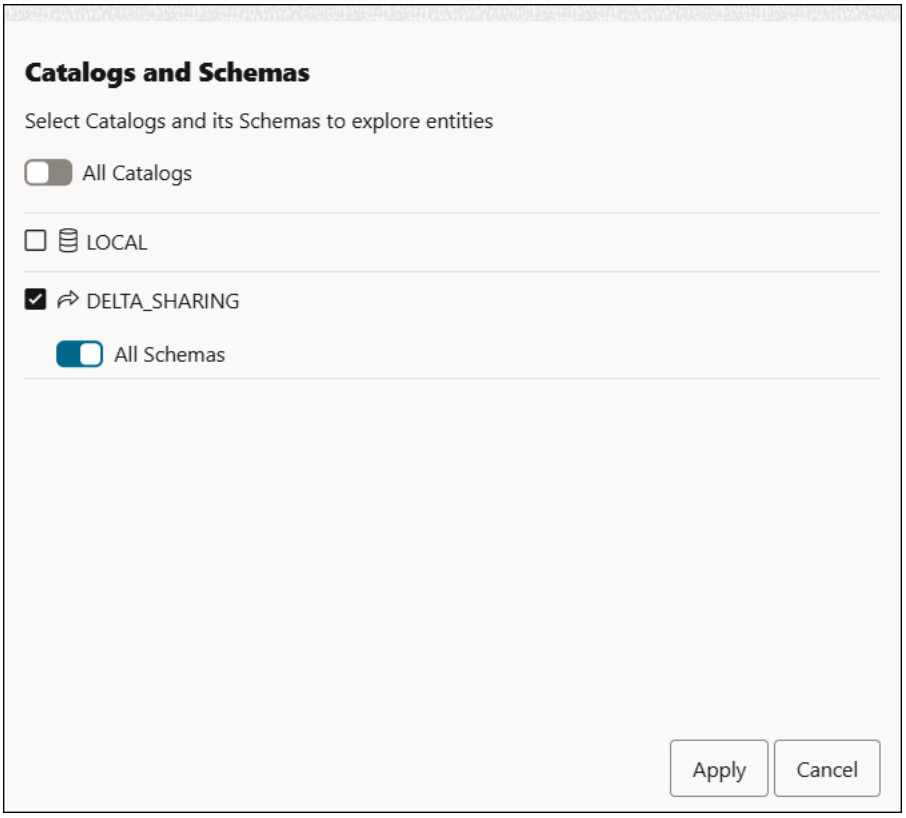
   To subscribe, you need to use the information contained in the uploaded JSON profile you received from the share provider.

From the **Manage Catalogs** feature of the Catalogs tool, you can upload the JSON profile and follow the Add Catalog wizard.

1. On the Data Studio menu of Database Actions, select the Catalog tool.
2. Select **Manage Catalogs**.
3. On the **Manage Catalogs** dialog, select **+Add**.
4. Mount Delta Share as catalog:

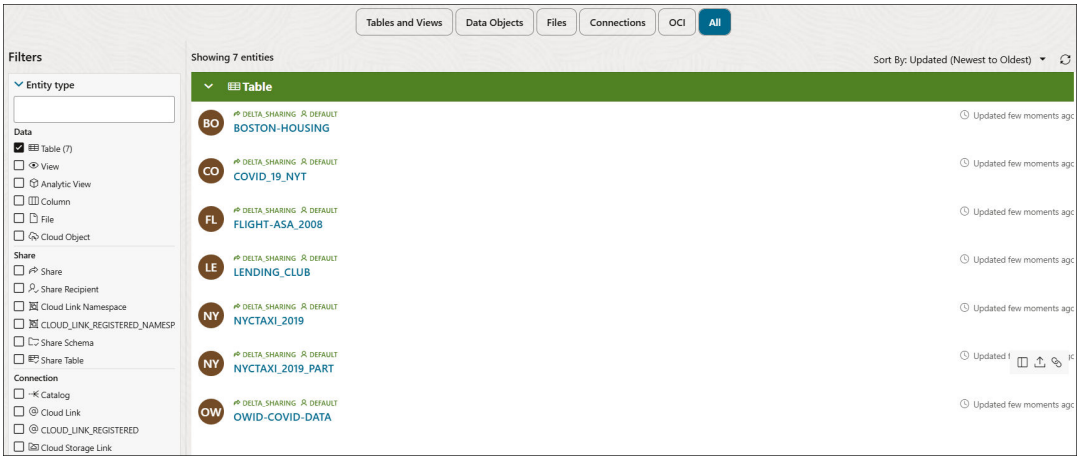
   Select **Data Share** as a Catalog Source on the Add Catalog dialog and click **Next**.
5. Select **Delta Share Provider JSON**, accept the **From File** default option, and click the **Delta Share Profile JSON** field.
   - Navigate to and select the **JSON profile** you downloaded in a previous step. Click **OPEN**.
   - Enter the **Provider Name**, under Share Provider Details and click **Next**.
6. Register shares made available to you.
   Move shares from **Available Shares** to **Selected Shares** and click **Next**.
7. Specify details of the Catalog:
   Enter the name of the Catalog and click **Create**.

   The Catalog page is displayed with a success message confirming that the Mounting Catalog for Shares was successful and that it has created the share provider.
8. View the newly created Delta share by filtering the results using the Select Catalogs…icon:
   - Click **Select Catalogs..**.
   - Select **Delta Sharing** from the **Catalogs and Schemas** wizard.

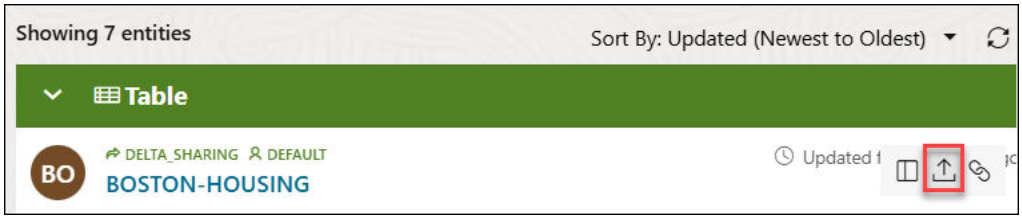


- Click **Apply** to apply the filter criteria in the search results.

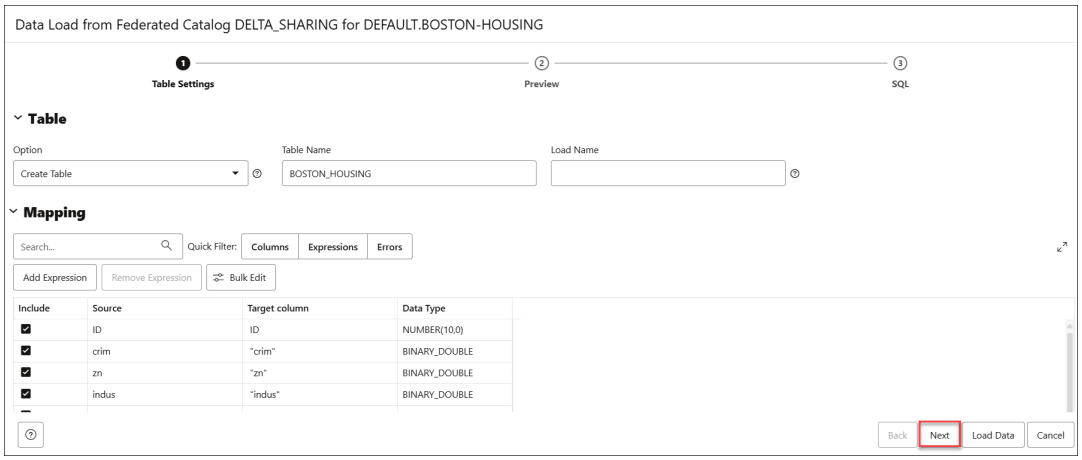
You will view all the tables listed from the Delta Share.



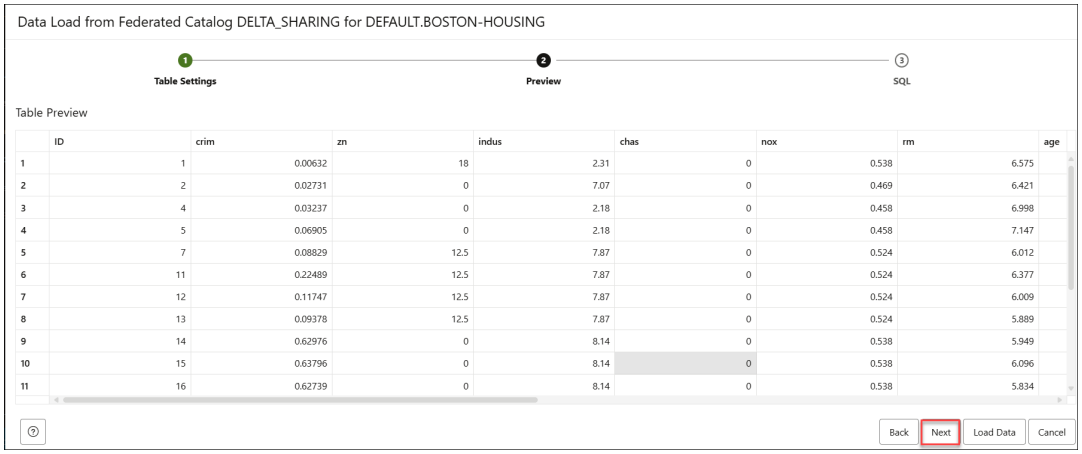
9. Select the table you want to load data from, then click **Load to table**.



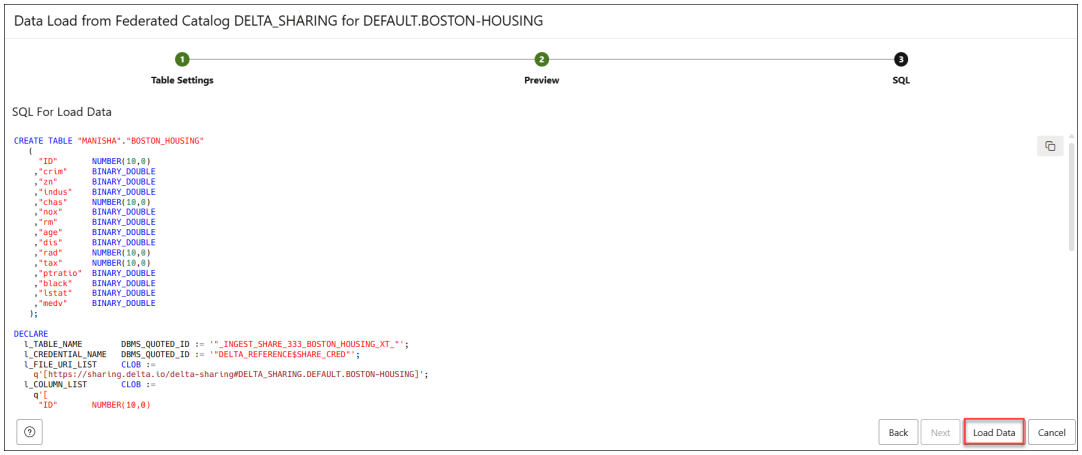
10. On the **Table Settings** of the Data Load dialog, you can specify processing options to specify settings for the data load job. You can also rename the newly created table from the catalog.



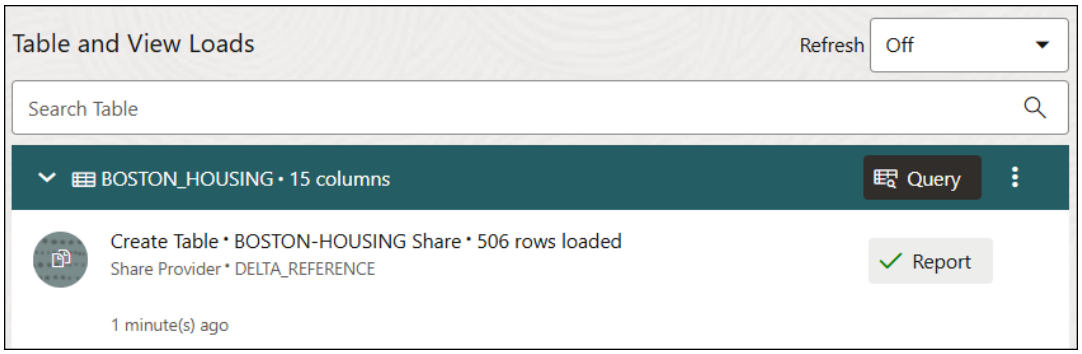
11. Click **Next** to view the data in the table.



12. Click **Next** to view the **Preview** tab to display the equivalent SQL code of the creating table job.



13. Select **Load Data**.
    A confirmation message is displayed that the data load job has started.
14. Navigate to the Data Load home page to view the progress of the data load job.



You have loaded data from Delta Share from the Catalog tool.

Query Remote Catalogs and Databases

You can access and query external data by mounting catalogs. This feature uses a simple SQL syntax that is similar to how you query data using database links.

The simplified SQL syntax makes it easy to discover and query external data without having to choose which external data to access, and without having to construct external tables or perform any other data integration tasks.

Refer to the Query External Data chapter for more information.

By querying remote catalogs and databases, you are accessing data directly from the external catalog in place, without having to move data into the Autonomous Database. This feature enables you to export data from the catalogs that the Autonomous Database can access without establishing multiple connections. You can query the Autonomous Database to see which external catalogs are available, what data is available for any given catalog, and you can query the data directly, from any SQL client. This involves providing a standard and simplified way to discover and query data sets connected to the Autonomous Database, leveraging the catalog's ability to present external data as standard Oracle schemas and tables.

Refer to the Manage Catalogs chapter for information on catalogs supported across various categories, including Databases, External Data Catalogs, Shares, and more.

Some of the key features are:

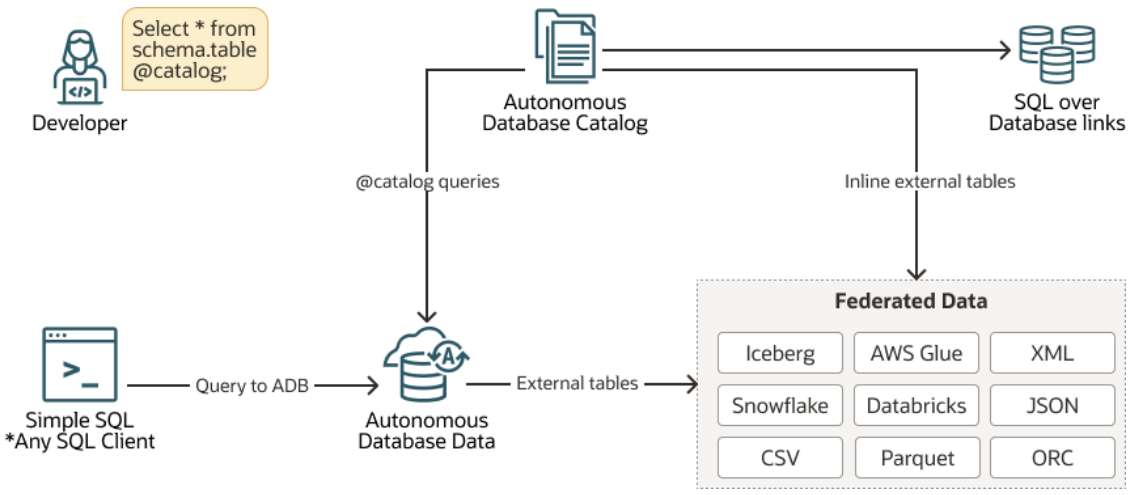
- **Discover External Data**: Using SQL, you can identify external catalogs along with the tables they contain. This will help you in understanding the available data structure.
- **Simplify Table Discovery**: Use existing query syntax and familiar views to easily list tables within an external catalog.
- **Integrate with data tools:** You can integrate the query data with Data Studio suite of tools to make informed decisions and gain deeper insights.

Topics:

- Conceptual Diagram of Querying Remote Catalogs
  The following diagram displays the conceptual network diagram of querying remote catalogs using Autonomous Database:
- Concepts Related to Querying Remote Catalogs and Databases
  An understanding of the following concepts is necessary for querying with Remote Catalogs:
- About Querying Remote Catalogs and Databases
  This topic describes how to query external data using the catalog metadata using an SQL syntax of the form: `[schema].object@catalog_name`.
- Examples: List the Mounted Catalogs
  You will use this syntax in a SQL query.
- Examples: List the Tables in a Catalog, Including Their Owning Schemas
  To list all tables within a specific schema (catalog) from the enabled catalogs, you can use:
- Examples: Query Data from a Table in a Catalog
  To query data from a table in a specific catalog (schema), you must prepend the schema name to the table name.
- Examples: Query a Table from a Delta Share and Analyze the View
  Let's assume that we are working on an inventory database where we have data about the information collected by the U.S. Census Service concerning housing in the Boston area, and we need a query that returns the maximum value of the full-value property tax.

Conceptual Diagram of Querying Remote Catalogs

The following diagram displays the conceptual network diagram of querying remote catalogs using Autonomous Database:



Data from all types of external catalog can be queried this way including:

- **Database Links**: Data in other databases where a catalog is mounted over a database link.

  These catalogs are built using traditional database links, which are set up in advance by system administrators. The catalog tool autonomous database uses standard SQL queries on these links to discover tables and other objects.

  There are two types of database links: Oracle and Heterogeneous.

  For Oracle links, use queries like `SELECT table_name FROM all_tables@link`.

  For Heterogeneous links, similar queries are run, such as `SELECT table_name FROM all_tables@link`.
- **Glue Catalogs**: Data made available through AWS Glue, where a catalog is mounted over an AWS Glue region.
- **OCI Data Catalogs**: Data made available via OCI Data Catalogs. These catalogs use customer-defined credential objects to securely connect to OCI endpoints protected by ACLs.
- **Delta Share Catalogs**: Data that has been shared using the Delta Sharing standard from any system that can use this standard to share data, where a catalog is mounted over the share. These catalogs leverage External Table Delta Sharing support, which employs customer-defined credential objects to access delta sharing endpoints. The credential includes a bearer token for authentication.
- **Iceberg Catalogs**: Data made available through an Iceberg Data Catalog, such as Databricks Unity Iceberg Catalog, or Snowflake Polaris.

Concepts Related to Querying Remote Catalogs and Databases

An understanding of the following concepts is necessary for querying with Remote Catalogs:

- **Catalog**: A catalog is a metadata layer that organizes and describes data assets such as schemas and tables. It acts as a bridge between the database and external data sources. A catalog contains one or more schemas, and each schema can have multiple tables. These schemas and tables represent the logical organization of data.
- **External Catalogs**: they are catalogs that stores metadata for systems other than the locally connected ADB, including other ADBs, remote databases, shares, external data catalogs.
- **DB Links**: DB Links are a mechanism to connect to a remote database from an Oracle database. They allow you to access tables, views, and other objects in a remote database as if they were local objects. They enable querying data in remote databases. This feature aims to treat DB Links as a type of "catalog." The data dictionary of a remote database connected via a DB Link can be considered a catalog.
- **Mounting a Catalog**: A **mounted catalog** is one where the Autonomous Database has an active connection to the catalog, allowing the database to access and query its datasets. Oracle Autonomous Database supports mounting catalogs from external data sources like AWS Glue or Databricks.

About Querying Remote Catalogs and Databases

This topic describes how to query external data using the catalog metadata using an SQL syntax of the form: `[schema].object@catalog_name`.

Here the syntax `[schema].object@catalog_name` refers to a representation of accessing database objects

The table describes the parameters and their descriptions used in the syntax:

Parameters	Description
Schema	This field is optional, and it defaults to the current user. A schema is a namespace within a catalog that organizes related database objects such as tables, views, procedures, or functions. It acts as a logical container for these objects.
Object	This field refers to any database object contained within a schema, such as tables, views, procedures, or functions.
Catalog	The catalog field displays the name of the catalog with the referenced table. You can create and manage catalogs using the `DBMS_CATALOG` package or the Catalog Tool of the Data Studio suite of tools.

> **Note:**
>
> You can replace catalogs with **dbLinks** as well.

Accessing Remote Tables

To access a table from a remote schema using a catalog, you must follow these steps:

1. **Mount a Catalog**: Begin by mounting a catalog over any supported type of data store (e.g., Iceberg, Delta Lake, etc.). Mounting a catalog standardizes the metadata, enabling SQL access to external data as if it were tables within a schema in the catalog. Refer to Manage Catalogs for detailed steps on mounting catalogs. Once the catalog is mounted and metadata is standardized, you can directly query external data using SQL. Even if the external data does not physically exist as traditional tables, the catalog abstracts it into a tabular format for easy querying.
2. **Reference the table**: Once mounted, you can reference the remote table by specifying the schema and catalog (if applicable) followed by the table name using the following syntax: `[remote_schema].table@catalog`.

Here the `schema` in this identifier is the remote schema that contains the remote table. `Table`, in this case, is the table name in the remote catalog.

> **Note:**
>
> The user must have the appropriate credentials to access the specific table and catalog.

Key Points to Consider

- The feature is designed using the `DBMS_CATALOG` package installed in an ADB-S environment.
- DDL operations (e.g., `CREATE, ALTER, DROP`) are not allowed on remote databases via catalogs. If you do so, you may encounter an *"ORA-02021: DDL operations are not allowed on a remote database"* error.
- **Same name Restriction**: This feature enforces that for a given user, there cannot exist a DB link and a catalog with the same name.

In the subsequent chapters, you will learn how to:

- Display the list of Mounted Catalogs
- Display the list of tables in a Catalog, Including Their Owning Schemas
- Query data from a Table in a Catalog
- Query data using DBLinks

You will use SQL Worksheet Editor of Database Actions to run SQL queries. You must first connect to your Database Actions instance and select **SQL** under **Development** menu.

Prerequisites:

You will require an OCI account.

How are DBlinks different from Catalogs?

DBLinks: They are primarily used to connect and query another database directly using a syntax.

For example:

```
SELECT * FROM <schema_name>.<table_name>@<dblink_name>;
```

When a catalog is mounted over another database via a DBLink, there is no need to separately define metadata in the catalog because the DBLink syntax directly accesses the remote database objects. Essentially, the DBLink acts as a transparent connection to the remote database, allowing queries to run as if the data were local.

Catalogs: They serve as a metadata layer that can be mounted over different types of data stores, not only databases connected via DBLinks but also external data catalogs or data shares like Delta Share or Live Share providers. When a catalog is mounted over a non-database store, the `@catalog` syntax adapts queries to generate inline external tables, enabling access to external data without requiring a DBLink. This means catalogs abstract the metadata management and provide a unified interface for querying diverse data sources, including external ones.

Examples: List the Mounted Catalogs

You will use this syntax in a SQL query.

For example, to list the catalogs you mount, you could run the following command:

```
SELECT CATALOG_NAME,CATALOG_TYPE,IS_ENABLED FROM USER_MOUNTED_CATALOGS;
```

The output of the query lists the catalogs that you connect with.

After the connection, you can query the catalogs using the syntax mentioned above:

```
SELECT * FROM "DEFAULT"."BOSTON-HOUSING"@CATALOGNAME;
```

In the above example, *Catalogname* is a catalog you have mounted, and *BOSTON-HOUSING* is a table in that catalog.

Examples: List the Tables in a Catalog, Including Their Owning Schemas

To list all tables within a specific schema (catalog) from the enabled catalogs, you can use:

```
SELECT OWNER, TABLE_NAME FROM ALL_TABLES@CATALOGNAME;
WHERE OWNER = '<SCHEMA_NAME>';
```

Here, *ALL_TABLES* is the table name.

For example, to list tables in the `SH` schema (a sample schema available in Oracle Autonomous Database):

```
SELECT OWNER, TABLE_NAME FROM ALL_TABLES@CATALOGNAME;
WHERE OWNER = 'SH';
```

This query returns all tables owned by the specified schema.

> **Note:**
>
> In this example, you are listing tables in the `SH` schema of an external database.

Examples: Query Data from a Table in a Catalog

To query data from a table in a specific catalog (schema), you must prepend the schema name to the table name.

For example, querying data from the `SALES` table in the SH schema:

```
SELECT * FROM "SH"."SALES"@catalogname;
```

> **Note:**
>
> In this example, you are querying data in the `SH` schema of an external database.

If you are working with other schemas like `SSB`, you can similarly query their tables by specifying the schema name:

```
SELECT * FROM "SSB"."SALES"@catalogname;
```

> **Note:**
>
> In this example, you are querying data in the `SSB` schema of an external database.

Examples: Query a Table from a Delta Share and Analyze the View

Let's assume that we are working on an inventory database where we have data about the information collected by the U.S. Census Service concerning housing in the Boston area, and we need a query that returns the maximum value of the full-value property tax.

1. **Create an Oracle Delta Share provider and credential**:
   For creating a Delta Share provider, you typically need to interact with the Delta Sharing platform, which involves creating a credential suitable for use with delta share providers. In this example, you can create `DATABRICKS_PROVIDER`.

```
BEGIN
      dbms_share.create_or_replace_share_provider(
             provider_name=>   'DATABRICKS_PROVIDER',
             endpoint=>        'https://sharing.delta.io/delta-sharing');
      dbms_cloud.create_credential(
             credential_name => 'CREDENTIAL_NAME',
             username =>        'BEARER_TOKEN',
             password =>        '**********************************');
      dbms_share.set_share_provider_credential(
             'DATABRICKS_PROVIDER', 'CREDENTIAL_NAME');
   END;
   //
```

   The output of this procedure successfully creates a delta share provider.

2. **Mount a DB Share Catalog**:
   After creating credentials, you can mount the registered share provider from step 1 as a catalog to enable the catalog users to search and discover data from the share, and to query that data using SQL.

```
BEGIN
      dbms_catalog.mount_share(
                               catalog_name=>'databricks',
                               share_provider=>'databricks_provider',
                               share_name=>'delta_sharing');
      END;
    /
```

   The output of this procedure successfully creates a catalog integration that enables the Catalog tool to query tables stored in Databricks.

3. **Run a simple query**:
   You can now run a simple query on the table:

```
select max(tax) from "DEFAULT"."BOSTON-HOUSING"@catalogname;
```

   The output returns the maximum value of the full value of the property tax rate present in the table.

```
MAX(TAX)
----------
711
```

Viewing Entity Details

When you have found an entity you are interested in, you can click the entity to view its details, or click one of the other available actions from the menu that appears on the right hand side when you hover over that row.

The way you select the entity details can be set using the **Customize Themes** tab in the Catalog Settings menu.

To view details about an entity, click the **Actions** icon at the top right of the Table and View Loads header entry, then click **Table** and select **View Details**.

For all entities, the details include Lineage and Impact sections. The inclusion of other details, such as Preview and Statistics, varies by entity type.

Preview

Preview displays the data of the entity. For a table, the Preview displays the columns of the table and the data in those columns. You can sort the data in the column into ascending or descending order by clicking the up or down arrow to the right of the column name.

Describe

For an analytic view, the Describe tab has information about the entity, and displays the hierarchies, levels, level depth, dimension tables, level columns, and number of distinct values for a level.

Lineage

Lineage displays all known information about the upstream dependencies of the entity, and therefore how the entity was created and how it is linked to other entities.

For example, for a table that you created in your database, the lineage is just the table. For a table that you created by loading a CSV file from cloud storage, the lineage includes the ingest directive for the data load and the CSV file that is the source of the data.

Pointing to the name of an item in the lineage displays the table name, the application that created it, the type of entity, the path to it, and the schema it is in.

Arrows point from an entity to the entity that it derives from. For example, for a table created in a data load job, an arrow points from the table to the ingest job and an arrow points from the ingest job to the CSV file. If you point to an arrow, then a Links Information box appears that shows information about the relationship between the two entities.

To view more details about an item, click the Actions icon for the item, then click **Expand**. For a table, the columns of the table are displayed. Pointing to the name of the table or of a column displays the name, application, type, path, and schema of the table or column. To collapse the display, click the Actions icon, then click **Collapse**.

You can increase or decrease the size of the displayed objects by using the + (plus) and - (minus) keys. You can reposition the objects by grabbing a blank spot in the display and dragging vertically or horizontally.

The lineage for some entities, such as analytic views, is more complex. An analytic view entity deployed for a business model has links to columns in a fact table and to hierarchies. The fact table has links to attribute dimensions and, for a data load job, to an ingest directive for the job. The ingest directive has a link to the source file. The attribute dimensions have links to tables for the dimensions. Those tables have links to ingest directives that have link to source files.

Impact

Impact shows all known information about the downstream use of an entity, and therefore how a change in the definition of an entity may affect other entities that depend on it. For example, if a table is used in an analytic view, a change to one of the column definitions in the table may affect the mapping from that column to the analytic view.

Classifications

For an analytic view and its attribute dimensions, hierarchies, and measures, the Classifications tab displays classifications and their values. Classifications are metadata that applications can use to present information about analytic views. When creating a business model, you may specify the values for the Caption and Description classifications.

Optimize

For an analytic view, the Optimize tab has information about caches created for the analytic view. A cache may exist if the advanced option Enable Autonomous Aggregate Cache was selected for the business model for which the analytic view is deployed.

Statistics

Statistics display information about the entity. For example, the statistics for a table include the size of the table and the numbers of rows and columns. They also include the names of the columns, their data types, the number of distinct values and null values, the maximum and minimum values, and other information.

The data is represented in the form of histogram which is column statistic which provides more detailed information about data distribution in a table's columns.

The histograms in the statistics pane can be representative of the following types:

- Frequency: In a frequency histogram, each distinct column value corresponds to a single bucket of the histogram. Since each value has its own dedicated bucket, some buckets may have many values, whereas others have few.
- Top-frequency: A top frequency histogram is a variation on a frequency histogram that ignores non-popular values that are statistically insignificant
- Height-Balanced: In this histogram, column values are divided into buckets so that each bucket contains approximately the same number of rows.
- Hybrid: A hybrid histogram combines characteristics of both height-based histograms and frequency histograms. This approach enables the optimizer to obtain better selectivity estimates in some situations.

Refer to Gathering Statistics for Tables section for more details.

Job Report

Displays a report of the total rows loaded and failed for a specific table.

The report displays of the total rows loaded and failed.

You can view the name of the table, the time the table was loaded and the time taken to process the load.

Data Definition

Data Definition displays the Oracle Autonomous Database DDL that created the entity.

You can view the following **Entity Types** from the drop-down:

- All
- Table
- View
- Analytic View
- Table,View,Analytic View
- Analytic Dashboard
- Attribute (Attribute Dimension)
- Attribute (Hierarchy)
- Attribute Dimension
- Cloud Link (Cloud Link Namespace)
- Cloud Link Namespace
- Cloud Object (Cloud Storage Link)
- Cloud Object (Data Catalog Asset)
- Cloud Object (Share Link)
- Cloud Storage Link
- Cloud Virtual Object (Cloud Storage Link)
- Column (Table)
- Column (View)
- Data Catalog Asset
- Data Catalog Link
- Database Link
- Directory
- External Location
- File (Directory)
- Function
- Hierarchy
- Ingest Job
- Level (Attribute Dimension)
- Level (Hierarchy)
- Live Table Feed
- Measure (Analytic View)
- Mining Model
- Package
- Procedure
- Schema
- Share
- Share Link

- Share Provider
- Share Recipient
- Share Schema (Share)
- Share Table (Share Schema)

Table Entity Details
Click the **Actions** icon next to **Table** entity to perform the following operations:

- Select **View Details** to view details about the table.
- Select Gather Statistics to display new statistics after you modify the table's structure. Refer to the Gathering Statistics for Tables for more detailed information. You can also gather statistics using the **Actions** drop-down field while viewing the table details in the Table Entity wizard.
- Select **Add to Share** to share the table. Refer to Provide Share for detailed information. You can also share the selected table using the **Actions** drop-down field while viewing the table details in the Table Entity wizard.
- Select **Register to Cloud Link** to register the table for remote access for a selected audience you define. See Registering Cloud Links to Access Data section for more information. You can also register the selected table using the **Actions** drop-down field while viewing the table details in the Table Entity wizard.
- Select **Create Analytic View** to create an Analytic View from the selected associated table. See Creating Analytic Views. You can also create an Analytic View from selected table using the **Actions** drop-down field while viewing the table details in the Table Entity wizard.
- Select **Query** to analyze the selected table in the Data Analysis tool. Refer to The Data Analysis Tool for more information. You can also analyze the selected table using the **Actions** drop-down field while viewing the table details in the Table Entity wizard.
- Select Export Data to Cloud to export data to a cloud object store. You can also export data from the **Actions** drop-down field while viewing the table details in the Table Entity wizard.
- Select **Edit** to edit the properties of the table. Refer to Editing Tables for more information. You can also edit properties of the selected table from the **Actions** drop-down field while viewing the table details in the Table Entity wizard.
- Select **Drop** to delete the table.

View Entity Details

Click the **Actions** icon next to **View** entity to perform the following operations:

- Select **View Details** to view details about the view.
- Select **Add to Share** to share the table. Refer to Provide Share for detailed information.
- Select **Create Analytic View** to create an Analytic View from the selected associated table. See Creating Analytic Views.
- Select **Query** to analyze the selected table in the Data Analysis tool. Refer to The Data Analysis Tool for more information.

Share Entity Details

Click the **Actions** icon next to **Share** entity to perform the following operations:

- Select **View Details** to view details about the share.

- Select **Objects** to view the **Select Tables** dialog of the specific share where you can remove tables from the list of Shared Tables list or add tables from the list of Available Tables list to the list of Shared Tables.
- Select **Recipients and Profiles** to view the **Recipients and Profiles** dialog which enables you to create a new share recipient. You can also revoke the recipient right from the share by clicking X next to the recipient name. The Recipients and Profiles dialog enables you to update the list of recipients, copy profile activation link, send an activation mail to the selected recipient.
- Select **Manage Versions** to
- Select **Jobs** to view the **Jobs** table and see information related to it. A job is created whenever something needs to happen to a share. This includes the following:
  - PUBLISH VERSIONED SHARE
  - PUBLISH LIVE SHARE
  - DROP SHARE VERSIONS
- Select **Scheduling** to
- Select **Publish**
- Select **Unpublish**
- Select **Delete**

Share Recipient Entity Details

Click the **Actions** icon next to the **Share Recipient** entity to view the share recipient entity details.

- Select **Granted Shares** to open **Manage Granted Shares** dialog where you can grant access to a share from list of available shares or you can revoke access from the current selection of share..
- Select **Get Bearer Token** to open an Auth Token dialog to authenticate the recipient. Use this option when your token expires. This enables you to generate a new token. A Bearer Token is part of the JSON file.
- Select **Copy Profile Activation Link** to copy the JSON profile activation link to clipboard.
- Select **Rename** to rename the Recipient.
- Select **Delete** to delete the recipient you created entirely. By doing so, access to all shares will end instantly and any existing delta credentials will become invalid.

Cloud Storage Link Entity Details

Click the **Actions** icon next to **Cloud Storage Link** entity to perform the following operations:

- Select **View Details** to view details about the table.
- Select **Objects** to view objects available in the selected storage link. You can click a file on the navigator pane to view it on the display area.
- Selecting **Link Tables** opens the Link Data page on the Data Load tool with the selected cloud storage link on the Cloud Location URL field. You can link data present in the cloud storage to the Autonomous Database. See Linking to Objects in Cloud Storage.
- Selecting **Load Tables** opens the Load Data page on the Data Load tool with the selected cloud storage link on the Cloud Location URL field. You can load data present in the cloud storage to the Autonomous Database. See Loading Data from Cloud Storage.

- Selecting **Create Live Table Feed** opens the Create Live Table feed wizard with he selected cloud storage link on the Cloud Location URL field. See Feeding Data to view more details.
- Select **Edit** to update any details on the cloud storage location. See Managing Connections to view details on creation of cloud storage location.
- Select **Rename** to rename the cloud store location to a different name.
- Select **Delete** to delete the cloud store location.

Catalog Entity Details

Click the **Actions** icon next to **Catalog** entity to perform the following operations:

- Select **View Details** to view more details about the catalog.
- Select **Disable** to disable the catalog.
- Select **Delete** to delete the catalog.

Data Catalog Link Entity Details

Click the **Actions** icon next to **Data Catalog Link** entity to perform the following operations:

- Select **View Details** to view more details about the data catalog link.
- Select **Edit** to edit the data catalog.
- Select **Rename** to rename the data catalog.
- Select **Delete** to delete the data catalog.

Database Link Entity Details

Click the **Actions** icon next to **Database Link** entity to perform the following operations:

- Select **View Details** to view more details about the database link.
- Select **Load to table** to load the database link to a table.
- Select **Link to table** to link the database link to a table.
- Select **Mount as catalog** to mount the database link as catalog.
- Select **Delete** to delete the database link entity.

Live Table Feed Entity Details

Click the **Actions** icon next to **Live Table Feed** entity to perform the following operations:

- Select **View Details** to view more details about the live table feed.
- Select **Edit Live Table Feed** to view the Edit Live Table Feed wizard where you can edit the live table.
- Select **Delete Live Table Feed** to delete the Live Table Feed.
- Select **Run Live Table Feed Immediately(Once)** to run the live table feed.

Share Provider Entity Details

Click the **Actions** icon next to **Share Provider** entity to perform the following operations:

- Select **View Details** to view the share provider entity details..
- Select **Manage Shares** for Share provider dialog box. This lists the shares you select to share with the recipient. You can edit the list of shares you wish to share with the recipient.

- Select **Rename** to renames the Provider Name.
- Select **Delete** to remove the Share Provider Entity.
- Select **Load to table** to view the Load Data page with the Share tab selected.
- Select **Link to table** to view the Link Data page with the Share tab selected to view and run the related link object. Drag and drop the shared data to add it to the data link job.
- Select **Copy Endpoint** to copy the share endpoint.

Bucket Entity Details

Click the **Actions** icon next to **Bucket** entity to perform the following operations:

- Select **View Details** to view details of the bucket entity.
- Select **Create Cloud Storage Link** to create cloud storage link.

Compartment Entity Details

Click the **Actions** icon next to **Compartment** entity to perform the following operations:

- Select **View Details** to view details of the compartment entity.
- Select **Create Share Recipient** to create a share recipient.

> **Note:**
>
> **Some catalog entity types do not have details screens**
>
> There are certain items in the catalog that do not have detail screens yet, thus clicking on them might not provide more details.

- Autonomous Database Entity Details
  You can mount additional catalogs from other Autonomous Databases in the same tenancy.

Autonomous Database Entity Details

You can mount additional catalogs from other Autonomous Databases in the same tenancy.

Before selecting whether or not to mount a catalog, you can browse the databases using the catalog to get additional information about each one.

Click **All** filter from the top and select **Autonomous Database** from OCI faceted filter panel in the left.

Selecting the above options lists all the autonomous databases in your tenancy.

You can view details about the database and launch the following applications from the **Database actions** drop-down field in the **Autonomous Database Entity** dialog box:

- Database Transforms
- Graph Studio
- APEX
- Machine learning notebook

- Machine learning user management
- Mongo Database
- ORDS
- SQL Developer

You can hover over the autonomous database whose details you want to view.

Click **Actions** next to the Autonomous Database entity and select **View Details** to view the following details:

General

The General tab displays the tenancy details, backup and recovery information, and network and maintenance details of the selected autonomous database.

Under **General Information**, you can view:

- **Database Name**: You set the database name when you provision a database or when you rename a database.
- **Workload Type**: The Autonomous Database supports different workload types, including: Data Warehouse, Transaction Processing, JSON Database, and APEX Service.
- **Compartment**: When creating a compartment, you must provide a name for it (maximum 100 characters, including letters, numbers, periods, hyphens, and underscores) that is unique within its parent compartment. Oracle will also assign the compartment a unique ID called an Oracle Cloud ID.
- **OCID**: It displays the Oracle Cloud ID. See Resource Identifiers for more information.
- **Created**: Displays the date when the autonomous database was created.
- **License type**: The License type field shows on the Autonomous Database Information tab when you enable the Bring Your Own License (BYOL) option. This field shows your license and Oracle Database Edition.
- **Database Version**: Specifies the version of the Database.
- **Lifecycle State**: You can view availability information for an autonomous database instance in this field.
- **Instance type**: You can view the instance details of the autonomous database.
- **Character Set**: The Autonomous Database default database character set is Unicode `AL32UTF8` and the default national character set is `AL16UTF16`. When you provision a database, depending on the workload type, you can select a database character set and a national character set.
- **Mode**: You can select any of the following autonomous database operation modes:
  - Read/Write
  - Read-Only
  - Restricted

  The default mode is Read/Write.

Under **Resources** you can view:

- **Compute**: This field displays the autonomous database instance ECPU count (OCPU count if your database uses OCPUs).

Under **APEX** you can view, **instance name** that displays the autonomous database instance you use to create an APEX Service instance using an OCI.

Under **Data Safe** you can view, **Status** field that displays if you have registered the autonomous database with Oracle Data Safe. Oracle Data Safe, which is included with Autonomous Database, provides a unified control center that helps you manage the day-to-day security and compliance requirements of Oracle Databases.

Under **Disaster Recovery**, you can view:

- **Role**: Each autonomous database has a designated role: primary, standby, or snapshot standby.
- **Local**: If there is a local peer database, Autonomous Data Guard standby, this field shows the local standby's availability domain. See Use Standby Databases with Autonomous Data Guard for Disaster Recovery for more information.
- **Cross-region**: Displays if cross-region Autonomous Data Guard is enabled by provisioning a remote standby instance.

Under **Backup**, you can view:

- **Automatic backup retention period**: The backup retention period specifies how many days automatic backups are retained and the period for which you are billed for the automatic backups that are retained. If you change the backup retention period to a shorter period and automatic backups exist that are older, you are no longer billed for the automatic backups with timestamps beyond the specified retention period.
- **Total backup storage**: Displays the total storage being billed, including for automatic backups and if there are long-term backups, this also includes the long term backup storage.
- **Last Automatic backup**: Shows the timestamp of the last automatic backup.
- **Next long term backup**: When you configure scheduled long-term backups, this field displays the timestamp for the next long-term backup.
- **Long term backup schedule**: Displays the long-term backup schedule, with links to edit or delete a long-term backup schedule.

Under **Network**, you can view:

- **Access type**: Displays the specified access type for the Autonomous Database instance.
- **Access control list**: This field displays if Access Control Lists (ACLs) of an Autonomous Database Instance are defined. It can display the following two values:
  - **Enabled**: The field displays Enabled when ACLs of the Autonomous Database Instance are defined.
  - **Disabled**: The field displays Disabled when ACLs of the Autonomous Database Instance are not defined.
- **Mutual TLS(mTLS) authentication**: Displays if your Autonomous Database instance is configured to only allow mTLS connections.

Under **Maintenance**, you can view:

- **Patch level**:You can view Autonomous Database patch information.
- **Next maintenance**: This field in the Oracle Cloud Infrastructure Console reflects a maintenance window date and time based on the patch level.

- **Customer contacts**: When customer contacts are set, Oracle sends notifications to the specified email addresses for Autonomous Database service-related issues.

Under **Encryption**, you can view the **Encryption key** field that displays the master encryption key the Oracle Autonomous Database creates.

Lineage

Lineage displays all known information about the upstream dependencies of the entity, and therefore how the entity was created and how it is linked to other entities.

For example, for an autonomous database that you create, the lineage is just the cloud. For an autonomous database you create from cloud storage, the lineage includes the ingest directive for the load and the cloud object store path that is the source of the database.

Pointing to the name of an item in the lineage displays the database name, the application that created it, the type of entity, the path to it, and the schema it is in.

Arrows point from an entity to the entity that it derives from.

You can increase or decrease the size of the displayed objects by using the + (plus) and - (minus) keys. You can reposition the objects by grabbing a blank spot in the display and dragging vertically or horizontally.

Impact

Impact shows all known information about the downstream use of an entity, and therefore how a change in the definition of an entity may affect other entities that depend on it.

Click **Actions** next to the Autonomous Database entity and select **+ Create Database Link** to create database links from the Autonomous Database. See Load Data from Oracle and Non-Oracle Databases using Database Links.

Click **Actions** next to the Autonomous Database entity and select **Create Share Recipient** to create share recipient from the Autonomous Database. See Create Share Recipient to create a share recipient from the Autonomous Database.

Registering Cloud Links to Access Data

Cloud Links enable you to remotely access read only data on an Autonomous Database instance.

You can register a table for remote access for a selected audience. Scope indicates who can remotely access the data. Scope can be set to various levels, including to the region where the database resides, to individual tenancies, or to compartments. You can provide a namespace and a name other than the original schema and object names. For example, you can register a table under the namespace FOREST and for security purposes or for naming convenience, you can provide a namespace and a name other than the original schema and object names.

To register a table to cloud link, click Actions next to the Table entity and select **Register to Cloud Link**.

Specify the following field values:

- **Cloud Link Namespace**: Enter a namespace to register the table. For example, ADMIN.
- **Cloud Link Name**: Enter a name for the cloud link. For example, Sales.
- **Description**: Specify the description of your cloud link.
- **Scope**: Specifies who and from where a user is allowed to access the registered table.

You can choose from any of the following available options:

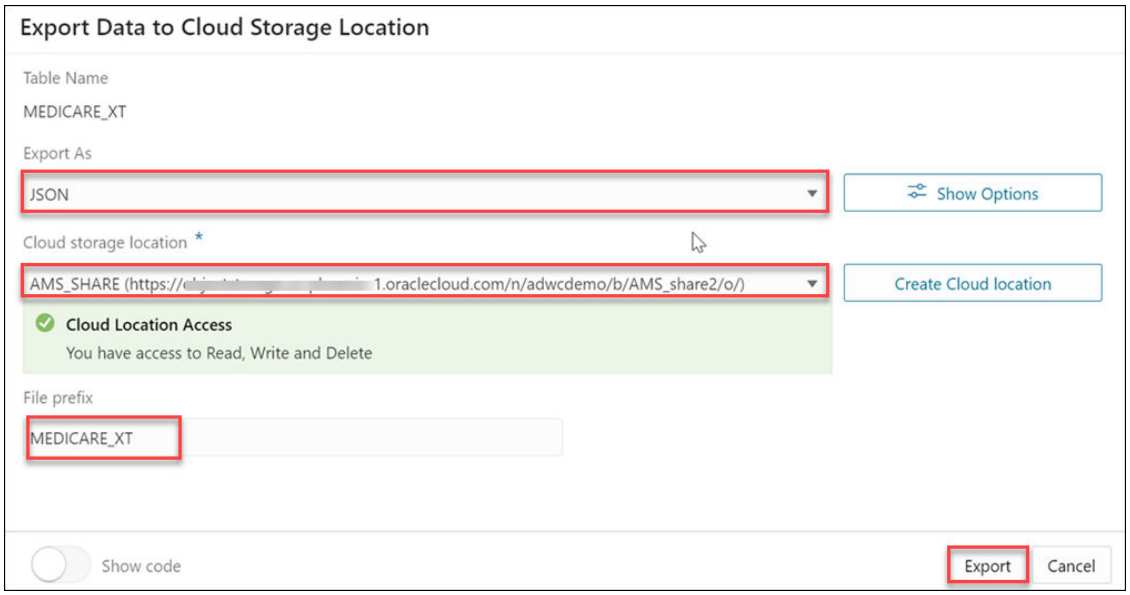
- **MY$REGION**: You can grant remote data access to other tenancies in the region of the Autonomous Database instance that is registering the data set. This is the least restrictive scope.
- **MY$TENANCY**: You can grant remote data access to any resource, tenancy, compartment, or database in the tenancy of the Autonomous Database instance that is registering the data set. This scope is more restrictive than MY$REGION scope.
- **MY$COMPARTMENT**: You can grant remote data access to any resource, compartment, or database in the compartment of the Autonomous Database instance that is registering the data set.
- **OCID**: Access to the data set is allowed for the specific Autonomous Database instances identified by OCID.

Click **OK** to finish the registration of the selected table with the cloud link.

Export Data to Cloud

Use the **Export Data to Cloud** menu in the table actions of the Browse Catalog page to export data as text from an Autonomous Database to a cloud Object Store. The text format export options are CSV, JSON, Parquet, or XML.

To export data to a cloud storage location, click **Actions** next to the Table entity on the Browse Catalog page and select **Export Data to Cloud**.



The Export Data to Cloud Storage Location displays the following fields:

- **Table Name**: This field displays the table name.
- **Export As**: This drop-down displays the list of format options you wish to export the data. You can select CSV, JSON, Parquet, or XML. When you export the data as **CSV**, and click **Show Options**, you can:
  - Set the maximum file size value using the slider. The maximum file size you can export is from 10MB to 1GB.
  - Select the **header** option to write column names as the first line in output files.

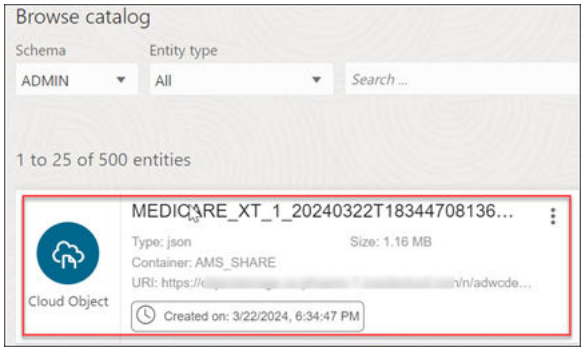
  - Select the **Escape** option to specify the occurrence of the quote character in the field value using the “\” character.
  - **Text Enclosure and Field delimiter**: These options are visible only when the selected file is in CSV format. Select or enter the character used in the source file for text enclosure and field delimiter.
  - **Compression**: Select **No Compression** to disable compression or **GZIP** to enable GZIP compression for selected file.
- **Cloud storage Location**: Select the uri value to the URL for an existing bucket on your Cloud Object Storage from the drop-down. If you do not have a cloud storage location, select **Create Cloud Location** to create a cloud storage location to export your data. For more details, see Create Credentials.

> **Note:**
>
> The Cloud storage location is highlighted in red if the tool doesn’t accept the selected cloud location. In this case, you will receive a Cloud Location Access Error that says “`You do not have access to Read, Write and Delete`”.

- **File prefix**: Specify the File prefix. This will be the prefix of the file you export.

Click **Export**.

After the file export is complete, you will view the newly created cloud object in the entity list of the Browse Catalog page.



Gathering Statistics for Tables

You can generate statistics that measure the data distribution and storage characteristics of tables.

You must gather statistics periodically for tables where the statistics become stale over time because of changing data volumes or changes in column values. The Gather Statistics icon gathers new statistics after the table’s structure are modified in ways that make the previous statistics inaccurate. For example, after loading a significant number of rows into a table, collect new statistics on the number of rows. After updating data in a table, you do not need to collect new statistics on the number of rows, but you might need new statistics on the average row length.

Table Statistics also include column statistics. The statistics you generate include the following:

Table statistics:

- **Table Size**: Specifies the size of the table in bytes.
- **Number of rows**: Displays the number of rows in the table.
- **Number of columns**: Displays the number of columns in the table
- **Compressed size**: Displays the size of compressed table in bytes.
- **Statistics gathered on**: Displays date and time of last statistics gathered.

Column Statistics

- Number of distinct values (NDV) in column
- Number of nulls in column
- Data distribution (histogram)

The above information is displayed in the statistics details of tables.

Editing Tables

You can create and edit objects using Edit Table wizard available from the **Edit** menu in **Actions** (three vertical dots) besides the table entity.

Clicking **Edit** from the Actions menu opens the **Edit Table** wizard. You can visit the panes in any order to edit a table. The table properties are grouped in several panes.

- **Schema**: Database schema in which the table exists.
- **Name**: Name of the table.

The different panes in the dialog are described in the following sections:

- Columns Pane
- Primary Key Pane
- Unique Keys Pane
- Indexes Pane
- Foreign Keys Pane
- Table Constraints Pane
- Comments Pane
- Storage Pane
- External Table Properties Pane
- Materialized View Pane
- DDL Pane
- Output Pane

- Columns Pane
  Specifies properties for each column in the table.
- Primary Key Pane
  Specifies the primary key for the table.
- Unique Keys Pane
  Specifies one or more unique constraints for the table.
- Indexes Pane
  Lists the indexes defined for the table.

- Foreign Keys Pane
  Specifies one or more foreign keys for the table.
- Table Constraints Pane
  Specifies one or more check constraints for the table.
- Comments Pane
  Enter descriptive comments in this pane.
- Storage Pane
  Enables you to specify storage options for the table.
- External Table Properties Pane
  Specifies options for an external table.
- Materialized View Pane
  Specifies options for a materialized view.
- DDL Pane
  You can review and save the SQL statements that are generated when creating or editing the object. If you want to make any changes, go back to the relevant panes and make the changes there.
- Output Pane
  Displays the results of the DDL commands.

Columns Pane

Specifies properties for each column in the table.

General tab

Lists the columns available in the table.To add a column, click **Add Column (+)**. A new row is added to the table below. Select the row and enter the details for the column. To delete a column, select the row and click **Remove Column (-)**. To move a column up or down in the table, select it and use the up-arrow and down-arrow icons.

- **Name**: Name for the column.
- **Datatype**: Data type for the column.
- **Default**: If no value is specified, the default value is null.
- **Default on NULL**: Applicable for Oracle Database 12c and later releases. If this option is selected, when a row is inserted into the table and the value specified for the column is NULL, the default value is inserted into the column.
- **Expression**: Expression for computing the value in the column.
- **Comments**: Optional descriptive comments about the column. Use this field to provide descriptions for the attributes.

In the table:

- **PK**: If this option is selected, the column becomes the primary key.
- **Identity Column** : If this option is selected, the column becomes an identity column. This is applicable only for Oracle Database 12c and later releases. For more details, see the Identity Column tab.

Constraints tab

Displays the Not Null and Check Constraints for a column. A check constraint requires values in a column to comply with a specified condition.

- **Not Null Constraint: Name**: Name for the Not Null constraint.

- **Not Null Constraint: Not Null**: If this option is selected, the column must contain data. You cannot specify no value or an explicit null value for this column when you insert a row. If this option is not checked, the column can contain either data or no data. A primary key column cannot be null.
- **Check Constraint: Name**: Name for the check constraint definition.
- **Check Constraint: Constraint**: Condition that must be met for a column to fulfill the check constraint. You can use any valid CHECK clause (without the CHECK keyword). For example, to indicate that the value in a numeric column named RATING must be from 1 to 10, you can specify: rating >=1 and rating <= 10.
- **Enabled**: If this option is selected, the constraint is checked when data is entered or updated in the column.
- **Deferrable**: If this option is selected, you can defer checking the validity of the constraint until the end of a transaction.
- **Initially Immediate**: If this option is selected, the constraint is checked whenever you add, update, or delete data from the column.
- **Validate**: If this option is selected, the existing data is checked to see if it conforms to the constraint.

Primary Key Pane

Specifies the primary key for the table.

The primary key is the column, or set of columns, that uniquely identifies each row in the table. If the Primary Key checkbox is selected for a column in the General tab, the corresponding fields are automatically populated in the Primary Key pane. You can make changes to the properties as required.

An index is automatically created on the primary key.

- **Name**: Name of the constraint to be associated with the primary key definition.
- **Enabled**: If this option is checked, the primary key constraint is enforced: that is, the data in the primary key column (or set of columns) must be unique and not null.
- **Index**: Name of the index to which the primary key refers.Tablespace: Name of the tablespace associated with the index.
- **Available Columns**: Lists the columns that are available to be added to the primary key definition. You can select multiple attributes, if required, for the primary key.
- **Selected Columns**: Lists the columns that are included in the primary key definition. To add a column to the primary key definition, select it in Available Columns and click the Add (>) icon; to remove a column from the primary key definition, select it in Selected Columns and click the Remove (<) icon. To move all columns from available to selected (or the reverse), use the Add All (>>) or Remove All (<<) icon. To move a column up or down in the primary key definition, select it in Selected Columns and use the arrow buttons.

Unique Keys Pane

Specifies one or more unique constraints for the table.

A unique constraint specifies a column, or set of columns, whose data values must be unique: each data value must not be null, and it must not be the same as any other value in the column.

To add a unique constraint, click the Add button; to delete a unique constraint, select it and click the Remove button.

- **Name**: Name of the unique constraint.
- **Enabled**: If this option is selected, the unique constraint is enforced.
- **Rely**: If this option is selected, the constraint in NOVALIDATE mode is taken into account during query rewrite.
- **Deferrable**: If this option is selected, in subsequent transactions, constraint checking can be deferred until the end of the transaction using the SET CONSTRAINT(S) statement.
- **Initially Immediate**: If this option is selected, the constraint is checked at the end of each subsequent SQL statement.
- **Validate**: If the option is selected, the existing data is checked to see if it conforms to the constraint.
- **Index**: Name of the index to which the unique key refers.
- **Tablespace**: Name of the tablespace associated with the index.
- **Available Columns**: Lists the columns that are available to be added to the unique constraint definition.
- **Selected Columns**: Lists the columns that are included in the unique constraint definition. To add a column to the unique constraint definition, select it in Available Columns and click the Add (>) icon; to remove a column from the unique constraint definition, select it in Selected Columns and click the Remove (<) icon. To move all columns from available to selected (or the reverse), use the Add All (>>) or Remove All (<<) icon. To move a column up or down in the unique constraint definition, select it in Selected Columns and use the arrow buttons.

Indexes Pane

Lists the indexes defined for the table.

To add an index, click Add Index (+); to delete an index, select it and click Remove Index (-).

- **Name**: Name of the index.
- **Type**: The type of Oracle index. Non-unique means that the index can contain multiple identical values; Unique means that no duplicate values are permitted; Bitmap stores rowids associated with a key value as a bitmap.
- **Tablespace**: Name of the tablespace for the index.
- **Expression**: A column expression is an expression built from columns, constants, SQL functions, and user-defined functions. When you specify a column expression, you create a function-based index.
- **Available Columns** and **Selected Columns**: Columns selected for the index. To select a column, click the column in the Available Columns box, and then click the click the Add Selected Columns icon to move it to the Selected Columns box.

Foreign Keys Pane

Specifies one or more foreign keys for the table.

A foreign key specifies a column ("local column"), whose data values match values in the primary key or unique constraint of another table.

- **Name**: Name of the foreign key definition.
- **Enabled**: If this option is checked, the foreign key is enforced.

- **Rely**, **Deferrable**, **Initially Immediate**, **Validate**: See the description of these fields in the Unique Keys pane.
- **Referenced Constraint: Schema**: Name of the schema containing the table with the primary key or unique constraint to which this foreign key refers.
- **Referenced Constraint: Table**: Name of the table with the primary key or unique constraint to which this foreign key refers.
- **Referenced Constraint: Constraint**: Name of the primary key or unique constraint to which this foreign key refers.
- **Referenced Constraint: On Delete**: Action to take automatically when a row in the referenced table is deleted and rows with that value exist in the table containing this foreign key: NO ACTION (shown by a crossing line in diagrams) performs no action on these rows; CASCADE (shown by an "X") deletes these rows; SET NULL (shown by a small circle) sets null all columns in those rows that can be set to a null value.
- **Associations: Local Column**: Lists the column in the currently selected (local) table that is included in the foreign key definition. For each referenced column in the foreign key definition, select the name of a column in the edited table.
- **Associations: Referenced Column**: For each local column, identifies the column in the other (foreign) table that must have a value matching the value in the local column.

Table Constraints Pane

Specifies one or more check constraints for the table.

A check constraint specifies a condition that must be met when a row is inserted into the table or when an existing row is modified.

- **Name**: Name of the check constraint definition.
- **Check Condition**: Condition that must be met for a row to fulfil the check constraint. You can use any valid CHECK clause (without the CHECK keyword). For example, to indicate that the value in a numeric column named RATING must be from 1 to 10, you can specify rating >=1 and rating <= 10.
- **Enabled**: If this option is checked, the check constraint is enforced.

Comments Pane

Enter descriptive comments in this pane.

This is optional.

Storage Pane

Enables you to specify storage options for the table.

When you create or edit a table or an index, you can override the default storage options.

- **Organization**: Specifies that the table is stored and organized with (Index) or without an index (Heap) or as an external table (External).
- **Tablespace**: Name of the tablespace for the table or index.
- **Logging**: ON means that the table creation and any subsequent INSERT operations against the table are logged in the redo log file. OFF means that these operations are not logged in the redo log file.

- **Row Archival**: YES enables in-database archiving, which allows you to archive rows within the table by marking them as invisible.

External Table Properties Pane

Specifies options for an external table.

An external table is a read-only table whose metadata is stored in the database but whose data in stored outside the database.

External Table

- **Access Driver Type**: Specifies the type of external table.
  - **ORACLE_LOADER**: Extracts data from text data files. This is the default access driver, which loads data from external tables to internal tables.
  - **– ORACLE_DATAPUMP**: Extracts data from binary dump files. This access driver can perform both loads and unloads.
  - **– ORACLE_BIGDATA**: Extracts data from Oracle Big Data Appliance.
  - **– ORACLE_HDFS**: Extracts data stored in a Hadoop Distributed File System (HDFS).
  - **– ORACLE_HIVE**: Extracts data stored in Apache HIVE.
- **Default Directory**: Specifies the default directory to use for all input and output files that do not explicitly name a directory object. The location is specified with a directory object, not a directory path.
- **Access Params**: Assigns values to the parameters of the specific access driver for the external table. Access parameters are optional.
  - **OPAQUE_FORMAT_SPEC**: The opaque_format_spec specifies all access parameters for the ORACLE_LOADER, ORACLE_DATAPUMP, ORACLE_HDFS, and ORACLE_HIVE access drivers. For descriptions of the access parameters, see *Oracle Database Utilities*. Field names specified in the opaque_format_spec must match columns in the table definition, else Oracle Database ignores them.
  - **USING CLOB**: Enables you to derive the parameters and their values through a subquery. The subquery cannot contain any set operators or an ORDER BY clause. It must return one row containing a single item of data type CLOB.
- **Reject Limit**: The number of conversion errors that can occur during a query of the external data before an Oracle Database error is returned and the query is aborted.
- **Project Column**: Determines how the access driver validates the rows of an external table in subsequent queries.
  - **ALL**: Processes all column values, regardless of which columns are selected, and validates only those rows with fully valid column entries. If any column value raises an error, such as a data type conversion error, the row is rejected even if that column was not referenced in the select list of the query.
  - **REFERENCED**: Processes only those columns in the select list of the query. The ALL setting guarantees consistent result sets. The REFERENCED setting can result in different numbers of rows returned, depending on the columns referenced in subsequent queries, but is faster than the ALL setting. If a subsequent query selects all columns of the external table, then the settings behave identically.
- **Location**: Specifies the data files for the external table. Use the Add (+) icon to add each location specification.
  - ORACLE_LOADER and ORACLE_DATAPUMP, the files are named in the form directory:file. The directory portion is optional. If it is missing, then the default directory

is used as the directory for the file. If you are using the ORACLE_LOADER access driver, then you can use wildcards in the file name. An asterisk (*) signifies multiple characters and a question mark (?) signifies a single character.

- For ORACLE_HDFS, LOCATION is a list of Uniform Resource Identifiers (URIs) for a directory or for a file. There is no directory object associated with a URI.
- For ORACLE_HIVE, LOCATION is not used. Instead, the Hadoop HCatalog table is read to obtain information about the location of the data source (which could be a file or another database).

Opaque Format Spec

Specifies all access parameters for the ORACLE_LOADER, ORACLE_DATAPUMP, ORACLE_HDFS, and ORACLE_HIVE access drivers.

CLOB Subquery

Type or copy and paste the query.

Materialized View Pane

Specifies options for a materialized view.

Query: Contains the SQL code for the query part of the view definition. Type or copy and paste the query.

General

- **On Pre-built Table**: If **Yes**, an existing table is registered as a preinitialized materialized view. This option is particularly useful for registering large materialized views in a data warehousing environment. The table must have the same name and be in the same schema as the resulting materialized view, and the table should reflect the materialization of a subquery.
- **Reduced Precision**: **Yes** authorizes the loss of precision that will result if the precision of the table or materialized view columns do not exactly match the precision returned by the subquery. If **No**, the precision of the table or materialized view columns must exactly match the precision returned by the subquery, or the create operation will fail.
- **For Update**: Select **Yes** to allow a subquery, primary key, object, or rowid materialized view to be updated. When used in conjunction with Advanced Replication, these updates will be propagated to the master.
- **Real Time MV**: Select **Yes** to create a real-time materialized view or a regular view. A real-time materialized view provides fresh data to user queries even when the materialized view is not in sync with its base tables due to data changes. Instead of modifying the materialized view, the optimizer writes a query that combines the existing rows in the materialized view with changes recorded in log files (either materialized view logs or the direct loader logs). This is called on-query computation.
- **Query Rewrite**: If **Enable**, the materialized view is enabled for query rewrite, which transforms a user request written in terms of master tables into a semantically equivalent request that includes one or more materialized views.
- **Build**: Specifies when to populate the materialized view. **Immediate**indicates that the materialized view is to be populated immediately. **Deferred** indicates that the materialized view is to be populated by the next refresh operation. If you specify Deferred, the first (deferred) refresh must always be a complete refresh; until then, the materialized view has a staleness value of unusable, so it cannot be used for query rewrite.
- **Use Index**: If **Yes**, a default index is created and used to speed up incremental (fast) refresh of the materialized view. If **No**, this default index is not created. (For example, you

might choose to suppress the index creation now and to create such an index explicitly later.)

- **Index Tablespace**: Specifies the tablespace in which the materialized view is to be created. If a tablespace is not selected, the materialized view is created in the default tablespace of the schema containing the materialized view.
- **Cache**: If **Yes**, the blocks retrieved for this table are placed at the most recently used end of the least recently used (LRU) list in the buffer cache when a full table scan is performed. This setting is useful for small lookup tables. If **No**, the blocks are placed at the least recently used end of the LRU list.

Refresh Clause

- **Refresh**: Select **Yes** to enable refresh operations.
- **Refresh Type**: The method of refresh operation to be performed:
  - Complete Refresh: Executes the defining query of the materialized view, even if a fast refresh is possible.
  - Fast Refresh: Uses the incremental refresh method, which performs the refresh according to the changes that have occurred to the master tables. The changes for conventional DML changes are stored in the materialized view log associated with the master table. The changes for direct-path INSERT operations are stored in the direct loader log.
  - Force Refresh: Performs a fast refresh if one is possible; otherwise, performs a complete refresh.
- **Action**: The type of refresh operation to be performed:
  - On Demand: Performs a refresh when one of the DBMS_MVIEW refresh procedures are called.
  - On Commit: Performs a fast refresh whenever the database commits a transaction that operates on a master table of the materialized view. This may increase the time taken to complete the commit, because the database performs the refresh operation as part of the commit process.
  - Specify: Performs refresh operations according to what you specify in the Start on and Next fields.
- **Start Date**: Starting date and time for the first automatic refresh operation. Must be in the future.
- **Next Date**: Time for the next automatic refresh operation. The interval between the Start on and Next times establishes the interval for subsequent automatic refresh operations. If you do not specify a value, the refresh operation is performed only once at the time specified for Start on.
- **With**: Refresh type, which determines the type of materialized view:
  - Primary Key: Creates a primary key materialized view, which allows materialized view master tables to be reorganized without affecting the eligibility of the materialized view for fast refresh.
  - Row ID: Creates a rowid materialized view, which is useful if the materialized view does not include all primary key columns of the master tables.
- **Default Storage**: If Yes, DEFAULT specifies that Oracle Database will choose automatically which rollback segment to use. If you specify DEFAULT, you cannot specify the `rollback_segment`. DEFAULT is most useful when modifying, rather than creating, a materialized view.

- **Storage Type**: MASTER specifies the remote rollback segment to be used at the remote master site for the individual materialized view. LOCAL specifies the remote rollback segment to be used for the local refresh group that contains the materialized view. This is the default.
- **Rollback Segment**: Enter the name of the rollback segment.
- **Using Constraint**: If this option is checked, more rewrite alternatives can be used during the refresh operation, resulting in more efficient refresh execution. The behavior of this option is affected by whether you select Enforced or Trusted.
  - Enforced: Causes only enforced constraints to be used during the refresh operation.
  - Trusted: Enables the use of dimension and constraint information that has been declared trustworthy by the database administrator but that has not been validated by the database. If the dimension and constraint information is valid, performance may improve. However, if this information is invalid, then the refresh procedure may corrupt the materialized view even though it returns a success status.

DDL Pane

You can review and save the SQL statements that are generated when creating or editing the object. If you want to make any changes, go back to the relevant panes and make the changes there.

For a new table, click **CREATE** to view the generated DDL statements.

When you edit table properties, click **UPDATE** to view the generated ALTER statements.

For a new table, the UPDATE tab will not be available. When you are finished, click **Apply**.

Output Pane

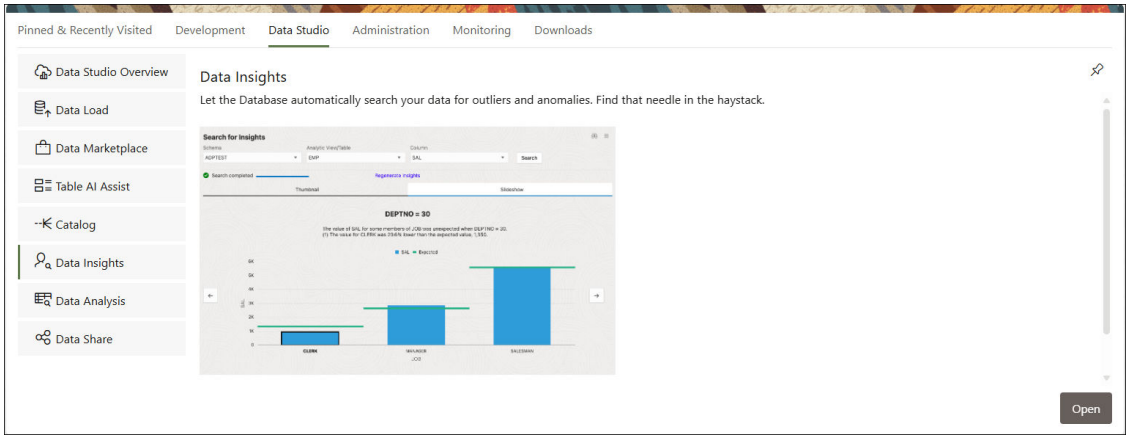
Displays the results of the DDL commands.

If there are any errors, go to the appropriate pane, fix the errors, and run the commands again. You can save to a text file or clear the output.

The Data Insights Page

The Data Insights page displays information about patterns and anomalies in the data of entities in your Oracle Autonomous Database.

To reach the Data Insights page, click the **Data Studio** tab in the Database Actions page, and select the **Data Insights** menu



or click the Selector icon and select **Data Insights** from the Data Tools menu in the navigation pane.

The following topics describe insights and how to generate and use them.

- About Insights
  You can generate insights for a table or for the analytic view deployed for data analysis.
- Generate Insights and View Reports
  Use these procedures to generate Insights and view reports about them.

About Insights

You can generate insights for a table or for the analytic view deployed for data analysis.

The insights that Data Insights generates for the analytic view of a business model can be more useful than those for a table because of the additional metadata that an analytic view provides.

Insights highlight data points as potentially anomalous if the actual value for a measure when filtering on pairs of analytic view hierarchy values or table column values is considerably higher or lower than the expected value, calculated across all hierarchy or column values. Insights highlight unexpected patterns, which you may want to investigate.

Insights are automatically generated by various analytic functions built into the database. The results of the insight analysis appear as a series of bar charts in the Data Insights dashboard.

Data Insights uses the following steps to generate insights:

1. Finds the values of a measure, for example Sales, across all of the distinct pairs of the hierarchy or column values for the measure. If Sales has the hierarchies or columns Marital Status, Age Band, Income Level, and Gender, then the pairs would be the values of each distinct value of each hierarchy or column paired to each distinct value of each of the other hierarchies or columns. For example, if the values of Marital Status are Married and Single, and the values of Age Band are A, B, and C, then the pairs would be Married and A, Married and B, Married and C, Single and A, Single and B, and Single and C. Each distinct value of Marital Status would also be paired with each distinct value of Income Level and Gender, and so on.
2. Estimates an expected value for the measure for each hierarchy or column pair.
3. Calculates the actual value for the measure for each hierarchy or column pair, for example Marital Status = S, Age Band = C, and then the difference between the actual value and the expected value.

4. Scores all of the differences and selects the largest variations between the actual and expected values to highlight as potential insights.

The resulting insights highlight cases where the measure value is significantly larger or smaller for a given hierarchy or column value pair than expected, for example much higher Sales where Marital Status = S and Age Band = C.

Insights for analytic views tend to use the higher levels of a hierarchy because the differences between the estimated and actual values are generally larger than they are for lower level attributes. For example, the difference in dollars between the estimated and actual sales for the entire USA are generally larger than the difference between the estimated and actual sales for a town with a population under 1000. The difference is calculated in absolute values, not percentages.

Insights for tables categorize columns as dimension columns or measure columns based on their data types and cardinality. A VARCHAR2 column is always categorized as a dimension, but a NUMBER column may be either a dimension or a measure. For example, a NUMBER column for YEAR values that has only 10 distinct values in a table with 1 million rows is assumed to be a dimension.

Generate Insights and View Reports

Use these procedures to generate Insights and view reports about them.

Generate Insights

To generate insights for a table or business model, do the following:

1. In the **Schema** field, select a schema.
2. In the **Analytic View/Table** field, select an analytic view or a table.
3. In the **Column** field, select a column that contains data about which you want gain insights.
4. Click **Search**.

A confirmation notice announces that the request for insights has been successfully submitted. Dismiss the notice by clicking the Close (X) icon in the notice.

A progress bar indicates that the search is in progress and when it has completed. The insights appear in the Data Insights dashboard as a series of bar charts.

To refresh the display of the insights, click **Refresh**. To have refreshes occur automatically, click **Enable Auto Refresh**.

Click **Recent Searches** to view the list of previous insights search.

Select **View Errors** to see any log of error that occurs while its creation. The results appear in a new browser tab.

View the Report

The charts in the Data Insights dashboard show the data that contain anomalous results. The bars in a chart show the actual values. The expected values are indicated by green horizontal lines. The bars that are outlined in black contain the most significant differences between the expected and the actual values.

For example, if the fact table for the insights records values about an insurance program, and the measures of the fact table are AGE_CODE, GENDER_CODE, INCOME_CODE, NUM_INSURED, NUM_UNINSURED, and YEAR, then insights might be generated for the NUM_INSURED measure. In that case, the dashboard would have a series of charts labeled

YEAR and INCOME_CODE. Each chart would have a value of the related dimension in the upper left corner. For example, an INCOME_CODE chart that has a related AGE_CODE might have the AGE_CODE value 2 in the upper left corner.

Clicking a chart displays more details about it. At the top of the expanded view of the chart is the dimension name and value and a short textual analysis of notable insights. Below the analysis is the chart showing the values and insights about them.

For example, a chart for INCOME_CODE might have at AGE_CODE = 2 at the top, plus the textual analysis. In the chart, the INCOME_CODE values would be on the x-axis and the NUM_INSURED values would be on the y-axis. Pointing to a bar on the expanded chart displays the actual and the expected NUM_INSURED value for that INCOME_CODE and AGE_CODE.

Click the **Back** button to return to the Data Insights dashboard.

View Previous Reports

To see the results of a previous search, click the Recent Searches icon at the upper right. In the Recent Searches panel, click anywhere in the box for the insights search that you want to see.

To filter the previous searches, enter a value in the search field at the top of the Recent Searches panel.

To close the Recent Searches panel without selecting a search, click the X at the upper right of the panel.

Use Oracle OLAP with Autonomous Database

Oracle OLAP is available on Autonomous Database. Oracle OLAP supports uses such as planning, forecasting, budgeting, and analysis.

- Enable Oracle OLAP on Autonomous Database
  Describes the steps to use Oracle OLAP on Autonomous Database.

Enable Oracle OLAP on Autonomous Database

Describes the steps to use Oracle OLAP on Autonomous Database.

1. Run `DBMS_CLOUD_ADMIN.ENABLE_FEATURE` to enable Oracle OLAP.

   ```
   EXEC DBMS_CLOUD_ADMIN.ENABLE_FEATURE('OLAP');
   ```

   Running this command also enables Java, if it is not already enabled.

   See ENABLE_FEATURE Procedure for more information.
2. Restart the Autonomous Database instance.

   See Restart Autonomous Database for more information.
3. After the restart operation completes, wait for about 20 minutes to allow the Oracle OLAP installation jobs to complete.
4. Verify that Oracle OLAP is installed.

For example:

```
SELECT comp_name, version, status FROM dba_registry
             WHERE comp_id IN ('APS','XOQ','JAVAVM');



COMP_NAME                      VERSION    STATUS
------------------------------ ---------- ------
JServer JAVA Virtual Machine 19.0.0.0.0 VALID
OLAP Analytic Workspace      19.0.0.0.0 VALID
Oracle OLAP API              19.0.0.0.0 VALID
```

5. Verify that Oracle OLAP is working.

```
set serveroutput on
EXEC dbms_aw.execute('show eversion');
```

This should show results like:

```
Oracle OLAP Release 19.23.0.1.0
```

For more information, see the following:

- Getting Started with Oracle OLAP
- Oracle OLAP

Manage the Service

Describes basic tasks for managing an Autonomous Database instance.

- Lifecycle Operations
  You can start, stop, restart, rename, and terminate an Autonomous Database instance.
- Update Compute and storage limits
  Autonomous Database has the ability to increase its storage or CPU cores. You can also activate auto scaling.
- Update Billing Model, Choose Bring Your Own License, or Update Instance to a Paid Instance
  Describes how to choose to bring your own license to your Autonomous Database instance. Also shows how to update the billing model to the ECPU compute model and shows how to update your instance from free (Always Free) to a paid instance.
- Clone, Move, or Upgrade an Autonomous Database Instance
  Autonomous Database provides cloning where you can choose to create a full clone, create a metadata clone, or create a refreshable clone. You can also move a database to a different compartment. By creating a clone you can upgrade an instance with Oracle Database 19c to create an Oracle Database 23ai clone (either a full clone or a refreshable clone).
- Use Refreshable Clones with Autonomous Database
  Autonomous Database provides cloning where you can choose to create a full clone of the active instance, create a metadata clone, or create a refreshable clone. With a refreshable clone the system creates a clone that can be easily updated with changes from the source database.

- Manage Autonomous Database Built-in Tools
  Autonomous Database includes built-in tools that you can enable and disable when you provision or clone a database, or at any time for an existing database.
- Use Autonomous Database Events
- Manage Time Zone File Updates on Autonomous Database
- Monitor Autonomous Database Availability
  Shows you the steps to view availability information for an Autonomous Database instance.
- Monitor Regional Availability of Autonomous Databases
  The regional availability is represented as a list of daily uptime percentages of Autonomous Databases across a set of services in each available region.
- Manage Resource Availability with Compartment Quotas
  Autonomous Database supports the compartment quotas feature of Oracle Cloud Infrastructure.

Lifecycle Operations

You can start, stop, restart, rename, and terminate an Autonomous Database instance.

- Provision an Autonomous Database Instance
  Follow these steps to provision a new Autonomous Database instance using the Oracle Cloud Infrastructure Console.
- Start Autonomous Database
  Describes the steps to start an Autonomous Database instance.
- Stop Autonomous Database
  Describes the steps to stop an Autonomous Database instance.
- Restart Autonomous Database
  Describes the steps to restart an Autonomous Database instance.
- Rename Autonomous Database
  Shows you the steps to change the database name for an Autonomous Database instance.
- Update the Display Name for an Autonomous Database Instance
  Shows you the steps to change the display name for an Autonomous Database instance.
- Terminate an Autonomous Database Instance
  Describes the steps to terminate an Autonomous Database instance.
- Undelete a Terminated Autonomous Database Instance
  Describes the steps to undelete an Autonomous Database instance.
- Change Autonomous Database Operation Mode to Read/Write Read-Only or Restricted
  You can select an Autonomous Database operation mode: Read/Write, Read-Only, or Restricted. The default mode is Read/Write.
- Schedule Start and Stop Times for an Autonomous Database Instance
  Describes the steps to schedule daily start and stop times for an Autonomous Database instance.
- Concurrent Operations on Autonomous Database
  Describes the actions that cause Oracle Cloud Infrastructure Console to ask you to confirm pausing or canceling a concurrent operation.

Provision an Autonomous Database Instance

Follow these steps to provision a new Autonomous Database instance using the Oracle Cloud Infrastructure Console.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the  next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- Choose your region. See Switching Regions for information on switching regions and working in multiple regions.
- Choose your **Compartment**. See Compartments for information on using and managing compartments.

On the Autonomous Databases page, perform the following steps:

1. Click **Create Autonomous Database**.
2. Provide basic information for the Autonomous Database.
   - **Compartment**. See Compartments for information on using and managing compartments.
   - **Display name** Specify a user-friendly description or other information that helps you easily identify the resource. The display name does not have to be unique.

     The default display name is a generated 16-character string that matches the default database name.
   - **Database name** Specify the database name; it must consist of letters and numbers only. The maximum length is 30 characters. The same database name cannot be used for multiple Autonomous Databases in the same tenancy in the same region.

     The default database name is a generated 16-character string that matches the default display name.
3. Choose a workload type. Select the workload type for your database from the choices:
   - **Data Warehouse**
   - **Transaction Processing**
   - **JSON**
   - **APEX**
4. Configure the database (ECPU compute model)
   - **Always Free**: Select to show Always Free options.

     You can only create a free instance in the tenancy's Home region.
   - **Developer**: Select to show Autonomous Database for Developers options.

     An Autonomous Database for Developers instance is a fixed shape database that is suited for development and testing use cases.

     Selecting this option sets the **Compute** resources to a fixed value of 4 ECPUs and the **Storage** to 20 GB.

- **Choose database version**: Select the database version. The available database versions are Oracle Database 23ai and Oracle Database 19c.

  > **Note:**
  >
  > – In regions where Oracle Database 23ai is not available, Oracle Database 19c is the only choice.
  >
  > – Autonomous Database with Oracle Database 23ai in the Paid tier is available in all commercial public cloud regions.
  >
  > – Always Free Autonomous Database with Oracle Database 23ai is available in all commercial public cloud regions *except* the following regions: Colombia Central: Bogota (BOG), Saudi Arabia Central (RUH), Singapore West: Singapore (XSP)

- **ECPU count**: Specify the number of CPUs for your database. The minimum value is 2.
- **Compute auto scaling**: By default compute auto scaling is enabled to allow the system to automatically use up to three times more CPU and IO resources to meet workload demand. If you do not want to use compute auto scaling then deselect this option.

  See Use Auto Scaling for more information.
- **Storage**: Specify the storage you wish to make available to your database. Depending on your workload type you have these options:
  - **Data Warehouse**: Specify your storage in Terabytes (TB).
  - **Transaction Processing** or **JSON**: Specify your storage in Gigabytes (GB) or Terabytes (TB). Enter the size in the **Storage** field. Select **GB** or **TB** for the **Storage unit size**.
- By default, the IO capacity of your database depends on the number of ECPUs you provision. When you provision 384 TB of storage, your database is entitled to the full IO capacity of the Exadata infrastructure, independent of the number of ECPUs you provision.

  Autonomous Database uses Exadata Smart Flash Cache to automatically cache frequently accessed data, delivering the high I/O rates and fast response times of flash. The amount of flash cache for your database depends on the amount of storage you provision, or the amount of allocated storage if you enable storage auto scaling.

  If you want to provision more than 384 TB of storage, file a Service Request at Oracle Cloud Support.
- **Storage auto scaling**: By default storage auto scaling is disabled. Select if you want to enable storage auto scaling to allow the system to automatically expand to use up to three times more storage.

  With storage auto scaling disabled, the guaranteed minimum flash cache size is 10% of your database's provisioned storage size.

  With storage auto scaling enabled, the guaranteed minimum flash cache size is 10% of your database's provisioned base storage size or its allocated storage size, whichever is higher.

  See Use Auto Scaling for more information.
- **Show advanced options**: Click to show additional options

- **Enable elastic pool:**

  See Create or Join an Elastic Pool While Provisioning or Cloning an Instance for more information.

- **Compute model**: Shows the selected compute model.

  * **ECPU**
    Use the ECPU compute model for Autonomous Database. ECPUs are based on the number of cores elastically allocated from a pool of compute and storage servers.

  * **OCPU**
    Use the legacy OCPU compute model if your tenancy is using the OCPU model and you want to continue using OCPUs. The OCPU compute model is based on physical core of a processor with hyper threading enabled.

  > **Note:**
  >
  > OCPU is a legacy billing metric and has been retired on Autonomous Database. ECPUs are the replacement billing metric for all new and existing Autonomous Database deployments. See Oracle Support Document 2998742.1 for more information.

  Click **Change compute model** to change the compute model. After you select a different compute model, click **Save**.

  See Compute Models in Autonomous Database for more information.

- **Bring your own license**: If you want to Bring Your Own License to the database, click **Enable** to show the **Update license and Oracle Database Edition** page.

  See Choose Bring Your Own License Option While Provisioning or Cloning for more information.

5. Backup retention

   Automatic backup retention period in days You have the option to select the automatic backup retention period, in a range from 1 to 60 days. You can restore and recover your database to any point-in-time in this retention period.

   Select **Immutable backup retention** to lock the backup retention period.

   After setting the immutable backup retention option, you cannot disable this option or change the retention period. To disable immutable backup retention or to change the backup retention period, file a Service Request at Oracle Cloud Support.

   See About Backup and Recovery on Autonomous Database for more information.

6. Create administrator credentials. Set the password for the Autonomous Database Admin user.

   - **Username** This is a read only field.
   - **Password** Set the password for the Autonomous Database Admin user.
   - **Confirm password** Enter the same password again to confirm your new password.

   The password must meet the strong password complexity criteria based on Oracle Cloud security standards. For more information on the password complexity rules, see About User Passwords on Autonomous Database.

7. Choose network access

> **Note:**
>
> After you provision your Autonomous Database you can change the network access option you select for the instance.

- **Secure access from everywhere**

  By default, secure connections are allowed from everywhere.

- **Secure access from allowed IPs and VCNs only**

  This option restricts connections to the database according to the access control lists (ACLs) you specify. To add multiple ACLs for the Autonomous Database, click **Add access control rule**.

  See Configure Access Control Lists When You Provision or Clone an Instance for more information.

- **Private endpoint access only**

  This option assigns a private endpoint, private IP, and hostname to your database. Specifying this option allows traffic only from the VCN you specify; access to the database from all public IPs or VCNs is blocked. This allows you to define security rules, ingress/egress, at the Network Security Group (NSG) level and to control traffic to your Autonomous Database.

  See Configure Private Endpoints When You Provision or Clone an Instance for more information.

See IP Address Ranges for information about the public IP address ranges in Oracle Cloud Infrastructure. You must allow traffic to these CIDR blocks to ensure access to an Autonomous Database instance on a public endpoint.

8. (Optional) Provide contacts for operational notifications and announcements

   Click **Add contact** and in the **Contact email** field, enter a valid email address. To enter multiple **Contact email** addresses, repeat the process to add up to 10 customer contact emails.

   See View and Manage Customer Contacts for Operational Issues and Announcements for more information.

9. (Optional) Click Show advanced options to select advanced options.

   - **Encryption key**
     Encrypt using an Oracle-managed key: By default Autonomous Database uses Oracle-managed encryption keys. Using Oracle-managed keys, Autonomous Database creates and manages the encryption keys that protect your data and Oracle handles rotation of the TDE master key.

     Encrypt using a customer-managed key in this tenancy: If you select this option, a master encryption key from a Oracle Cloud Infrastructure Vault in the same tenancy is used to generate the TDE master key on Autonomous Database.

     Encrypt using a customer-managed key located in a remote tenancy: If you select this option, a master encryption key in the Oracle Cloud Infrastructure Vault located in a remote tenancy is used to generate the TDE master key on Autonomous Database.

     See Use Customer-Managed Encryption Keys on Autonomous Database for more information.

   - **Maintenance**

Patch level By default the patch level is **Regular**. Select **Early** to configure the instance with the early patch level.

See Set the Patch Level for more information.

- **Management**
  Choose a character set and a national character set for your database.

  See Choose a Character Set for Autonomous Database for more information.

- **Tools**
  If you want to view or customize the tools configuration, select the tools tab.

  See Configure Autonomous Database Built-in Tools when you Provision or Clone an Instance for more information.

- **Security attributes**
  Add a security attribute to control access for your resources using Zero Trust Packet Routing (ZPR) policies. To enter security attributes during provisioning you must already have set up security attributes with Zero Trust Packet Routing. You also can add security attributes after provisioning.

  > **Note:**
  >
  > You can apply Oracle Zero Trust Packet Routing (ZPR) policies to a private endpoint.

  Specify a **Namespace**, **Key**, and **Value** security attribute.

  Click **Add security attribute** to add additional security attributes.

  See Overview of Zero Trust Packet Routing for more information.

- **Tags**
  If you want to use Tags, enter the **Tag key** and **Tag value**. Tagging is a metadata system that allows you to organize and track resources within your tenancy. Tags are composed of keys and values which can be attached to resources.

  See Tagging Overview for more information.

10. Optionally, you can save the resource configuration as a stack by clicking **Save as Stack**. You can then use the stack to create the resource through the Resource Manager service.

    Enter the following details on the **Save as Stack** dialog, and click **Save**.

    - **Name:** Optionally, enter a name for the stack.
    - **Description:** Optionally, enter a description for this stack.
    - **Save in compartment:** Select a compartment where this Stack will reside.
    - **Tag namespace**, **Tag key**, and **Tag value**: Optionally, apply tags to the stack.

    For requirements and recommendations for Terraform configurations used with Resource Manager, see Terraform Configurations for Resource Manager. To provision the resources defined in your stack, apply the configuration.

11. Click **Create Autonomous Database**.

On the Oracle Cloud Infrastructure console the Lifecycle state shows **Provisioning** until the new database is available.

Start Autonomous Database

Describes the steps to start an Autonomous Database instance.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

1. On the **Details** page, from the **More actions** drop-down list, select **Start**.

   Start is only shown for a stopped instance.
2. Click **Start** to confirm.

Notes for starting a stopped Autonomous Database instance:

- When an Autonomous Database instance is started, Autonomous Database CPU billing is initiated, billed by the second with a minimum usage period of one minute.
- When an Autonomous Database remains stopped for over 7 days it may take a few minutes or longer to start up, depending on the size of the database. To avoid this extra startup time, Oracle recommends starting your Autonomous Database instance once every 7 days.

Stop Autonomous Database

Describes the steps to stop an Autonomous Database instance.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

1. On the **Details** page, from the **More actions** drop-down list, select **Stop**.
2. Click **Stop** to confirm.

   When concurrent operations such as scaling are active, the confirmation also confirms either pausing or canceling the concurrent operation. See Concurrent Operations on Autonomous Database for more information.

Notes for a stopped Autonomous Database instance:

- Autonomous Database Tools are no longer able to connect to a stopped instance.
- Autonomous Database in-flight transactions and queries are stopped.
- Autonomous Database CPU billing is halted.

- When an Autonomous Database remains stopped for over 7 days it may take a few minutes or longer to start up, depending on the size of the database. To avoid this extra startup time, Oracle recommends starting your Autonomous Database instance once every 7 days.

Restart Autonomous Database

Describes the steps to restart an Autonomous Database instance.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the  next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

1. On the **Details** page, from the **More actions** drop-down list, select **Restart**.
2. In the confirmation dialog, select **Restart** to confirm.

   When concurrent operations such as scaling are active, the confirmation also confirms either pausing or canceling the concurrent operation. See Concurrent Operations on Autonomous Database for more information.

The Autonomous Database instance state changes to "Restarting". After the restart is successful, Autonomous Database instance state is "Available".

> **Note:**
>
> When an Autonomous Database instance is restarted, Autonomous Database CPU billing is initiated, billed by the second with a minimum usage period of one minute.

Rename Autonomous Database

Shows you the steps to change the database name for an Autonomous Database instance.

Before you rename your database, note the following:

- A single tenancy in the same region cannot contain two Autonomous Databases with the same name.
- The database rename operation changes the connection strings required to connect to the database. Thus, after you rename a database you must download a new wallet for any existing instance wallets that you use to connect to the database (and any applications also must be updated to use the new connection string and the new wallet to connect to the database).

  If you are using a regional wallet, you can continue to use the existing wallet to connect to any databases that were not renamed. However, if you want to connect to a renamed database, you must download a new regional wallet.

  See Download Client Credentials (Wallets) for more information.

> **Note:**
>
> The rename operation terminates all connections to the database. After the rename operation completes, you can reconnect to the database.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To rename an Autonomous Database instance, do the following:

1. On the **Details** page, from the **More actions** drop-down list, select **Rename Database**.

   This shows the Rename Database dialog.
2. On the Rename Database dialog, provide values for the fields:
   - New database name: the name you enter must consist of letters and numbers only. The maximum length is 30 characters.
3. Click **Rename**.

After validating the fields, while the system renames the database the lifecycle state changes to **Updating**. You can start using the database after the rename operation completes and the lifecycle state shows **Available**.

Notes for Autonomous Database rename:

- Renaming your database does not change global references to your database from existing database links on remote databases. Changing such references is the responsibility of the administrator of the remote databases.
- When you rename a database, the Autonomous Database OCID does not change.
- If your database has Autonomous Data Guard enabled, the rename operation is not available.
- You cannot use the rename operation on a refreshable clone instance or on a database that is the source for a refreshable clone.
- The rename operation restarts the database.
- You can rename your Autonomous Database using the API. See UpdateAutonomousDatabase for more information.

Update the Display Name for an Autonomous Database Instance

Shows you the steps to change the display name for an Autonomous Database instance.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.

- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To update the display name on an Autonomous Database instance:

1. On the **Details** page, from the **More actions** drop-down list, select **Update display name**.
2. Enter a display name in the **New display name** field.
   - The display name minimum length is 1 character and maximum length is 255 characters.
   - The display name must include only letters, numbers, underscores "_", and hyphens "-".
   - The display name must start with a letter or an underscore "_", and cannot contain two successive hyphens "-".
3. Click **Update display name**.

You can see the updated display name on the Oracle Cloud Infrastructure Console.

Terminate an Autonomous Database Instance

Describes the steps to terminate an Autonomous Database instance.

> **Note:**
>
> Terminating an Autonomous Database deletes the instance and it will no longer be available to use. You may undelete a terminated database instance and its retained backups using the Undelete database option. See Undelete a Terminated Autonomous Database Instance for more information.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

1. On the **Details** page, from the **More actions** drop-down list, select **Terminate**.
2. On the **Terminate Autonomous Database** page enter the database name to confirm that you want to terminate the database.

   When concurrent operations such as scaling are active, the confirmation also confirms either pausing or canceling the concurrent operation. See Concurrent Operations on Autonomous Database for more information.
3. Click **Terminate**.

Notes for terminating a database:

- There are limitations for terminating a database when Autonomous Data Guard is enabled with a cross-region standby database. See Terminate a Cross-Region Standby Database for more information.
- You must first terminate the alastic pool before you terminate a database that is an elastic pool leader. See Terminate an Elastic Pool for more information.
- After a database is terminated, you can list the terminated databases. To show terminated databases, on the Autonomous Databases list page, in the Search box enter **Terminated databases**, select **Include**, and click **Apply filter**.

Undelete a Terminated Autonomous Database Instance

Describes the steps to undelete an Autonomous Database instance.

You can undelete a database up to 72 hours after the database was terminated, or up to the backup retention period, whichever is lower. Using undelete is not possible after more than three days (72 hours) from the time you terminated the database.

Before you undelete your database, note the following:

- To undelete an Autonomous Database instance the instance must use the ECPU compute model.

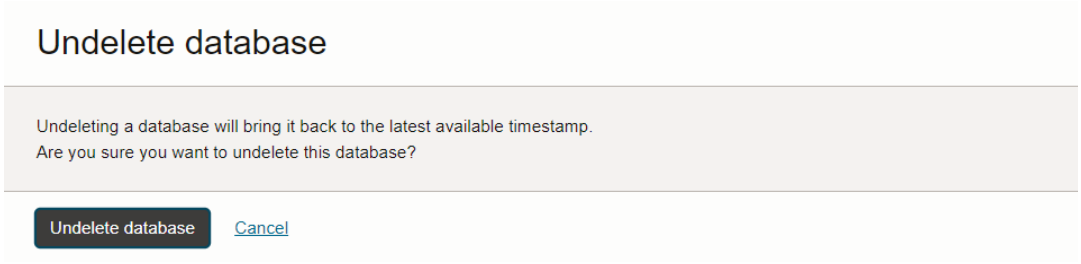
  See Compute Models in Autonomous Database for more information.
- If a different database with the same name exists in the same region, you must first rename or terminate the existing database before you can successfully undelete the database.
- If a database was configured with a Private Endpoint within a Virtual Cloud Network (VCN), a VCN must exist before you can undelete the database. If VCN is not available, use Oracle Cloud Support to file a Service Request to recover the VCN.
- The undelete operation is not available when the terminated database is an Autonomous Database for Developers instance.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the  next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.

To undelete an Autonomous Database instance:

1. On the Autonomous Databases page in the **Resources** area, from the **State** drop-down list, select **Terminated**.
2. In a row that shows your terminated database, click  at the end of the row and select **Undelete**. This shows the Undelete database dialog.



3. Click **Undelete database**.

   The Autonomous Database Lifecycle state changes to **Provisioning**.

While the system undeletes the database, the Lifecycle state changes from **Provisioning** to **Updating**, and then to **Available**.

You can start using the database after the undelete operation completes and the lifecycle state shows **Available**.

Notes for Autonomous Database undelete:

- If you are unable to undelete a database instance within 72 hours, file a service request at Oracle Cloud Support or contact your support representative for assistance.
- When an Autonomous Database instance is undeleted, it will be billed as a running database for the time-period that it was terminated. See Oracle Autonomous Database Serverless Features Billing for more information.
- If you were using a wallet to connect to your database, you must download the wallet again after you undelete the database.

  See Download Client Credentials (Wallets) for more information.
- Any additional features enabled after the initial creation of the database must be re-enabled after you undelete a database. For example, you need to re-enable Data Safe registration, Autonomous Data Guard, or Flashback Time Travel after you undelete.
- When you undelete your database, the following are automatically restored:
  - All the backups before the restore timestamp, including automatic backups and if there are any long-term backups.
  - The most recent Network Access options that were specified before the database was deleted.

    See About Network Access Options for more information.
  - The most recent Backup retention period specified for automatic backups and long-term backups before the database was deleted.
- You cannot perform the undelete operation if the terminated database is:
  - A refreshable clone
  - A Backup-Based Disaster Recovery peer
  - An Autonomous Data Guard standby database

  However, if you delete one of these, you can provision them again from the source database.

Change Autonomous Database Operation Mode to Read/Write Read-Only or Restricted

You can select an Autonomous Database operation mode: Read/Write, Read-Only, or Restricted. The default mode is Read/Write.

- Change Autonomous Database Operation Mode for a Session
  You can set the Autonomous Database operation mode for a session to Read-Only. When you set the session mode to Read-Only users in the session can only run queries.
- Change Autonomous Database Operation Mode to Read/Write Read-Only or Restricted
  From the Oracle Cloud Infrastructure Console you can select an Autonomous Database operation mode: Read/Write, Read-Only, or Restricted. The default mode is Read/Write.
- Autonomous Database Operation in Read-Only Mode
  When the database is in Read-Only mode, each of the following is disabled:

Change Autonomous Database Operation Mode for a Session

You can set the Autonomous Database operation mode for a session to Read-Only. When you set the session mode to Read-Only users in the session can only run queries.

To enable Read-Only mode for a session:

1. Connect to Autonomous Database.

   You can connect as the ADMIN user or as the user whose session you want to set to Read-only mode.
2. Run the following SQL statement:

```
ALTER SESSION SET READ_ONLY = TRUE;
```

To disable Read-Only mode for a session:

```
ALTER SESSION SET READ_ONLY = FALSE;
```

> **Note:**
>
> The `READ_ONLY` session parameter applies when the database is in Read/Write mode. If the database is in Read-Only mode, setting the session parameter value to `TRUE` or to `FALSE` does not change the operation mode for the session.

Change Autonomous Database Operation Mode to Read/Write Read-Only or Restricted

From the Oracle Cloud Infrastructure Console you can select an Autonomous Database operation mode: Read/Write, Read-Only, or Restricted. The default mode is Read/Write.

If you select Read-Only mode users can only run queries. In addition, for Read-Only or Read/Write mode, you can restrict access to only allow users with the `RESTRICTED SESSION` privilege to connect to the database. The ADMIN user has this privilege. You can use the restricted access mode to perform administrative tasks such as indexing, data loads, or other planned activities.

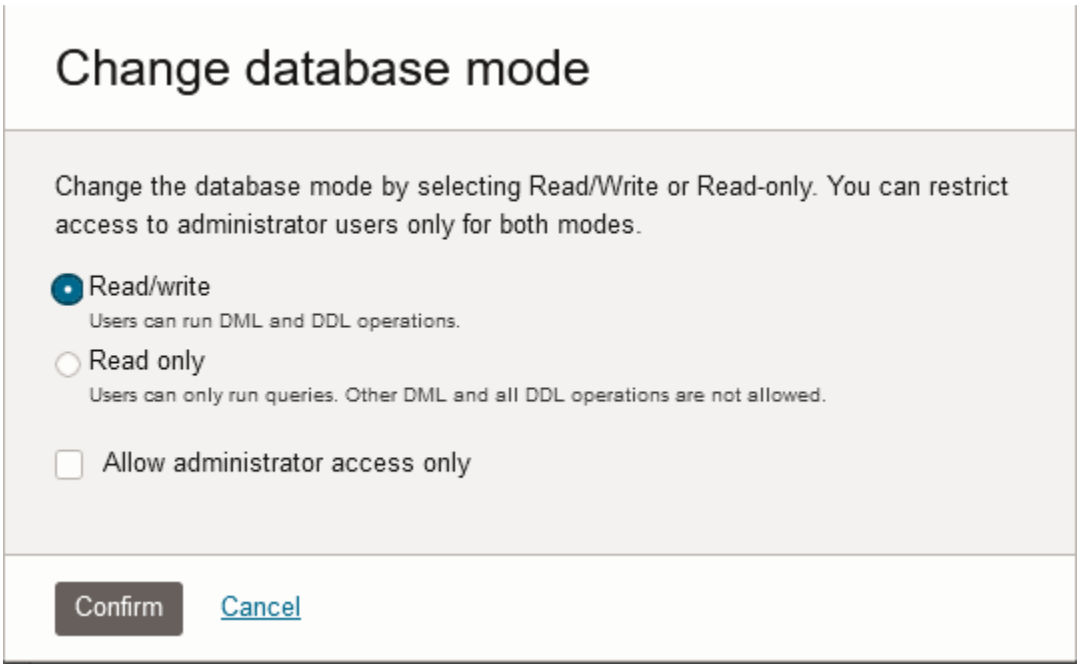
Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To change the Autonomous Database operation mode, do the following:

1. On the **Details** page, in the **Mode** field under General Information, click **Edit**.

   This shows the Change database mode panel.



2. Select a mode, either **Read/write** or **Read only**.
3. If you want to restrict access to only allow the ADMIN or privileged users, then select **Allow administrator access only**.
4. Click **Confirm**.

   When concurrent operations such as scaling are active, the confirmation also confirms either pausing or canceling the concurrent operation. See Concurrent Operations on Autonomous Database for more information.

   While the mode changes, the Lifecycle state field shows **Updating**. After the mode change completes, the Lifecycle state field shows **Available**.

Notes for changing the database mode:

- The database must be available with the **Lifecycle state** field showing **Available** to change the database mode.
- When you change the database mode to restricted mode or when you change from restricted mode to unrestricted mode, the database terminates existing database connections as follows:

  - When the mode is restricted, that is the Mode field shows **(Admin-only access)**, if you deselect **Allow administrator access only** and click **Confirm**, then users and applications need to reestablish database connections after the mode change completes.
  - When the mode is unrestricted, that is the Mode field does not show **(Admin-only access)**, if you select **Allow administrator access only** and click **Confirm**, then users and applications need to reestablish database connections after the mode change completes.

Autonomous Database Operation in Read-Only Mode

When the database is in Read-Only mode, each of the following is disabled:

- On the Oracle Cloud Infrastructure Console, the **Administrator Password** action is not allowed. Also, in Database Actions, under Administration, the **Database Users** card is disabled.
- Oracle APEX URLs are not available:
  - In Database Actions, under Development, the APEX link is disabled.
  - On the Oracle Cloud Infrastructure Console Tools tab, the **Open APEX** link is disabled.
- Oracle Machine Learning User Administration is disabled:
  - In Database Actions, under Development, the Oracle Machine Learning link is disabled.
  - On the Oracle Cloud Infrastructure Console Tools tab, the **Open Oracle ML User Administration** link is disabled.
- Setting resource management rules is disabled. In Database Actions, under Administration, the **Set Resource Management Rules** link is disabled.
- Upgrading the database is disabled when the database is in Read-Only mode.

Schedule Start and Stop Times for an Autonomous Database Instance

Describes the steps to schedule daily start and stop times for an Autonomous Database instance.

When you enable Auto Start/Stop Schedule on an Autonomous Database instance, the instance automatically starts and stops according to the schedule you specify. This allows you to reduce costs by scheduling shutdown periods for times when a system is not in use.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To set the Auto Start/Stop Schedule, do the following:

1. On the **Details** page, from the **More actions** drop-down list, select **Auto start/stop schedule**.
2. For each day you want to schedule, select a start time and a stop time.

The drop-down list allows you to schedule start times and stop times hourly, on the hour, specified in Coordinated Universal Time (UTC).

Selecting both a start time and a stop time is not required. You can select just a start time or just a stop time.

3. Click **Apply** to apply the schedule.
4. In the schedule confirmation, click **Apply**.

   While the system applies the Auto Start/Stop Schedule, the lifecycle state changes to **Updating**. When the operation completes the lifecycle state shows **Available**.

   The **Auto start/stop schedule** field under **General Information** on the Autonomous Database details page lists the days with an Auto Start/Stop Schedule enabled, and includes an **Edit** link that you can click to view the details or make changes.

To remove start and stop times for a day from an Auto Start/Stop Schedule, edit the schedule and choose the Select time (UTC) value for either the start or the stop time for that day.

Notes for Autonomous Database Auto Start/Stop Schedule:

- A scheduled stop works the same way as a manual stop; when the scheduled stop occurs, any running database sessions are stopped.
- You can perform a manual start for a database, if necessary, after a database has stopped due to a scheduled stop.

  See Start Autonomous Database for more information.
- You can perform a manual stop for a database, if necessary, after a database has started due to a scheduled start.

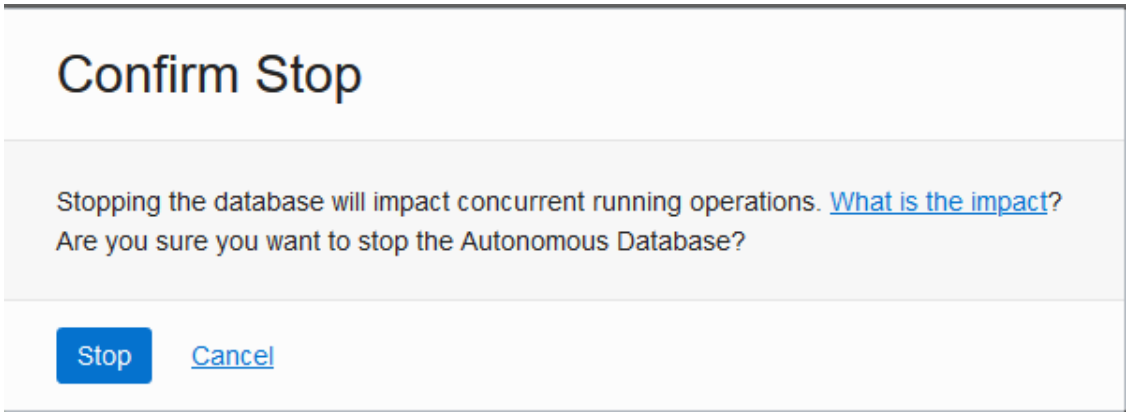
  See Stop Autonomous Database for more information.
- The Lifecycle state must be **Available** to view or to edit the Auto Start/Stop Schedule.
- Always Free Autonomous Database does not support Auto Start/Stop Scheduling.
- A database that has been scheduled stopped, remains stopped until the next scheduled start time or until you start the database. Likewise, a database that has been scheduled started, remains running until the next scheduled stop time or until you stop the database.
- When an Autonomous Database remains stopped for over 7 days it may take a few minutes or longer to start up, depending on the size of the database. To avoid this extra startup time, Oracle recommends starting your Autonomous Database instance once every 7 days.

Concurrent Operations on Autonomous Database

Describes the actions that cause Oracle Cloud Infrastructure Console to ask you to confirm pausing or canceling a concurrent operation.

When you initiate certain operations that take some time to complete, including scaling the system, these operations do not prevent you from performing other operations. For example, database connections, database operations, and most Autonomous Database actions proceed normally while you scale the system. However, certain database lifecycle management actions such as stopping the database have an impact on concurrent long-running operations.

For some actions a confirmation dialog displays to let you know when a concurrent operation is active. This allows you to proceed or to cancel the operation. For example, while stopping Autonomous Database you see a message such as the following:



- Manage Scaling with Concurrent Operations
  Describes the impact on scaling operations when you initiate certain actions during an ongoing scaling request.
- Long-Term Backup with Concurrent Operations
  Describes the impact on a long-term backup operation when you initiate certain actions during an ongoing long-term backup.

Manage Scaling with Concurrent Operations

Describes the impact on scaling operations when you initiate certain actions during an ongoing scaling request.

Action	Description
Stop	Stop pauses scaling. The scaling restarts when the database starts.
Restart	Restart pauses scaling. The scaling restarts when the database starts.
Restore	Restore pauses scaling. The scaling restarts when the restore completes if resources allow.
Switchover	Switchover pauses scaling. The scaling restarts when the switchover completes. The scaling also occurs on the standby database after the switchover.
Failover	Failover pauses scaling. The scaling restarts when the failover completes. The scaling also occurs on the standby database if it is available after the failover.
Terminate	Scaling stops and the database is terminated.
Change Database Mode Read-Write or Read-Only	A mode change pauses scaling. The scaling restarts when the database starts after the mode change completes.

Long-Term Backup with Concurrent Operations

Describes the impact on a long-term backup operation when you initiate certain actions during an ongoing long-term backup.

Action	Description
Stop	Stop cancels an ongoing long-term backup. The long-term backup does not restart when the database starts.
Restart	Restart cancels an ongoing long-term backup. The long-term backup does not restart when the database starts.

Action	Description
Restore	Restore cancels an ongoing long-term backup.
Switchover	Switchover cancels an ongoing long-term backup. The long-term backup does not restart after the switchover completes.
Failover	Failover cancels an ongoing long-term backup. The long-term backup does not restart after the failover completes.
Terminate	Terminate cancels an ongoing long-term backup.

Update Compute and storage limits

Autonomous Database has the ability to increase its storage or CPU cores. You can also activate auto scaling.

- Add CPU or Storage Resources or Enable Auto Scaling
  Describes how to scale your Autonomous Database on demand by adding CPU cores or storage. Also describes how to enable auto scaling.
- Remove CPU or Storage Resources or Disable Auto Scaling
  Describes how to scale your Autonomous Database on demand by removing CPU cores or storage. Also describes how to disable auto scaling.
- Use Auto Scaling
  When you create an Autonomous Database instance, by default compute auto scaling is enabled and storage auto scaling is disabled. You can manage auto scaling from the Oracle Cloud Infrastructure Console to enable or disable compute auto scaling or storage auto scaling.

Add CPU or Storage Resources or Enable Auto Scaling

Describes how to scale your Autonomous Database on demand by adding CPU cores or storage. Also describes how to enable auto scaling.

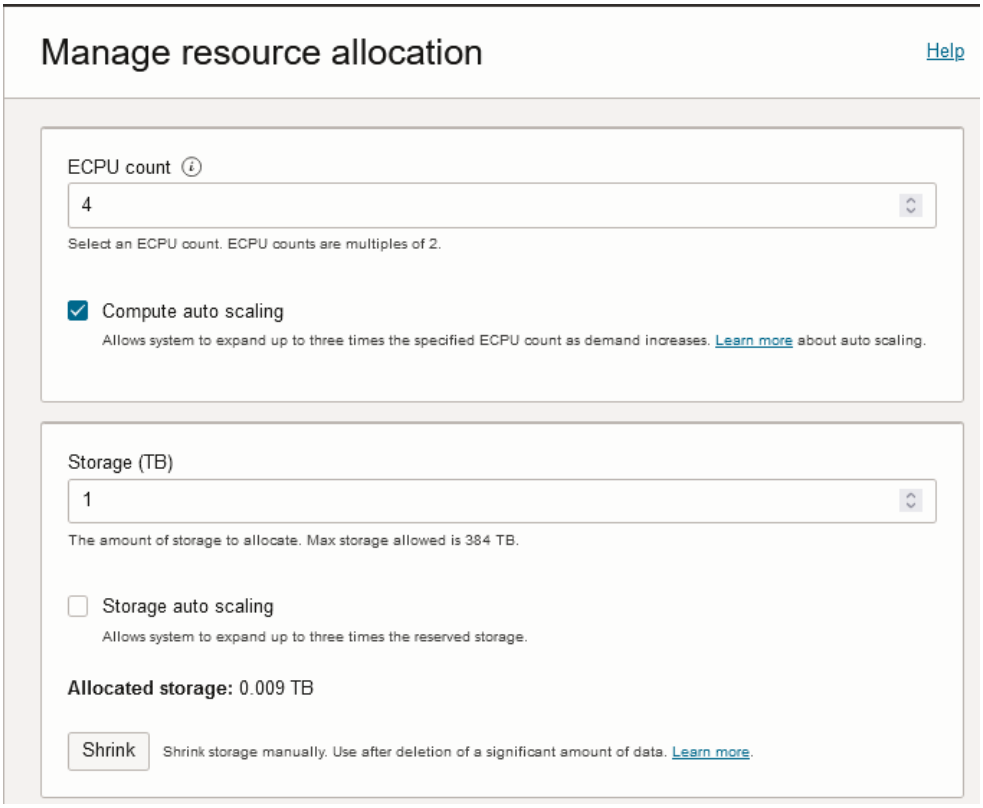
Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

1. On the **Details** page, click **Manage resource allocation**.
2. In the **Manage resource allocation** area, add resources for your scale request.
   - Enter a value or click the up arrow to select a value for **ECPU count** (**OCPU count** if your database uses OCPUs). The default is no change.
   - **Storage**: Specify the storage you wish to make available to your database. Depending on your workload type and your compute model, you have these options:
     – **Data Warehouse**: Specify your storage in Terabytes (TB).
     – **Transaction Processing** or **JSON**: Specify your storage in Gigabytes (GB) or Terabytes (TB). Enter the size in the **Storage** field. Select **GB** or **TB** for the

Storage unit size. GB units are only available when the Workload type is Transaction Processing and the Compute model is **ECPU**.

The default is no change.



3. Enable compute auto scaling.

   Select **Compute auto scaling** to allow the system to automatically use up to three times more CPU and IO resources to meet workload demand, compared to the database operating with compute auto scaling disabled. The additional resources are available until you disable compute auto scaling.

   If compute auto scaling is disabled while more CPUs are in use than the specified **ECPU count** (**OCPU count** if your database uses OCPUs), then Autonomous Database scales the number of CPUs in use down to the **ECPU count** (**OCPU count**).

   See Use Auto Scaling for more information.

4. To enable Storage auto scaling, select **Storage auto scaling**.

   When you select **Storage auto scaling** the database can expand to use up to three times the reserved base storage.

   If you disable **Storage auto scaling** and the used storage is greater than the reserved base storage, as specified by the storage shown in the **Storage** field on the Oracle Cloud Infrastructure Console, Autonomous Database shows a warning on the disable storage auto scaling confirmation dialog. The warning lets you know that the reserved base storage value will be increased to the nearest TB greater than the actual storage usage on a Data Warehouse workload instance, or to the nearest GB greater than the actual storage usage on a Transaction Processing workload instance, and shows the new reserved base storage value.

   See Use Auto Scaling for more information.

> **Note:**
>
> Clicking **Shrink** initiates the Shrink storage operation. See Shrink Storage for more information.

5. Click **Apply** to change your resources.

When you click **Apply** with a resource change, the Lifecycle state changes to **Scaling in Progress...**. After the Lifecycle state changes to **Available** the changes apply immediately.

Note the following for adding storage resources or changing the **Storage auto scaling** value:

- Autonomous Database uses Exadata Smart Flash Cache to automatically cache frequently accessed data, delivering the high I/O rates and fast response times of flash. The amount of flash cache for your database depends on the amount of storage you provision, or the amount of allocated storage if you enable storage auto scaling.

  With storage auto scaling disabled, the guaranteed minimum flash cache size is 10% of your database's provisioned storage size.

  With storage auto scaling enabled, the guaranteed minimum flash cache size is 10% of your database's provisioned base storage size or its allocated storage size, whichever is higher.

Remove CPU or Storage Resources or Disable Auto Scaling

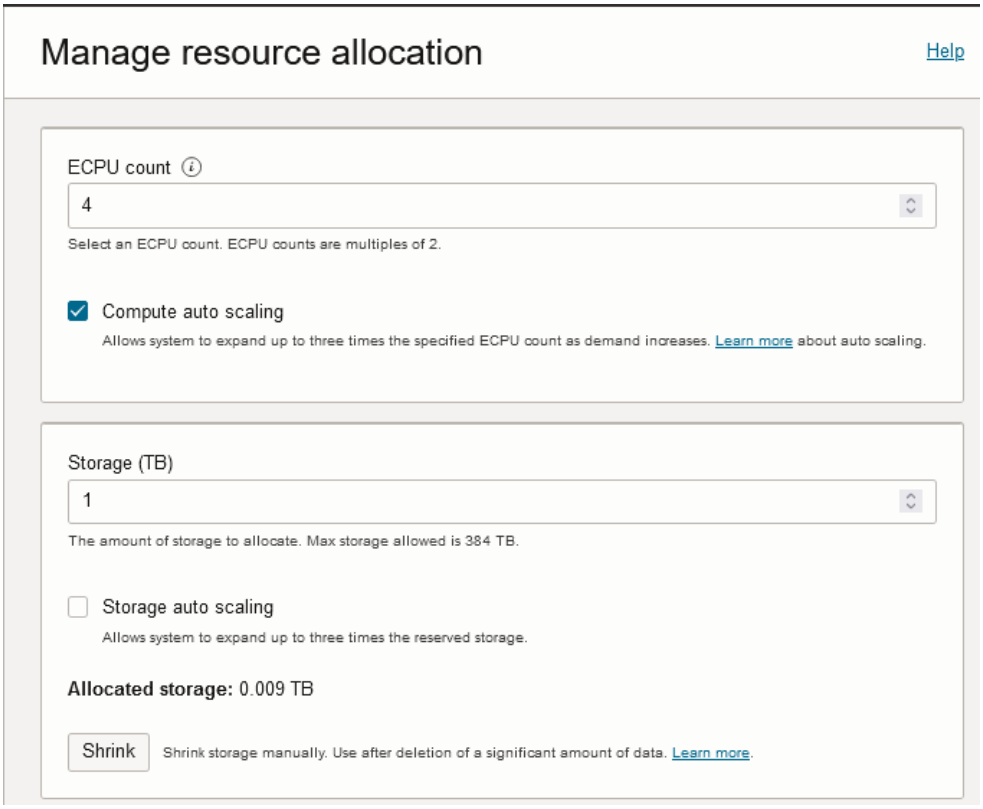
Describes how to scale your Autonomous Database on demand by removing CPU cores or storage. Also describes how to disable auto scaling.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

1. On the **Details** page, click **Manage resource allocation**.
2. In the **Manage resource allocation** area, select the change in resources for your scale request:
   - Enter a value or click the down arrow to select a value for the **ECPU count** (**OCPU count** if your database uses OCPUs). The default is no change.
   - **Storage**: Specify the storage you wish to make available to your database. Depending on your workload type and your compute model, you have these options:
     - **Data Warehouse**: Specify your storage in Terabytes (TB).
     - **JSON**: Specify your storage in Terabytes (TB).
     - **Transaction Processing**: Specify your storage in Gigabytes (GB) or Terabytes (TB). Enter the size in the **Storage** field. Select **GB** or **TB** for the **Storage unit size**. GB units are only available when the Workload type is Transaction Processing and the Compute model is **ECPU**.

     The default is no change.



3. When compute auto scaling is enabled, deselect **Compute auto scaling** to disable compute auto scaling.

   If compute auto scaling is disabled while more CPUs are in use than the specified **ECPU count** (**OCPU count** if your database uses OCPUs), then Autonomous Database scales the number of CPUs in use down to the **ECPU count** (**OCPU count** if your database uses OCPUs).

   See Use Auto Scaling for more information.

4. When Storage auto scaling is enabled, deselect **Storage auto scaling** to disable storage auto scaling.

   If you disable **Storage auto scaling** and the used storage is greater than the reserved base storage, as specified by the storage shown in the **Storage** field on the Oracle Cloud Infrastructure Console, Autonomous Database shows a warning on the disable storage auto scaling confirmation dialog. The warning lets you know that the reserved base storage value will be increased to the nearest TB greater than the actual storage usage on a Data Warehouse workload instance, or to the nearest GB greater than the actual storage usage on a Transaction Processing workload instance, and shows the new reserved base storage value.

   See Use Auto Scaling for more information.

> **Note:**
>
> Clicking **Shrink** initiates the Shrink storage operation. See Shrink Storage for more information.

5. Click **Apply** to change your resources.

When you click **Apply** with a resource change, the Lifecycle state changes to **Scaling in Progress...**. After the Lifecycle state changes to **Available** the changes apply immediately.

Note the following when you scale down storage or change the **Storage auto scaling** value:

- Scaling down storage is not allowed if the Autonomous Database instance contains the following:
  - `MEMOPTIMIZE FOR WRITE` tables
- If you have columns with the `ROWID` data type, the `ROWID`s that these column values point to may change during the scale down storage operation.
- Tables that contain the following may be moved offline during a scale down operation. DML operations on these tables may be blocked for the duration of the move and the table indexes for these tables may become unusable until the scale down operation completes:
  - Tables with bitmap join indexes
  - Nested tables
  - Object tables
  - Partitioned tables with domain indexes
- Autonomous Database uses Exadata Smart Flash Cache to automatically cache frequently accessed data, delivering the high I/O rates and fast response times of flash. The amount of flash cache for your database depends on the amount of storage you provision, or the amount of allocated storage if you enable storage auto scaling.

  With storage auto scaling disabled, the guaranteed minimum flash cache size is 10% of your database's provisioned storage size.

  With storage auto scaling enabled, the guaranteed minimum flash cache size is 10% of your database's provisioned base storage size or its allocated storage size, whichever is higher.

Use Auto Scaling

When you create an Autonomous Database instance, by default compute auto scaling is enabled and storage auto scaling is disabled. You can manage auto scaling from the Oracle Cloud Infrastructure Console to enable or disable compute auto scaling or storage auto scaling.

- Compute Auto Scaling
- Storage Auto Scaling
  When you create an Autonomous Database instance, by default **Storage auto scaling** is disabled. You can manage scaling and enable storage auto scaling from the Oracle Cloud Infrastructure Console or using the API.
- Shrink Storage
  When the storage used in the database is significantly lower than the allocated storage, the shrink operation reduces the allocated storage.

Compute Auto Scaling

With compute auto scaling enabled the database can use up to three times more CPU and IO resources than specified by the number of ECPUs (OCPUs if your database uses OCPUs) as shown in the **ECPU count** or **OCPU count** field on the Oracle Cloud Infrastructure Console.

When auto scaling is enabled, if your workload requires additional CPU and IO resources, the database automatically uses the resources without any manual intervention required. For example:

- In the ECPU compute model, when the **ECPU count** is 512, this allows the database to use up to 512 x 3 ECPUs (1536 ECPUs) when auto scaling is enabled.

  To see the average number of ECPUs used during an hour you can use the "Number of ECPUs allocated" graph on the Overview tab on the **Database Dashboard** card in Database Actions. See Database Dashboard Overview for more information.

- In the OCPU compute model, when the **OCPU count** is 128, this allows the database to use up to 128 x 3 OCPUs (384 OCPUs) when auto scaling enabled.

  To see the average number of OCPUs used during an hour you can use the "Number of OCPUs allocated" graph on the Overview tab on the **Database Dashboard** card in Database Actions. See Database Dashboard Overview for more information.

Enabling compute auto scaling does not change the concurrency and parallelism settings for the predefined services. See Manage Concurrency and Priorities on Autonomous Database for more information.

> **Note:**
>
> Your license type determines the **ECPU count** maximum. For example, if your license type is Bring your own license (BYOL) with Oracle Database Standard Edition (SE), the **ECPU count** maximum is 32.For this license type the maximum allowed value for **ECPU count** is 32. With compute auto scaling enabled you can use up to **ECPU count** x 3 ECPUs. This license restricts the number of ECPUs you can use to a maximum of 32 ECPUs, with or without compute auto scaling enabled.

When compute auto scaling is enabled your database may use and you may be billed for additional CPU consumption as needed by your workload, up to three times (3x) the number of base CPUs (as shown in the **ECPU count** or **OCPU count** field on the Oracle Cloud Infrastructure Console). See Oracle Autonomous Database Serverless Features Billing for details on compute auto scaling billing.

See Add CPU or Storage Resources or Enable Auto Scaling for the steps to enable compute auto scaling.

Storage Auto Scaling

When you create an Autonomous Database instance, by default **Storage auto scaling** is disabled. You can manage scaling and enable storage auto scaling from the Oracle Cloud Infrastructure Console or using the API.

With **Storage auto scaling** enabled, the Autonomous Database can expand to use up to three times the reserved base storage, as specified by the storage shown in the **Storage** field on the Oracle Cloud Infrastructure Console. If you need additional storage, the database automatically uses the reserved storage without any manual intervention required.

You specify the base storage when you provision or clone your database, or you can change the storage at any time by clicking **Manage resource allocation** and changing the storage size. Depending on your workload type and the compute model selection, you have these options to specify the reserved base storage units:

- **Data Warehouse**: Specify your storage in Terabytes (TB).
- **Transaction Processing**: Specify your storage in Gigabytes (GB) or Terabytes (TB). GB units are only available when the Workload type is Transaction Processing and the Compute model is ECPU.

Using the Oracle Cloud Infrastructure Console or the APIs you can provision or auto-scale an Autonomous Database instance's storage up to 384 TB. For storage requirements larger than 384 TB, Oracle recommends that you file a Service Request at Oracle Cloud Support.

For example, if your storage is 100 TB and storage auto scaling is enabled, you have access to a maximum of 300 TB of storage and if your storage is 200 TB, you have access to a maximum of 384 TB (if you requested a larger maximum by filing a service request, then the maximum would be your custom maximum size).

As data flows in, depending on your workload type, you are billed as follows:

- For storage usage below your reserved base storage, you are billed based on your base storage.
- After your allocated storage exceeds your reserved base storage, storage usage is billed based on your allocated storage rounded up to the nearest TB for workload type Data Warehouse and to the nearest GB for workload types Transaction Processing, APEX, or JSON, in a given hour.

For example, when the Workload type is Data Warehouse, if your reserved base storage is 4 TB, until your allocated storage exceeds 4TB of storage, you are billed based on your base storage (4 TB). After you exceed 4 TB, storage is billed based on the allocated storage rounded up to the nearest TB, in a given hour. In this example, if the allocated storage grows over 4 TB in a given hour, say to 4.9 TB, you are billed for 5 TB of storage from that hour onward.

If you then delete 1 TB of data, your allocated storage remains at 4.9 TB and you are billed for 5 TB until you perform a shrink operation. When you perform a shrink operation, Autonomous Database may be able to reduce your allocated storage back to 3.9TB (shrinking the data and undo tablespaces). After the shrink operation completes and your allocated storage (3.9TB) is once again below your reserved base storage (4 TB), you will once again be billed for your reserved base storage of 4 TB. See Shrink Storage for more information.

> **Note:**
>
> Reducing temp tablespace requires a database restart.

If you disable **Storage auto scaling** and the used storage is greater than the reserved base storage, as specified by the storage shown in the **Storage** field on the Oracle Cloud Infrastructure Console, Autonomous Database shows a warning on the disable storage auto scaling confirmation dialog. The warning lets you know that the reserved base storage value will be increased to the nearest TB greater than the actual storage usage on a Data Warehouse workload instance, or to the nearest GB greater than the actual storage usage on a Transaction Processing workload instance, and shows the new reserved base storage value.

To see the Autonomous Database instance storage usage, you can view the "Storage allocated" and "Storage used" graphs on the Overview tab by clicking the **Database**

Dashboard card in Database Actions. See Database Dashboard Overview for more information.

See Add CPU or Storage Resources or Enable Auto Scaling for the steps to enable storage auto scaling.

Note the following when you enable **Storage auto scaling**:

- Autonomous Database uses Exadata Smart Flash Cache to automatically cache frequently accessed data, delivering the high I/O rates and fast response times of flash. The amount of flash cache for your database depends on the amount of storage you provision, or the amount of allocated storage if you enable storage auto scaling.

  With storage auto scaling disabled, the guaranteed minimum flash cache size is 10% of your database's provisioned storage size.

  With storage auto scaling enabled, the guaranteed minimum flash cache size is 10% of your database's provisioned base storage size or its allocated storage size, whichever is higher.

Shrink Storage

When the storage used in the database is significantly lower than the allocated storage, the shrink operation reduces the allocated storage.

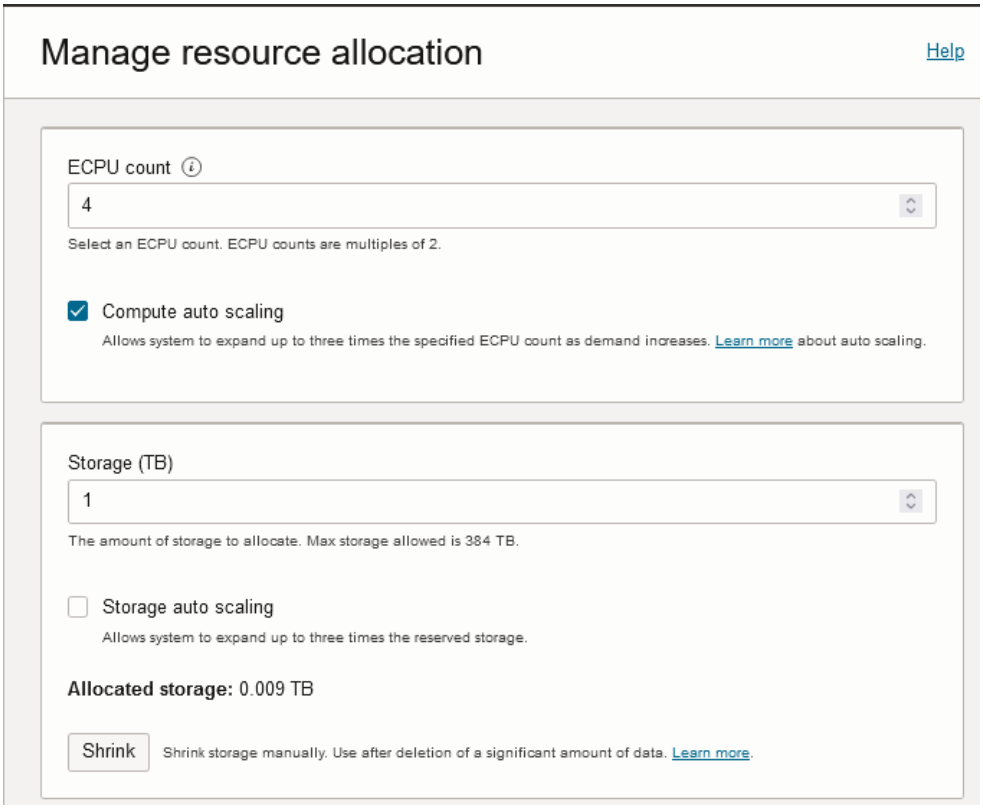
To understand storage allocation and the shrink operation, note the following:

- Reserved base storage: is the base amount of storage you select for the database when you provision or scale the database, excluding any auto-scaled value. The reserved base storage is shown in the **Storage** field on the Oracle Cloud Infrastructure Console.
- Allocated storage: is the amount of storage physically reserved for all database tablespaces (excluding sample schema tablespaces). This number also includes the free space in these tablespaces.
- Used storage: is the amount of storage actually used in all tablespaces (excluding the sample schema tablespaces). The used storage excludes the free space in these tablespaces. Used storage is the storage actually used by database objects, tables, indexes, and so on, including internally used temp space.
- Maximum storage: is the maximum storage reserved. When storage auto scaling is disabled, the maximum storage equals the reserved base storage. When storage auto scaling is enabled, the maximum storage is three times the base storage (maximum = reserved base x 3).

> **Note:**
>
> The **Shrink** operation is not available with Always Free Autonomous Database.

To shrink storage:

1. On the **Details** page, click **Manage resource allocation**.
2. In the **Manage resource allocation** area, select **Shrink**.



3. Click **Confirm** in the Shrink Database dialog.

> **Note:**
>
> The **Shrink** operation is a long running operation.

The **Shrink** operation requires that all of the following apply:

- **Storage auto scaling** must be enabled.
- The allocated storage must be greater than the reserved base storage.
- The allocated storage, rounded up to the nearest 1TB, can be reduced by 1TB or more.
- The following must be true:

  Allocated storage - Used storage > 100 GB

When you click **Shrink** and these conditions are not met, Autonomous Database shows the **Action unavailable** dialog.

Note the following for the Shrink operation:

- The shrink operation runs an `alter table... move online` operation which uses the database's CPUs. In cases where the shrink operation is running slow or taking a very long time, Oracle recommends that you scale up the number of CPUs. See Add CPU or Storage Resources or Enable Auto Scaling for more information.
- The shrink operation is not allowed if the Autonomous Database instance contains the following:
  - Vector indexes

  - Advanced Queuing tables
  - `MEMOPTIMIZE FOR WRITE` tables
- If you have columns with the `ROWID` data type, the `ROWID`s that these column values point to may change during the shrink operation.
- Tables that contain the following may be moved offline during the shrink operation. DML operations on these tables may be blocked for the duration of the move and the table indexes for these tables may become unusable until the shrink operation completes:
  - Tables with bitmap join indexes
  - Nested tables
  - Object tables
  - Immutable tables
  - Blockchain tables
  - Partitioned tables with domain indexes
- If you perform a **Shrink** operation very soon after a data deletion operation, the **Shrink** operation may fail. This can be due to the delay required for Autonomous Database to recalculate storage values. In this case, Oracle recommends that you retry the **Shrink** operation (that is, wait for several minutes for the storage deletion and any associated storage usage updates to complete and perform the **Shrink** operation again).

Update Billing Model, Choose Bring Your Own License, or Update Instance to a Paid Instance

Describes how to choose to bring your own license to your Autonomous Database instance. Also shows how to update the billing model to the ECPU compute model and shows how to update your instance from free (Always Free) to a paid instance.

- Choose Bring Your Own License Option While Provisioning or Cloning
  Describes how to select the option to Bring Your Own License (BYOL) while provisioning or cloning an Autonomous Database instance.
- Choose Bring Your Own License on Autonomous Database (ECPU Compute Model)
- Choose Bring Your Own License on Autonomous Database (OCPU Compute Model)
  Describes how to Bring Your Own License (BYOL) to your Autonomous Database instance.
- Update to ECPU Billing Model on Autonomous Database
- Update Always Free Instance to Paid with Autonomous Database
  Describes how to update your instance to paid from free with Autonomous Database.
- Upgrade Autonomous JSON Database to Autonomous Transaction Processing
  You can promote an Autonomous JSON Database to an Autonomous Transaction Processing database at any time.
- Upgrade APEX Service to Autonomous Transaction Processing
  Easily upgrade APEX Service to Autonomous Transaction Processing.

Choose Bring Your Own License Option While Provisioning or Cloning

Describes how to select the option to Bring Your Own License (BYOL) while provisioning or cloning an Autonomous Database instance.

> **Note:**
>
> The Bring Your Own License (BYOL) option is only available for Autonomous Database instances with the Transaction Processing or Data Warehouse workload types.

During provisioning or cloning you can select **Enable Bring Your Own License (BYOL)** to bring existing database software licenses to the database cloud service when your organization already owns Oracle Database software licenses:

1. During provisioning, in the Configure the database area select **Show advanced options**.
2. In the Configure the database area, in the **Bring your own license** field, click **Enable**.

   This shows the Update license and Oracle Database Edition pane.
3. On the Update license and Oracle Database Edition pane, select **Enable Bring Your Own License (BYOL)**.

   Selecting **Enable Bring Your Own License (BYOL)** displays the Update license and Oracle Database Edition pane.
4. On the Update license and Oracle Database Edition pane, choose an Oracle Database Edition, one of: **Oracle Database Enterprise Edition (EE)** or **Oracle Database Standard Edition (SE).**

   When you select **Bring your own license (BYOL)** you also specify the Oracle Database Edition. The Oracle Database Edition you specify is based on the licenses you bring to Autonomous Database and determines the values that you can select for the **ECPU count**.

Oracle Database Edition	Description
Oracle Database Enterprise Edition (EE)	With this option you can select the **Enable the BYOL ECPU limit and set the number of ECPUs to be covered by BYOL** (we call this the **BYOL ECPU limit** option). • **BYOL ECPU limit** selected: When you select this option you set a **BYOL ECPU limit** for your BYOL licenses.When you select this option you set a **BYOL ECPU limit** for your BYOL licenses. • **BYOL ECPU limit** deselected: For this license type the maximum allowed value for **ECPU count** is 512, however you may contact your Oracle account team to request more ECPUs. With compute auto scaling enabled you can use up to **ECPU count** x 3 ECPUs. For example, if you set the **ECPU count** to 512, you can use up to 1,536 ECPUs.
Oracle Database Standard Edition (SE)	For this license type the maximum allowed value for **ECPU count** is 32. With compute auto scaling enabled you can use up to **ECPU count** x 3 ECPUs. This license restricts the number of ECPUs you can use to a maximum of 32 ECPUs, with or without compute auto scaling enabled.

5. Click **Save**.

6. Continue with any remaining provisioning or cloning steps.

   See Provision an Autonomous Database Instance and Clone an Autonomous Database Instance for more information.

Notes for when you Bring Your Own License (BYOL):

- See Elastic Pool Notes for information on BYOL usage with an instance that is in an elastic pool.
- The **BYOL ECPU limit** you set on an Autonomous Data Guard Primary database does not apply to a cross-region or cross-tenancy Autonomous Data Guard Standby database. On a cross-region or cross-tenancy Standby you can independently set the **BYOL ECPU limit**, as required.
- For information on Bring Your Own License (BYOL) licensing options for Oracle Cloud Infrastructure, see the following:
  - Lower your TCO for Oracle Database Standard Edition with the Bring Your Own License (BYOL) program
  - For Bring Your Own License (BYOL) policies, see Oracle Cloud Services contracts. BYOL policies are described in the **Oracle PaaS and IaaS Universal Credits Service Descriptions** document under **Cloud Service Descriptions**.
  - Frequently asked questions: Oracle Bring Your Own License (BYOL)
  - Cloud pricing

Choose Bring Your Own License on Autonomous Database (ECPU Compute Model)

Describes how to choose Bring Your Own License to your Autonomous Database instance.

> **Note:**
>
> See Choose Bring Your Own License on Autonomous Database (OCPU Compute Model) if you are using the OCPU compute model.

The **License type** field on the Oracle Cloud Infrastructure Console Autonomous Database Information tab shows when you enable the Bring Your Own License (BYOL) option. This field shows your license type and Oracle Database Edition. For example:

License type: Bring Your Own License (BYOL), Enterprise Edition

The Bring Your Own License (BYOL) option is only available for Autonomous Database instances with the Transaction Processing or Data Warehouse workload types.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To choose Bring Your Own License for your Autonomous Database instance:

1. On the **Details** page, from the **More actions** drop-down list, select **Update license and Oracle Database edition**.
2. On the **Update license Oracle Database edition** page enable or disable Bring Your Own License (BYOL) option from the choices:

   Enable **Bring Your Own License (BYOL)** to bring existing database software licenses to the database cloud service when your organization already owns Oracle database software licenses.

   When you select **Bring your own license (BYOL)** you also specify the Oracle Database Edition. The Oracle Database Edition you specify is based on the licenses you bring to Autonomous Database and determines the values that you can select for the **ECPU count**.

   Choose one of: **Oracle Database Enterprise Edition (EE)** or **Oracle Database Standard Edition (SE)**.

Oracle Database Edition	Description
Oracle Database Enterprise Edition (EE)	You have two options, depending on whether you select **Enable BYOL ECPU limit and set the number of ECPUs to be covered by BYOL** (we call this the **BYOL ECPU limit** option and this option is either selected or deselected): • **BYOL ECPU limit** selected: When you select this option you set a **BYOL ECPU limit** for your BYOL licenses. The value you specify indicates how many licenses to use from your existing Enterprise Edition licenses. That is, the **BYOL ECPU limit** indicates how many of your existing licenses you are bringing to the Autonomous Database instance. The **BYOL ECPU limit** you specify indicates how many ECPUs will be covered by the Enterprise Edition licenses you bring. Any ECPU usage beyond the **BYOL ECPU limit** will not by covered by BYOL licenses. • **BYOL ECPU limit** deselected: For this license type the maximum allowed value for **ECPU count** is 512, however you may contact your Oracle account team to request more ECPUs. With compute auto scaling enabled you can use up to **ECPU count** x 3 ECPUs. For example, if you set the **ECPU count** to 512, you can use up to 1,536 ECPUs.
Oracle Database Standard Edition (SE)	For this license type the maximum allowed value for **ECPU count** is 32. With compute auto scaling enabled you can use up to **ECPU count** x 3 ECPUs. This license restricts the number of ECPUs you can use to a maximum of 32 ECPUs, with or without compute auto scaling enabled.

3. Click **Save**.

While the system applies the changes, the lifecycle state changes to **Updating**. The database remains up and accessible, there is no downtime while the license type updates. When the operation completes the lifecycle state shows **Available**.

For information on Bring Your Own License (BYOL) licensing options for Oracle Cloud Infrastructure, see the following:

- Lower your TCO for Oracle Database Standard Edition with the Bring Your Own License (BYOL) program
- For BYOL policies, see Oracle Cloud Services contracts. BYOL policies are described in the **Oracle PaaS and IaaS Universal Credits Service Descriptions** document under **Cloud Service Descriptions**.
- Frequently asked questions: Oracle Bring Your Own License (BYOL)

- Cloud pricing

Notes Update license and Oracle Database edition:

- If you try to switch to Bring Your Own License (BYOL) Standard Edition when the base ECPU count is above 32, you see the following message:
  ```
  You cannot select Oracle Database Standard Edition unless your ECPU count is a total of sixteen (32) or less ECPUs. Use the Manage Scaling option to adjust your ECPU count before changing to Standard Edition.
  ```
  In this case, lower your ECPU count before you enable Bring Your Own License (BYOL). See Remove CPU or Storage Resources or Disable Auto Scaling for more information.
- If you enable **Bring Your Own License (BYOL)** field and select Standard Edition when the base ECPU count is 32 and compute auto scaling is enabled, you see the following message:
  ```
  Compute auto scaling will be disabled with a base ECPU count of 32, under an Oracle Database Standard Edition license.
  ```
  In this case, **Compute auto scaling** is disabled and the **ECPU count** is set to 32.
- If you select **Enable Bring Your Own License (BYOL)** and Standard Edition when the base ECPU count is below 32 and above 8, with compute auto scaling enabled, you see the following message:
  ```
  Your compute auto scaling maximum will be set to 32 ECPUs, under an Oracle Database Standard Edition license.
  ```
  In this case, the maximum number of ECPUs with compute auto scaling is 32.
- See Elastic Pool Notes for information on BYOL usage with an instance that is in an elastic pool.
- The **BYOL ECPU limit** you set on an Autonomous Data Guard Primary database does not apply to a cross-region or cross-tenancy Autonomous Data Guard Standby database. On a cross-region or cross-tenancy Standby you can independently set the **BYOL ECPU limit**, as required.

  See Autonomous Data Guard Notes for information on BYOL usage with Autonomous Data Guard.
- For information on Bring Your Own License (BYOL) licensing options for Oracle Cloud Infrastructure, see the following:
  - Lower your TCO for Oracle Database Standard Edition with the Bring Your Own License (BYOL) program
  - For BYOL policies, see Oracle Cloud Services contracts. BYOL policies are described in the **Oracle PaaS and IaaS Universal Credits Service Descriptions** document under **Cloud Service Descriptions**.
  - Frequently asked questions: Oracle Bring Your Own License (BYOL)
  - Cloud pricing

Choose Bring Your Own License on Autonomous Database (OCPU Compute Model)

Describes how to Bring Your Own License (BYOL) to your Autonomous Database instance.

> **Note:**
>
> See Choose Bring Your Own License on Autonomous Database (ECPU Compute Model) if you are using the ECPU compute model.

The **License type** field shows on the Autonomous Database Information tab when you enable the Bring Your Own License (BYOL) option. This field shows your license and Oracle Database Edition. For example:

License type: Bring Your Own License (BYOL), Enterprise Edition

The Bring Your Own License (BYOL) option is only available for Autonomous Database instances with the Transaction Processing or Data Warehouse workload types.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To enable or disable Bring Your Own License (BYOL):

1. On the **Details** page, from the **More actions** drop-down list, select **Update license and Oracle Database edition**.
2. On the **Update license and Oracle Database edition** page enable or disable the Bring Your Own License (BYOL) option and select an Oracle Database Edition:

   When you enable **Bring Your Own License (BYOL)**, you also choose an Oracle Database Edition. The choices are:

   Oracle Database Enterprise Edition (EE): For this license type the maximum allowed value for **OCPU count** is 128, however you may contact your Oracle account team to request more than 128 OCPUs. With auto scaling enabled you can use up to **OCPU count** x 3 OCPUs. For example, if you set the **OCPU count** to 128, you can use up to 384 OCPUs.

   Oracle Database Standard Edition (SE): For this license type the maximum allowed value for **OCPU count** is 8. With auto scaling enabled you can use up to **OCPU count** x 3 OCPUs. This license restricts the number of OCPUs you can use to a maximum of 8 OCPUs, with or without auto scaling enabled.

   The edition you select is based on the licenses you bring to Autonomous Database and changes the maximum value that you can select for the **OCPU count**.

   See Use Auto Scaling for more information.
3. Click **Save**.

While the system applies the changes, the lifecycle state changes to **Updating**. The database remains up and accessible, there is no downtime while the license type updates. When the operation completes the lifecycle state shows **Available**.

For information on Bring Your Own License (BYOL) licensing options for Oracle Cloud Infrastructure, see the following:

- Lower your TCO for Oracle Database Standard Edition with the Bring Your Own License (BYOL) program
- For BYOL policies, see Oracle Cloud Services contracts. BYOL policies are described in the **Oracle PaaS and IaaS Universal Credits Service Descriptions** document under **Cloud Service Descriptions**.
- Frequently asked questions: Oracle Bring Your Own License (BYOL)
- Cloud pricing

Notes for Update license and Oracle Database Edition:

- If you enable **Bring Your Own License (BYOL)** with Standard Edition when the base OCPU count is above 8, you see the following message:
  ```
  You cannot select Oracle Database Standard Edition unless your OCPU count is a total of eight (8) or less OCPUs. Use the Manage Scaling option to adjust your OCPU count before changing to Standard Edition.
  ```
  In this case, lower your OCPU count before you enable Bring Your Own License (BYOL). See Remove CPU or Storage Resources or Disable Auto Scaling for more information.
- If you enable **Bring Your Own License (BYOL)** with Standard Edition when the base OCPU count is 8 and **Compute auto scaling** is enabled, you see the following message:
  ```
  Compute auto scaling will be disabled with a base OCPU count of 8, under an Oracle Database Standard Edition license.
  ```
  In this case, **Compute auto scaling** is disabled and the OCPU count is set to 8.
- If you select **Enable Bring Your Own License (BYOL)** with Standard Edition when the base OCPU count is below 8 and above 2, with **Compute auto scaling** enabled, you see the following message:
  ```
  Your compute auto scaling maximum will be set to 8 OCPUs, under an Oracle Database Standard Edition license.
  ```
  In this case, the maximum number of OCPUs with **Compute auto scaling** enabled is 8.

Update to ECPU Billing Model on Autonomous Database

Describes how to update an Autonomous Database instance from the OCPU billing model to the ECPU billing model.

> **Note:**
>
> If you update an Autonomous Database instance to use the ECPU compute model, you cannot revert back to the OCPU compute model.

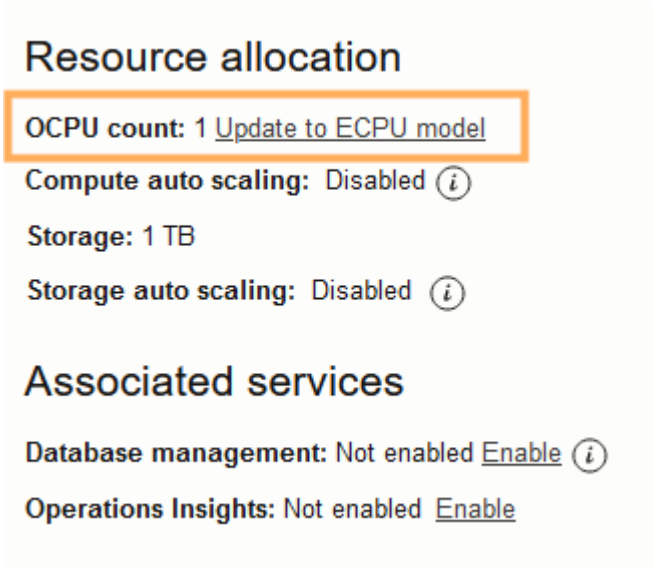
See Compute Models in Autonomous Database for details on the Autonomous Database compute models.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To update to the ECPU compute model:

1. On the **Details** page, under **Resource allocation**, in the **OCPU count field** click **Update to ECPU model**.



2. In the **Update billing model** area.
   - **Compute**: For the compute options **ECPU count** and **Compute auto scaling**, the values are read only and you must accept the defaults. You can change the values after you update to ECPU billing model.

     See Use Auto Scaling for more information.
   - **Backup retention**: The default backup retention is 60 days.

     To change the default specify a value for **Automatic backup retention period in days**.

     See Edit Automatic Backup Retention Period on Autonomous Database for more information.
3. Click **Save changes**.

While the system applies the changes, the Lifecycle state changes to **Updating**.

The database remains up and accessible, there is no downtime when you update the billing model. When the operation completes the Lifecycle state shows **Available** and on the details page under **Resource allocation**, the fields show **ECPU count** and **Compute auto scaling**.

Notes for Updating to the ECPU billing model:

- Your current storage selection is retained. You may scale up or down your storage after updating the billing model. See Add CPU or Storage Resources or Enable Auto Scaling for more information.

- Updating to the ECPU compute model impacts the cost of your compute, storage, and backups. See Compute Models in Autonomous Database for details on the ECPU compute model.
- Peer databases connected to an Autonomous Database instance that is updated to the ECPU compute model are also updated to the ECPU compute model.

  See Use Standby Databases with Autonomous Data Guard for Disaster Recovery and Use Backup-Based Disaster Recovery for more information on the Autonomous Database disaster recovery options.
- Any refreshable clones associated with an Autonomous Database instance that you update to use the ECPU compute model do not change their compute model (the associated refreshable clones continue to use the OCPU compute model). On a refreshable clone, you can update the instance to use the ECPU compute model. See Use Refreshable Clones with Autonomous Database for more information.

Update Always Free Instance to Paid with Autonomous Database

Describes how to update your instance to paid from free with Autonomous Database.

You can upgrade from an Always Free Autonomous Database account to a paid account at any time. If you have a free Oracle Cloud Infrastructure account this is a two step process:

First, upgrade to a paid Oracle Cloud Infrastructure account.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Always Free Autonomous Database instance from the links under the **Display Name** column.

> **Note:**
>
> Promotion of Always Free to a paid Autonomous Database is supported when the database version for the Always Free Autonomous Database is Oracle Database 19c or Oracle Database 23ai.

Upgrade your Always Free Autonomous Database instance as follows:

1. On the **Details** page, from the **More actions** drop-down list, select **Update instance to paid**.
2. On the **Update Instance** page confirm and proceed with the upgrade.

When you upgrade from an Always Free instance to a paid instance you get a paid instance of the same workload type with the minimum CPU and minimum storage available for that workload type. After you upgrade you can scale up CPU and storage resources to fit your needs.

For example:

- **Data Warehouse** workload type: When you upgrade to paid from an Always Free instance, the upgrade process provides a paid Autonomous Database instance with workload type

Data Warehouse, with 2 ECPUs and 1 TB of database storage (if your instance is not enabled to use ECPUs you get 1 OCPU and 1 TB of database storage).

- **Transaction Processing** workload type: When you upgrade to paid from an Always Free instance, the upgrade process provides a paid Autonomous Database instance with workload type **Transaction Processing**, with 2 ECPUs and 1 TB of database storage (if your instance is not enabled to use ECPUs you get 1 OCPU and 1 TB of database storage).

Upgrade Autonomous JSON Database to Autonomous Transaction Processing

You can promote an Autonomous JSON Database to an Autonomous Transaction Processing database at any time.

An Autonomous JSON Database is the same as an Autonomous Transaction Processing database, except that an Autonomous JSON Database is limited in these respects:

- You can store only up to 20 GB of data other than JSON document collections.[1]

  (All Autonomous Databases, including , limit the storage of JSON data to 128 TB.)

- Collections cannot be **heterogeneous**. That is, they can only contain JSON documents. For example, you cannot have a collection of image documents or a collection that contains both JSON documents and image documents.

These limitations are appropriate if your use is primarily development of applications that use JSON documents. If you have a greater need to use data other than JSON data then follow these steps to promote your database to an Autonomous Transaction Processing database:

1. Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
2. From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database**, and then click Autonomous JSON Database.
3. Choose your JSON Database from those listed in the compartment, by clicking its name in column **Display name**.
4. Do one of the following:
   - From the **More actions** drop-down list, select **Change workload type**.
   - In tab Autonomous Database information, under heading **General information**, item **Workload type**, click **Edit**.
5. Click **Convert** to confirm that you want to convert this database to Autonomous Transaction Processing.
6. If you were using the refreshable clone feature with your Autonomous JSON Database then re-create the clone after promotion to Autonomous Transaction Processing. See Using Refreshable Clones with Autonomous Database.

See Autonomous Database Billing Summary for more information.

Upgrade APEX Service to Autonomous Transaction Processing

Easily upgrade APEX Service to Autonomous Transaction Processing.

If you want to expand your capabilities to include the full features of Oracle Autonomous Database with APEX, you can easily upgrade an APEX Service instance to a full Oracle

[1] You can subscribe to information event AJDNonJsonStorageExceeded, to be informed when the 20 GB limit is exceeded. See About Information Events on Autonomous Database.

Autonomous Transaction Processing instance. You can upgrade from either the APEX Instances Details page or the Autonomous Database Details page.

To upgrade to Oracle Autonomous Transaction Processing:

1. From the Oracle Cloud Infrastructure left navigation list, select **Developer Services** and then click **APEX Application Development** or **APEX Instances**.

   The APEX Instances page appears.
2. Under **Filters**, narrow the display by selecting **APEX** from the Type list.
3. On the APEX Instances page, click the APEX instance name.

   The APEX Instance Details page appears.
4. You can upgrade in two ways:
   - On the APEX Instance Details page:
     - Under **APEX instance information**, find **Database type** and click **Edit**.
   - Navigate to the Autonomous Database Details page:
     - Under **APEX instance information**, click the Database name.
     - Under **General information**, find **Workload type** and click **Edit**.

     A dialog appears.
5. Review the information in the dialog and then click **Convert** to convert the database to Autonomous Transaction Processing.

See Autonomous Database Billing Summary for more information.

Clone, Move, or Upgrade an Autonomous Database Instance

Autonomous Database provides cloning where you can choose to create a full clone, create a metadata clone, or create a refreshable clone. You can also move a database to a different compartment. By creating a clone you can upgrade an instance with Oracle Database 19c to create an Oracle Database 23ai clone (either a full clone or a refreshable clone).

- About Cloning on Autonomous Database
  When you create a clone for an Autonomous Database instance, you have the option to select the clone type: a full clone, a metadata clone, or a refreshable clone:
- Prerequisites for Cloning
  Describes prerequisites for cloning an Autonomous Database instance.
- Clone an Autonomous Database Instance
  Shows you the steps to clone an Autonomous Database from the Oracle Cloud Infrastructure Console.
- Clone an Autonomous Database from a Backup
  Shows the options to select a backup as the clone source for cloning Autonomous Database.
- Cross Tenancy and Cross-Region Cloning
  You can clone an Autonomous Database instance from one tenancy, source tenancy, to a different tenancy (destination tenancy).
- Clone Autonomous Database to Change Workload Type
  When you create an Autonomous Database clone, by default the clone is of the same workload type. Optionally, you can select that the clone is of a different workload type than the source database.

- Notes for Cloning Autonomous Database
  Provides information about the cloning operation and the resulting cloned database.
- Move an Autonomous Database to a Different Compartment
  Shows you the steps to move a database to a different Oracle Cloud Infrastructure compartment.
- Upgrade an Autonomous Database Instance to Oracle Database 23ai
  When you clone an Autonomous Database instance and the source database version is Oracle Database 19c, in regions where Oracle Database 23ai is available, you can choose either Oracle Database 19c or Oracle Database 23ai for the clone.

About Cloning on Autonomous Database

When you create a clone for an Autonomous Database instance, you have the option to select the clone type: a full clone, a metadata clone, or a refreshable clone:

- **Full Clone**: creates a new database with the source database's data and metadata.
- **Refreshable Clone**: creates a read-only full clone that can be easily refreshed with the data from the source database.

  See Use Refreshable Clones with Autonomous Database for more information.
- **Metadata Clone**: creates a new database that includes all of the source database schema metadata, but not the source database data.

For a Full Clone or a Metadata Clone, you have the option to select the clone source:

- **Clone from a database instance**: This creates a clone of a running database.
- **Clone from a backup**: This creates a clone when you select a backup from a list of backups, or when you enter a backup point-in-time to clone.
- **Clone from the latest backup**: This creates a clone where Autonomous Database uses the most recent backup data that is available to create the clone.

> **Note:**
>
> Depending on the cloning options you choose, cloning an Autonomous Database instance copies your database files to a new instance. Depending on the size of your database and where you are cloning your database, expect the cloning operation to take significantly longer than provisioning a new Autonomous Database instance. There is no downtime associated with cloning and the cloning operation has no impact on applications running on the source.

Additional Cloning Options

There are several additional cloning options that allow you to clone to a different workload type, a different region, or to a different tenancy.

Cloning Option	Description
Cross-Region Cloning	You can clone a database where the source and the cloned database are in different regions (cross-region cloning). This allows you to clone an existing database or clone a backup to create a database in a different region.

Cloning Option	Description
Clone to Change the Workload Type	You can clone a database and select a different workload type for the cloned database. See Clone Autonomous Database to Change Workload Type for more information.
Source Unavailable Cloning	You can clone a database when the source Autonomous Database instance is in the unavailable state. This allows you to create a new database from a backup and use the new database in place of the unavailable database.
Preview Release Cloning	When an upcoming release is available for preview with Autonomous Database, you can clone your existing database to use the preview version. A preview version is only available when there is an upcoming major version to release. When no preview version is available, cloning does not provide this option.
Cross Tenancy Cloning	You can clone a database where the source and the cloned database are in different tenancies (cross-tenancy cloning). This allows you to clone an existing database or clone a backup to create a database in a different tenancy, where the clone is either in a different tenancy in the same region or in a different tenancy in a different region. **Note:** The cross tenancy cloning option is only available using the CLI or the Autonomous Database REST APIs. This option is not available using the Oracle Cloud Infrastructure Console. See Cross Tenancy and Cross-Region Cloning for more information.
Move to a Different Compartment	You can also move an Autonomous Database to a different Oracle Cloud Infrastructure Compartment. See Move an Autonomous Database to a Different Compartment for more information.

Prerequisites for Cloning

Describes prerequisites for cloning an Autonomous Database instance.

To clone an Autonomous Database instance you must define the required access using OCI Identity and Access Management policy statements written by an administrator, whether you're using the Console, the REST API, the CLI, or another tool.

The following policies allow you to create a clone:

```
Allow group Group_Name to read autonomous-databases in compartment
Compartment_Name
      where target.id = 'oc1.autonomousdatabase.oc1..unique_ID'


Allow group Group_Name to manage autonomous-databases in compartment
Compartment_Name
```

The where clause is optional and provides a more fine grained way to grant access to a specific database.

You can limit cloning permissions so that the group can only clone Autonomous Databases but cannot create Autonomous Databases, or further limit permission to only create a particular type of clone: Full Clone, Metadata Clone, or Refreshable Clone. See IAM Permissions and API Operations for Autonomous Database for more information and examples.

This shows the policies for cloning within a tenancy. See Prerequisites for Cross Tenancy Cloning for cross tenancy cloning policies.

See IAM Policies for Autonomous Database and Getting Started with Policies for more information.

Clone an Autonomous Database Instance

Shows you the steps to clone an Autonomous Database from the Oracle Cloud Infrastructure Console.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- Choose your region. See Switching Regions for information on switching regions and working in multiple regions.
- Choose your **Compartment**. See Compartments for information on using and managing compartments.
- Select an Autonomous Database instance from the list in your compartment.

To clone the Autonomous Database instance:

1. On the **Autonomous Database Details** page, from the **More actions** drop-down list, select **Create clone**.
2. On the **Create Autonomous Database clone** page, choose the clone type from the choices:
   - **Full clone**: creates a new database with the source database's data and metadata.
   - **Refreshable clone**: creates a read-only full clone that can be easily refreshed with the source database's data.
     See Use Refreshable Clones with Autonomous Database for more information.
   - **Metadata clone**: creates a new database with the source database's metadata without the data.
3. In the **Clone source** area, select one of:
   - **Clone from database instance**: This creates a clone from a running database.
   - **Clone from a backup**: This creates a database clone from a selected backup, from a point-in-time timestamp, or from the latest backup data that is available. See Clone an Autonomous Database from a Backup for more information.
4. Provide basic information for the Autonomous Database clone.
   - **Choose your preferred region**: from the list, select the region where you want to create the clone.

Note: the list only shows the regions that you are subscribed to.

For the clone type **Clone from database instance** as the clone source, when you choose another region other than the current region for your clone target, using either the Oracle Cloud Infrastructure CLI or Terraform you can only perform such a cross-region clone from the remote region. That is, call the create clone API from the remote region to which you want to clone, with the source database OCID as that of the source that you want to clone from.

- **Create in Compartment**: See Compartments for information on using and managing compartments.

  Note: the list only shows the compartments that you are subscribed to.

- **Display name**: Specify a user-friendly description or other information that helps you easily identify the resource.

  You can use the name provided, of the form: **Clone-of-***DBname* or change this to the name you want to use to identify the database. The supplied *DBname* is the name of the source database that you are cloning.

- **Database name**: Specify the database name; it must consist of letters and numbers only. The maximum length is 30 characters. The same database name cannot be used for multiple Autonomous Databases in the same tenancy in the same region.

  The default database name is a generated 16-character string.

5. Select a value for **Choose a workload type** from the following options.
   - **Data Warehouse**: This creates an Autonomous Data Warehouse type clone.
   - **Transaction Processing**: This creates an Autonomous Transaction Processing type clone.
   - **JSON**: This creates an Autonomous JSON Database type clone.
   - **APEX**: This creates an APEX type clone.

   > **Note:**
   >
   > The unavailable cloning options are grayed out. See Clone Autonomous Database to Change Workload Type for more information on cross workload cloning.

6. Configure the database (ECPU compute model)
   - **Always Free**: Select to show Always Free options.

     You can only create a free instance in the tenancy's Home region.

   - **Developer**: Select to show Autonomous Database for Developers options.

     This option is only available when the source is an Autonomous Database for Developers instance.

     An Autonomous Database for Developers instance is a fixed shape database that is suited for development and testing use cases.

     Selecting this option sets the **Compute** resources to a fixed value of 4 ECPUs and the **Storage** to 20 GB.

   - **Choose database version**: Select the database version. The available database versions are Oracle Database 23ai and Oracle Database 19c.

     Notes for selecting a database version:

  - In regions where Oracle Database 23ai is not available, Oracle Database 19c is the only choice.
  - Autonomous Database with Oracle Database 23ai in the Paid tier is available in all commercial public cloud regions.
  - Always Free Autonomous Database with Oracle Database 23ai is available in all commercial public cloud regions *except* the following regions: Colombia Central: Bogota (BOG), Saudi Arabia Central (RUH), Singapore West: Singapore (XSP)
  - If the source database version is Oracle Database 19c, in regions where Oracle Database 23ai is available, you can choose either Oracle Database 19c or Oracle Database 23ai for the clone.
  - If the source database version is Oracle Database 23ai, you can select Oracle Database 23ai for the clone.

- **ECPU count**: Specify the number of CPUs for your database. The minimum value for the number of ECPUs is 2.
- **Compute auto scaling**: By default compute auto scaling is enabled to allow the system to automatically use up to three times more CPU and IO resources to meet workload demand. If you do not want to use compute auto scaling then deselect this option.

  See Use Auto Scaling for more information.
- **Storage**: Specify the storage you wish to make available to your database. Depending on your workload type you have these options:
  - **Data Warehouse**: Specify your storage in Terabytes (TB).

    For a Full Clone, the minimum storage that you can specify is the source database's actual used space rounded to the next TB.
  - **Transaction Processing** or **JSON**: Specify your storage in Gigabytes (GB) or Terabytes (TB). Enter the size in the **Storage** field. Select **GB** or **TB** for the **Storage unit size**.

    For a Full Clone, the minimum storage that you can specify is the source database's actual used space rounded to the next GB.
- By default, the IO capacity of your database depends on the number of ECPUs you provision. When you provision 384 TB of storage, your database is entitled to the full IO capacity of the Exadata infrastructure, independent of the number of ECPUs you provision.

  Autonomous Database uses Exadata Smart Flash Cache to automatically cache frequently accessed data, delivering the high I/O rates and fast response times of flash. The amount of flash cache for your database depends on the amount of storage you provision, or the amount of allocated storage if you enable storage auto scaling.

  If you want to provision more than 384 TB of storage, file a Service Request at Oracle Cloud Support.
- **Storage auto scaling**: By default storage auto scaling is disabled. Select if you want to enable storage auto scaling to allow the system to automatically expand to use up to three times more storage.

  With storage auto scaling disabled, the guaranteed minimum flash cache size is 10% of your database's provisioned storage size.

  With storage auto scaling enabled, the guaranteed minimum flash cache size is 10% of your database's provisioned base storage size or its allocated storage size, whichever is higher.

See Use Auto Scaling for more information.

- **Show advanced options**: Click to show additional options:
  - **Enable elastic pool**:

    See Create an Elastic Pool While Provisioning or Cloning an Instance for more information.
  - **Compute model**: Shows the selected compute model.

    See Compute Models in Autonomous Database for more information.
  - **Bring your own license**: If you want to Bring Your Own License to the database, click **Enable** to show the **Update license and Oracle Database Edition** page.

    See Choose Bring Your Own License Option While Provisioning or Cloning for more information.

7. Backup retention

   By default the automatic backup retention is 60 days.

   Automatic backup retention period in days: You have the option to select an automatic backup retention period, in a range from 1 to 60 days. You can restore and recover your database to any point-in-time in this retention period.

   Select **Immutable backup retention** to lock the backup retention period.

   After setting the immutable backup retention option, you cannot disable this option or change the retention period. To disable immutable backup retention or to change the backup retention period, file a Service Request at Oracle Cloud Support.

   See About Backup and Recovery on Autonomous Database for more information.

8. Create administrator credentials
   - **Username** This is a read-only field.
   - **Password** Set the password for the Autonomous Database Admin user. The password must meet the strong password complexity criteria based on Oracle Cloud security standards. For more information on the password complexity rules see Create Users on Autonomous Database - Connecting with a Client Tool.
   - **Confirm password** Specify a value to confirm the password.

9. Choose network access

   > **Note:**
   >
   > After you clone your Autonomous Database you can change the network access option you select for the cloned instance.

   - **Secure access from everywhere**

     By default all secure connections are allowed from everywhere.
   - **Secure access from allowed IPs and VCNs only**

     This option restricts connections to the database according to the access control rules (ACLs) you specify. To add multiple ACLs for the Autonomous Database, click **Add access control rule**.

     See Configure Access Control Lists When You Provision or Clone an Instance for more information.

- **Private endpoint access only**

  This option assigns a private endpoint, private IP, and hostname to your database. Specifying this option allows traffic only from the VCN you specify; access to the database from all public IPs or VCNs is blocked. This allows you to define security rules, ingress/egress, at the Network Security Group (NSG) level and to control traffic to your Autonomous Database.

  See Configure Private Endpoints When You Provision or Clone an Instance for more information.

See IP Address Ranges for information about the public IP address ranges in Oracle Cloud Infrastructure. You must allow traffic to these CIDR blocks to ensure access to an Autonomous Database instance on a public endpoint.

10. (Optional) Provide contacts for operational notifications and announcements

    Click **Add contact** and in the **Contact email** field, enter a valid email address. If the database you are cloning has a customer contact list, the list is copied. To enter multiple **Contact email** addresses, repeat the process to add up to 10 customer contact emails.

    See View and Manage Customer Contacts for Operational Issues and Announcements for more information.

11. (Optional) Click Show advanced options to select advanced options.

    - **Encryption key**
      Encrypt using an Oracle-managed key: By default Autonomous Database uses Oracle-managed encryption keys. Using Oracle-managed keys, Autonomous Database creates and manages the encryption keys that protect your data and Oracle handles rotation of the TDE master key.

      Encrypt using a customer-managed key in this tenancy: If you this option, a master encryption key from a Oracle Cloud Infrastructure Vault in the same tenancy is used to generate the TDE master key on Autonomous Database.

      Encrypt using a customer-managed key located in a remote tenancy: If you this option, a master encryption key in the Oracle Cloud Infrastructure Vault located in a remote tenancy is used to generate the TDE master key on Autonomous Database.

      See Use Customer-Managed Encryption Keys on Autonomous Database for more information.

    - **Maintenance**
      Patch level By default the patch level is the patch level of the source database. Select **Early** to configure the instance with the early patch level. Select **Regular** to configure the instance with the regular patch level.

      See Set the Patch Level for more information.

    - **Management**
      Shows the character set and national character set for your database.

      See Choose a Character Set for Autonomous Database for more information.

    - **Tools**
      If you want to view or customize the tools configuration, select the tools tab.

      See Configure Autonomous Database Built-in Tools when you Provision or Clone an Instance for more information.

    - **Security attributes**
      Add a security attribute to control access for your resources using Zero Trust Packet Routing (ZPR) policies. To enter security attributes during provisioning you must

already have set up security attributes with Zero Trust Packet Routing. You also can add security attributes after provisioning.

> **Note:**
>
> You can apply Oracle Zero Trust Packet Routing (ZPR) policies to a private endpoint.

Specify a **Namespace**, **Key**, and **Value** security attribute.

Click **Add security attribute** to add additional security attributes.

See Overview of Zero Trust Packet Routing for more information.

- **Tags**
  If you want to use Tags, enter the **Tag key** and **Tag value**. Tagging is a metadata system that allows you to organize and track resources within your tenancy. Tags are composed of keys and values which can be attached to resources.

  See Tagging Overview for more information.

12. Click **Create Autonomous Database clone**.

On the Oracle Cloud Infrastructure console the **State** shows Provisioning... until the new database is available.

If you created a cross-region clone, a new tab shows the newly provisioned clone's Oracle Cloud Infrastructure Console and the region shown is the region you selected when you created the clone.

See Notes for Cloning Autonomous Database for additional information on cloning.

See Cloning an Autonomous Database for information on using the API.

Clone an Autonomous Database from a Backup

Shows the options to select a backup as the clone source for cloning Autonomous Database.

Perform the following prerequisite steps as necessary:

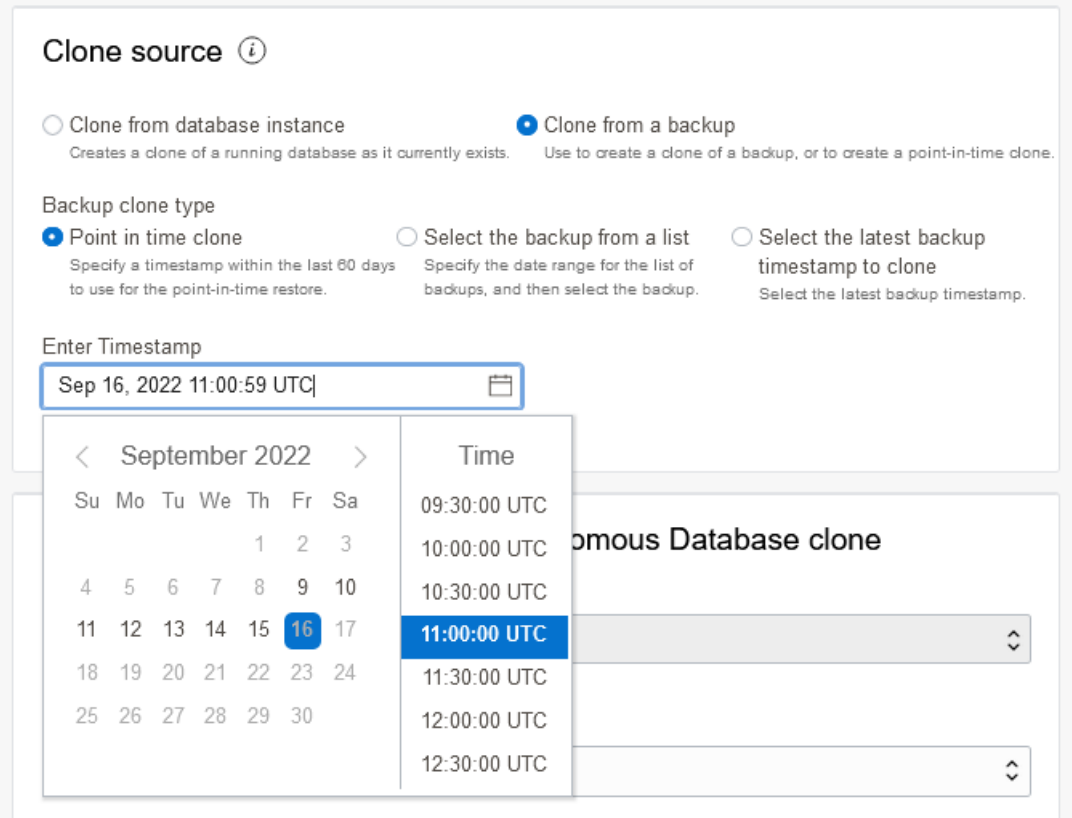
- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.

1. Choose your region. See Switching Regions for information on switching regions and working in multiple regions.
2. Choose your **Compartment**. See Compartments for information on using and managing compartments.
3. Select an Autonomous Database instance from the list in your compartment.
4. On the **Details** page, from the **More actions** drop-down list, select **Create clone**.
5. On the **Create Autonomous Database clone** page, choose the clone type from the choices:
   - **Full clone**: creates a new database with the source database's data and metadata.

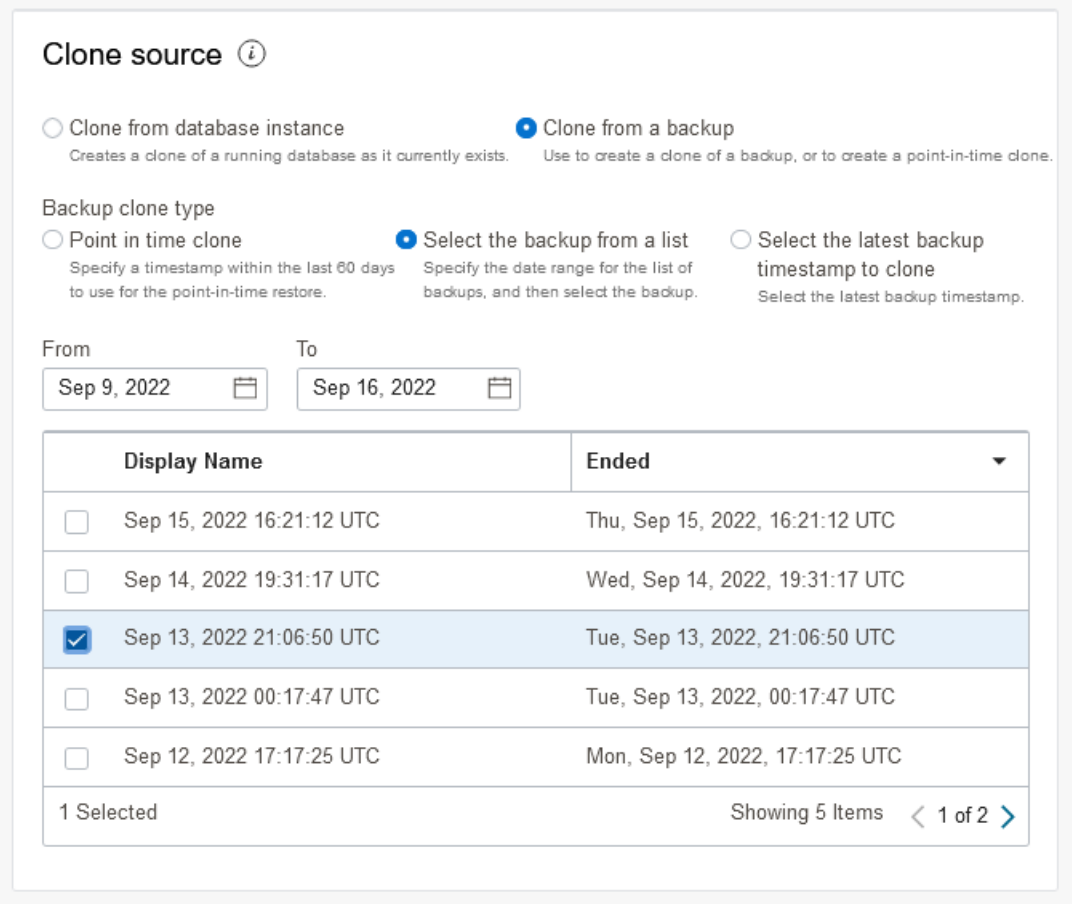
- **Refreshable clone**: creates a read-only full clone that can be easily refreshed with the source database's data.
  See Use Refreshable Clones with Autonomous Database for more information.
- **Metadata clone**: creates a new database with the source database's metadata without the data.

6. In the **Configure clone source** area, select the **Clone source** option:
   - **Clone from database instance**: This creates a clone from a running database. See Clone an Autonomous Database Instance for details and steps with this selection.
   - **Clone from a backup**: This selection creates a database clone using a backup. Select this option.
7. In the **Configure clone source** area, select the **Backup clone type**:
   - **Point in time clone**: Enter a timestamp to clone in the **Enter Timestamp** field.
   - **Select the backup from a list**: Enter **From** and **To** dates to narrow the list of backups then select a backup to use for the clone source.
   - **Latest backup timestamp**: See Clone an Autonomous Database from Latest Backup for information on this clone option.

Point in time clone



Select the backup from a list



8. Provide basic information for the Autonomous Database.
   - **Choose your preferred region**: from the list, select the region where you want to create the clone.

     Note: the list only shows the regions that you are subscribed to.

     For the clone type **Clone from a backup** as the clone source, when you choose another region other than the current region for your clone target, using either the Oracle Cloud Infrastructure CLI or Terraform you can only perform such a cross-region clone from the remote region. That is, call the create clone API from the remote region to which you want to clone, with the source database OCID as that of the source that you want to clone from.
   - **Create in Compartment**: See Compartments for information on using and managing compartments.
   - **Display name**: Specify a user-friendly description or other information that helps you easily identify the resource.

     You can use the name provided, of the form: **Clone-of-***DBname* or change this to the name you want to use to identify the database. The supplied *DBname* is the name of the source database that you are cloning.
   - **Database Name**: Specify the database name; it must consist of letters and numbers only. The maximum length is 30 characters. The same database name cannot be used for multiple Autonomous Databases in the same tenancy in the same region.

     The default database name is a generated 16-character string.

9. Select a value for **Choose a workload type** from the following options.
   - **Data Warehouse**: This creates an Autonomous Data Warehouse type clone.
   - **Transaction Processing**: This creates an Autonomous Transaction Processing type clone.
   - **JSON**: This creates an Autonomous JSON Database type clone.
   - **APEX**: This creates an APEX type clone.

   > **Note:**
   >
   > The unavailable cloning options are grayed out. See Clone Autonomous Database to Change Workload Type for more information on cross workload cloning.

10. Configure the database (ECPU compute model)
    - **Always Free**: Select to show Always Free options.

      You can only create a free instance in the tenancy's Home region.
    - **Choose database version**: Select the database version. The available database versions are Oracle Database 23ai and Oracle Database 19c.

      Notes for selecting a database version:
      - In regions where Oracle Database 23ai is not available, Oracle Database 19c is the only choice.
      - Autonomous Database with Oracle Database 23ai in the Paid tier is available in all commercial public cloud regions.
      - Always Free Autonomous Database with Oracle Database 23ai is available in all commercial public cloud regions *except* the following regions: Colombia Central: Bogota (BOG), Saudi Arabia Central (RUH), Singapore West: Singapore (XSP)
      - If the source database version is Oracle Database 19c, in regions where Oracle Database 23ai is available, you can choose either Oracle Database 19c or Oracle Database 23ai for the clone.
      - If the source database version is Oracle Database 23ai, you can select Oracle Database 23ai for the clone.
    - **ECPU count**: Specify the number of CPUs for your database. The minimum value for the number of ECPUs is 2.
    - **Compute auto scaling**: By default compute auto scaling is enabled to allow the system to automatically use up to three times more CPU and IO resources to meet workload demand. If you do not want to use compute auto scaling then deselect this option.

      See Use Auto Scaling for more information.
    - **Storage**: Specify the storage you wish to make available to your database. Depending on your workload type you have these options:
      - **Data Warehouse**: Specify your storage in Terabytes (TB).

        For a Full Clone, the minimum storage that you can specify is the source database's actual used space rounded to the next TB.

  - **Transaction Processing** or **JSON**: Specify your storage in Gigabytes (GB) or Terabytes (TB). Enter the size in the **Storage** field. Select **GB** or **TB** for the **Storage unit size**.

    For a Full Clone, the minimum storage that you can specify is the source database's actual used space rounded to the next GB.

- By default, the IO capacity of your database depends on the number of ECPUs you provision. When you provision 384 TB of storage, your database is entitled to the full IO capacity of the Exadata infrastructure, independent of the number of ECPUs you provision.

  Autonomous Database uses Exadata Smart Flash Cache to automatically cache frequently accessed data, delivering the high I/O rates and fast response times of flash. The amount of flash cache for your database depends on the amount of storage you provision, or the amount of allocated storage if you enable storage auto scaling.

  If you want to provision more than 384 TB of storage, file a Service Request at Oracle Cloud Support.

- **Storage auto scaling**: By default storage auto scaling is disabled. Select if you want to enable storage auto scaling to allow the system to automatically expand to use up to three times more storage.

  With storage auto scaling disabled, the guaranteed minimum flash cache size is 10% of your database's provisioned storage size.

  With storage auto scaling enabled, the guaranteed minimum flash cache size is 10% of your database's provisioned base storage size or its allocated storage size, whichever is higher.

  See Use Auto Scaling for more information.

- **Show advanced options**: Click to show the compute model options or if you want to create or join an elastic pool:

  - **Enable elastic pool**:

    See Create an Elastic Pool While Provisioning or Cloning an Instance for more information.

  - **Compute model**: Shows the selected compute model.

    See Compute Models in Autonomous Database for more information.

  - **Bring your own license**: If you want to Bring Your Own License to the database, click **Enable** to show the **Update license and Oracle Database Edition** page.

    See Choose Bring Your Own License Option While Provisioning or Cloning for more information.

11. Backup retention

    By default the automatic backup retention is 60 days.

    Automatic backup retention period in days You have the option to select the automatic backup retention period, in a range from 1 to 60 days. You can restore and recover your database to any point-in-time in this retention period.

    Select **Immutable backup retention** to lock the backup retention period.

    After setting the immutable backup retention option, you cannot disable this option or change the retention period. To disable immutable backup retention or to change the backup retention period, file a Service Request at Oracle Cloud Support.

    See About Backup and Recovery on Autonomous Database for more information.

12. Create administrator credentials.
    - **Username** This is a read-only field.
    - **Password** Set the password for the Autonomous Database Admin user. The password must meet the strong password complexity criteria based on Oracle Cloud security standards. For more information on the password complexity rules see Create Users on Autonomous Database - Connecting with a Client Tool.
    - **Confirm password** Specify a value to confirm the password.
13. Choose network access

    > **Note:**
    >
    > After you clone your Autonomous Database you can change the network access option you select for the cloned instance.

    - **Secure access from everywhere**

      By default all secure connections are allowed from everywhere.
    - **Secure access from allowed IPs and VCNs only**

      This option restricts connections to the database according to the access control rules (ACLs) you specify. To add multiple ACLs for the Autonomous Database, select this option and click **Add access control rule**.

      See Configure Access Control Lists When You Provision or Clone an Instance for more information.
    - **Private endpoint access only**

      This option assigns a private endpoint, private IP, and hostname to your database. Specifying this option allows traffic only from the VCN you specify; access to the database from all public IPs or VCNs is blocked. This allows you to define security rules, ingress/egress, at the Network Security Group (NSG) level and to control traffic to your Autonomous Database.

      See Configure Private Endpoints When You Provision or Clone an Instance for more information.

    See IP Address Ranges for information about the public IP address ranges in Oracle Cloud Infrastructure. You must allow traffic to these CIDR blocks to ensure access to an Autonomous Database instance on a public endpoint.
14. (Optional) Provide contacts for operational notifications and announcements

    Click **Add contact** and in the **Contact email** field, enter a valid email address. If the database you are cloning has a customer contact list, the list is copied. To enter multiple **Contact email** addresses, repeat the process to add up to 10 customer contact emails.

    See View and Manage Customer Contacts for Operational Issues and Announcements for more information.
15. (Optional) Click Show advanced options to select advanced options.
    - **Encryption Key**
      Encrypt using an Oracle-managed key: By default Autonomous Database uses Oracle-managed encryption keys. Using Oracle-managed keys, Autonomous Database creates and manages the encryption keys that protect your data and Oracle handles rotation of the TDE master key.

Encrypt using a customer-managed key in this tenancy: If you this option, a master encryption key from a Oracle Cloud Infrastructure Vault in the same tenancy is used to generate the TDE master key on Autonomous Database.

Encrypt using a customer-managed key located in a remote tenancy: If you this option, a master encryption key in the Oracle Cloud Infrastructure Vault located in a remote tenancy is used to generate the TDE master key on Autonomous Database.

See Use Customer-Managed Encryption Keys on Autonomous Database for more information.

- **Maintenance**
  Patch level By default the patch level is the patch level of the source database. Select **Early** to configure the instance with the early patch level. Select **Regular** to configure the instance with the regular patch level.

  See Set the Patch Level for more information.

- **Management**
  Shows the character set and national character set for your database.

  See Choose a Character Set for Autonomous Database for more information.

- **Tools**
  If you want to view or customize the tools configuration, select the tools tab.

  See Configure Autonomous Database Built-in Tools when you Provision or Clone an Instance for more information.

- **Security attributes**
  Add a security attribute to control access for your resources using Zero Trust Packet Routing (ZPR) policies. To enter security attributes during provisioning you must already have set up security attributes with Zero Trust Packet Routing. You also can add security attributes after provisioning.

  > **Note:**
  >
  > You can apply Oracle Zero Trust Packet Routing (ZPR) policies to a private endpoint.

  Specify a **Namespace**, **Key**, and **Value** security attribute.

  Click **Add security attribute** to add additional security attributes.

  See Overview of Zero Trust Packet Routing for more information.

- **Tags**
  If you want to use Tags, enter the **Tag key** and **Tag value**. Tagging is a metadata system that allows you to organize and track resources within your tenancy. Tags are composed of keys and values which can be attached to resources.

  See Tagging Overview for more information.

16. Click **Create Autonomous Database Clone**.

On the Oracle Cloud Infrastructure console the **State** shows Provisioning... until the new database is available.

See Notes for Cloning an Autonomous Database from a Backup for additional information on cloning.

See Cloning an Autonomous Database for information on using the API.

- Clone an Autonomous Database from Latest Backup
  When you choose to clone from latest backup, this selects the latest backup as the clone source. You can choose this option if a database becomes unavailable or for any reason when you want to create a clone based on the most recent backup.

Clone an Autonomous Database from Latest Backup

When you choose to clone from latest backup, this selects the latest backup as the clone source. You can choose this option if a database becomes unavailable or for any reason when you want to create a clone based on the most recent backup.

Follow the steps in Clone an Autonomous Database from a Backup and in the Clone source area, select **Clone from a backup** and **Select the latest backup timestamp to clone** to create a clone that recovers the most recent backup data that is available for the Autonomous Database instance (the clone source) from the database backup and logs backup.

After Autonomous Database has finished provisioning the clone, query the view `dba_pdbs` to see the `last_recover_time` and `last_recover_scn` columns. These columns provide the saved timestamp and the saved SCN of the data from which the cloned database was created.

For example:

```
SELECT last_recover_time, last_recover_scn FROM dba_pdbs;
```

Cross Tenancy and Cross-Region Cloning

You can clone an Autonomous Database instance from one tenancy, source tenancy, to a different tenancy (destination tenancy).

- About Cross Tenancy Cloning
  When you create a cross tenancy clone you can either select that the clone is created in the same region as the source tenancy or in a different region than the source tenancy (cross-region).
- Prerequisites for Cross Tenancy Cloning
  Describes prerequisites for creating a cross tenancy clone where the source database is in one tenancy and the cloned database is in a different tenancy.
- Create a Cross Tenancy or Cross-Region Clone
  Shows the steps to create a cross tenancy clone when the source database and the cloned database are in the same region, or when the source database and the cloned database are in different regions (cross-region).
- Create a Cross Tenancy Clone from a Backup
  Shows the steps to create a cross tenancy clone from a backup.

About Cross Tenancy Cloning

When you create a cross tenancy clone you can either select that the clone is created in the same region as the source tenancy or in a different region than the source tenancy (cross-region).

> **Note:**
>
> The cross tenancy cloning option is only available using the CLI or the Autonomous Database REST APIs. This option is not available using the Oracle Cloud Infrastructure Console.

Note the following for cross tenancy cloning:

- All clone types are supported: the cloned database can be a Full clone, a Metadata clone, or a Refreshable clone.
- A clone can be created from a source Autonomous Database instance or from a backup (using the latest backup, a specified backup, or by selecting a long-term backup).
- The source Autonomous Database instance can use either the ECPU or OCPU compute model. Depending on your workload type, you can clone from a source that uses the OCPU compute model to a clone that uses the ECPU compute model (this is allowed for the Data Warehouse and the Transaction Processing workload types).
- The cloned database can be in the same region or in a different region (cross-region).
- The cross tenancy cloning option does not support cloning with customer managed keys on the source. See Manage Encryption Keys on Autonomous Database for more information on customer managed keys.

Prerequisites for Cross Tenancy Cloning

Describes prerequisites for creating a cross tenancy clone where the source database is in one tenancy and the cloned database is in a different tenancy.

You must run the commands to create a cross tenancy clone on the destination tenancy. Before you create a cross tenancy clone you need to define OCI Identity and Access Management groups and policies on the source tenancy, the tenancy that contains the instance you are cloning, and on the destination tenancy. The groups and policies you define allow you to run commands to create the clone on the destination tenancy and allow the destination tenancy to contact the source tenancy where the source Autonomous Database instance resides.

The OCI Identity and Access Management groups and policies you add support the following:

- A member of a group in the source tenancy allows a group in the destination tenancy to access (read) the source Autonomous Database instance on the source tenancy.

  You do not need to allow other actions on the source Autonomous Database instance (for example, start, stop terminate, or any write operations).
- A member of a group in the destination tenancy is allowed to create a clone in the destination tenancy using the Autonomous Database instance in the source tenancy as the clone source.

  On the destination tenancy you also add a policy that allows a group to manage the Autonomous Database instance on the source tenancy. For example, this policy allows the group to create the clone database, and allows a refreshable clone to run commands that contact the source tenancy, such as **Refresh** and **Disconnect**.

To create a cross tenancy clone use OCI Identity and Access Management to create the required groups and to define the policies that authorize cross tenancy cloning:

1. Create a group on the destination tenancy that contains the user that will be allowed to create a clone.

a. On the destination tenancy, in the Oracle Cloud Infrastructure Console click **Identity & Security**.

b. Under **Identity** click **Domains** and select an identity domain (or create a new identity domain).

c. Under **Identity domain**, click **Groups**.

d. To add a group, click **Create group**.

e. On the Create group page, enter a Name and a Description.

   For example, enter the Name: **DestinationGroup**.

f. On the Create group page, click **Create**.

g. Click **Create** to save the group.

h. On the **Group** page, click **Assign user to groups** and select the users you want to add to the group.

i. Click **Add**.

j. On the **Group** page, from the **Group information** tab copy the OCID for use in Step 2.

2. On the source tenancy, define OCI Identity and Access Management policies for the source Autonomous Database instance.

   a. On the source tenancy, in the Oracle Cloud Infrastructure Console click **Identity & Security**.

   b. Under **Identity**, click **Policies**.

   c. To write a policy, click **Create Policy**.

   d. On the Create Policy page enter a Name and a Description.

   e. On the Create Policy page, select **Show manual editor**.

   f. In the policy builder, add policies so that the group in the destination tenancy is allowed to create a clone using an Autonomous Database instance on the source tenancy as the clone source.

      For example, define the following generic policies:

```
define tenancy DestinationTenancy as ocid1.tenancy.oc1..unique_ID
define group DestinationGroup as ocid1.group.region1..unique_ID
admit group DestinationGroup of tenancy DestinationTenancy to read
autonomous-database-family
       in compartment ocid1.compartment.region1..unique_ID
       where target.id = 'oc1.autonomousdatabase.oc1..unique_ID'
```

      This policy specifies the following:

      - Line 1: the OCID is the OCID of the destination tenancy. This is the tenancy where you are going to create the clone.
      - Line 2: the OCID is the OCID of the group to which the user who will create the clone belongs. This is the OCID you created in Step 1.
      - Line 3: The first OCID is the OCID of the compartment where the source database resides. The second OCID, after the `where` clause, is the OCID of the source Autonomous Database instance.

> **Note:**
>
> The where clause is optional and provides a more fine grained way to grant access to a specific database.

For example, set these policies on the source tenancy to allow cross tenancy cloning:

```
define tenancy DestinationTenancy as
ocid1.tenancy.oc1..aaa_example_rcyx2a
define group DestinationGroup as ocid1.group.oc1..aaa_example_6vctn6xsaq
admit group DestinationGroup of tenancy DestinationTenancy to read
autonomous-database-family in compartment
   ocid1.compartment.region1..bbb_example_rcyx2b where target.id =
'oc1.autonomousdatabase.oc1.aaaabbbbcccc'
```

This policy specifies a user in the `DestinationGroup` of the `DestinationTenancy` can read from a specific Autonomous Database instance in the specified compartment (on the source tenancy). To create a cross tenancy clone the policy only needs to allow read on the source Autonomous Database instance.

g. Click **Create** to save the policy.

3. Define policies on the destination tenancy.

   a. On the destination tenancy, in the Oracle Cloud Infrastructure Console click **Identity & Security**.

   b. Under **Identity**, click **Policies**.

   c. To write a policy, click **Create Policy**.

   d. On the Create Policy page enter a Name and a Description.

   e. On the Create Policy page, select **Show manual editor**.

   f. In the policy builder, add policies so that a group is endorsed to manage Autonomous Databases on the source tenancy.

   For example:

   ```
   Define tenancy SourceTenancy as ocid1.tenancy.oc1..unique_ID
   Endorse group DestinationGroup to manage autonomous-database-family in
   tenancy SourceTenancy
   ```

   This policy specifies the following:

   - Line 1: The OCID is the source tenancy OCID. This is the tenancy where the source Autonomous Database instance resides.
   - Line 2: Specifies that the DestinationGroup group can manage Autonomous Databases in the source tenancy.

   Notes for defining policies on the destination tenancy:

   - For the following policy:

     ```
     Endorse group DestinationGroup to manage autonomous-database-family in
     tenancy SourceTenancy
     ```

This policy allows the group `DestinationGroup` to create Autonomous Databases and Autonomous Database clones in the source tenancy. You can limit cloning permissions so that the group can only clone Autonomous Databases but cannot create Autonomous Databases, or further limit permission to only create a particular type of clone: Full Clone, Metadata Clone, or Refreshable Clone. See IAM Permissions and API Operations for Autonomous Database for more information and examples.

- If these polices are revoked, cross tenancy cloning is no longer allowed.

See Getting Started with Policies for more information.

Create a Cross Tenancy or Cross-Region Clone

Shows the steps to create a cross tenancy clone when the source database and the cloned database are in the same region, or when the source database and the cloned database are in different regions (cross-region).

Cross-region cloning is not supported in the free tier or for Autonomous Database for Developers instances.

These steps cover creating a Full clone or a Metadata clone. See Create a Cross Tenancy or Cross-Region Refreshable Clone for details on creating a cross tenancy refreshable clone.

> **Note:**
>
> The cross tenancy cloning option is only available using the CLI or the Autonomous Database REST APIs. This option is not available using the Oracle Cloud Infrastructure Console.

To create a cross tenancy clone:

1. Perform the prerequisite steps to define the OCI Identity and Access Management policies to authorize cross tenancy cloning.

   See Prerequisites for Cross Tenancy Cloning for details.

2. On the tenancy where you want to create the clone, on the destination tenancy in the destination region, use the CLI or call the REST API with a valid clone type FULL or METADATA and provide the OCID of the source database, where the source database resides in a different tenancy (the source tenancy).

   For example, with the CLI:

   ```
   oci db autonomous-database create-from-clone
         --clone-type metadata
         --compartment-id ocid1.tenancy.oc1..unique_ID
         --source-id ocid1.autonomousdatabase.oc1.iad.unique_ID
         --db-name dbnameclone
         --admin-password password
         --data-storage-size-in-tbs 1
         --compute-model ECPU
         --compute-count 4
   ```

   See create-from-clone for more information.

   Use the `CreateAutonomousDatabase` API to create a cross tenancy clone.

   See the following for additional information on the REST API:

- CreateAutonomousDatabase
- CreateAutonomousDatabaseCloneDetails

- See Notes for Cross Tenancy and Cross-Region Cloning for more information.
- For information about using the API and signing requests, see REST APIs and Security Credentials.
- For information about SDKs, see Software Development Kits and Command Line Interface.

Create a Cross Tenancy Clone from a Backup

Shows the steps to create a cross tenancy clone from a backup.

These steps cover creating a Full clone or a Metadata clone. See Create a Cross Tenancy or Cross-Region Refreshable Clone for details on creating a cross tenancy refreshable clone.

> **Note:**
>
> The cross tenancy cloning option is only available using the CLI or the Autonomous Database REST APIs. This option is not available using the Oracle Cloud Infrastructure Console.

To create a cross tenancy clone from a backup:

1. Perform the prerequisite steps to define the OCI Identity and Access Management policies to authorize cross tenancy cloning.

   See Prerequisites for Cross Tenancy Cloning for details.

2. On the tenancy where you want to create the clone, on the destination tenancy in the destination region, use the CLI or call the REST API with a valid clone type FULL or METADATA and provide the OCID of the backup (on the source tenancy), where the source database resides in a different tenancy (the source tenancy).

   > **Note:**
   >
   > See Create a Cross Tenancy or Cross-Region Refreshable Clone to create a cross tenancy refreshable clone.

   For example, with the CLI:

```
oci db autonomous-database create-from-backup-timestamp
     --autonomous-database-id ocid1.autonomousdatabase.oc1.iad.anuw_example
     --clone-type full
     --compartment-id ocid1.tenancy.oc1..fcue4_example
     --admin-password password
     --compute-model ECPU
     --compute-count 2
     --db-name ExampleTest1
     --timestamp 2023-12-15T19:30:00Z
     --data-storage-size-in-tbs 1
```

See create-from-backup-timestamp and create-from-backup-id for more information.

Use the `CreateAutonomousDatabase` API to create a cross tenancy clone by cloning from a backup of an existing Autonomous Database.

See the following for information on the REST API:

- CreateAutonomousDatabase
- CreateAutonomousDatabaseFromBackupDetails Reference
- CreateAutonomousDatabaseFromBackupTimestampDetails Reference

- For information about using the API and signing requests, see REST APIs and Security Credentials.
- For information about SDKs, see Software Development Kits and Command Line Interface.

Clone Autonomous Database to Change Workload Type

When you create an Autonomous Database clone, by default the clone is of the same workload type. Optionally, you can select that the clone is of a different workload type than the source database.

Cloning a database to a different workload type is useful when:

- You have provisioned a database of an incorrect workload type and now need to clone the database to the appropriate workload type.
- You need a copy of your database of a different workload type to branch out to run new types of workloads on your data.

Note the following when you clone to a different workload type:

- You can clone between the following workload types:
  - Autonomous Data Warehouse to Autonomous Transaction Processing
  - Autonomous Transaction Processing to Autonomous Data Warehouse
  - Autonomous JSON Database to Autonomous Transaction Processing or Autonomous Data Warehouse or APEX Service
  - APEX Service to Autonomous JSON Database or Autonomous Transaction Processing or Autonomous Data Warehouse
- You may also clone from a backup, when choosing to clone a database to a different workload type.
- The cloned database can be in the same region or in a different region (cross-region).
- Cloning a database to a different workload type does not affect the database's existing data dictionary, optimizer statistics, and object compression. New objects that are created after cloning are created using the cloned database's workload type parameters (for example compression and statistics).
- When you clone from Autonomous Data Warehouse to Autonomous Transaction Processing, the cloned database has the additional services TP and TPURGENT.
- When you clone from one of the following to Autonomous Data Warehouse the cloned database does not have the services TP or TPURGENT:
  - Autonomous Transaction Processing
  - Autonomous JSON Database

  - APEX Service

The following are not allowed when you clone and select a different workload type:

- Creating a cross workload refreshable clone.
- Cloning to a more restrictive workload type, such as from Autonomous Transaction Processing or Autonomous Data Warehouse database to Autonomous JSON Database or APEX Service.

See the following for steps to create an Autonomous Database clone:

- Clone an Autonomous Database Instance
- Clone an Autonomous Database from a Backup

Notes for Cloning Autonomous Database

Provides information about the cloning operation and the resulting cloned database.

- General Notes for Cloning on Autonomous Database
  Provides general information about the cloning operation and the resulting cloned database.
- Notes for Cloning an Autonomous Database from a Backup
- Notes for Cross Tenancy and Cross-Region Cloning
  Provides information about cross tenancy and cross-region cloning.
- Resource Management Rules and Performance Data for a Cloned Database
  Provides notes for resource management rules and performance data for a cloned database.
- Optimizer Statistics for a Cloned Database
  During the provisioning for either a Full Clone or a Metadata Clone, the optimizer statistics are copied from the source database to the cloned database.
- Disable Oracle Scheduler Jobs for a Cloned Database
  Oracle Scheduler jobs from a source database are copied to the clone when you perform a clone operation. You can assure that cloned Oracle Scheduler jobs are disabled on the clone.

General Notes for Cloning on Autonomous Database

Provides general information about the cloning operation and the resulting cloned database.

- If the source database version is Oracle Database 19c, in regions where Oracle Database 23ai is available, you can choose either Oracle Database 19c or Oracle Database 23ai for the clone.
- If the source database version is Oracle Database 23ai, you can select Oracle Database 23ai for the clone.
- If there is an ongoing scaling operation on a source database, you cannot initiate a clone operation until the ongoing operation completes. If you attempt such an operation you will see message such as the following:

  ```
  The operation cannot be performed because the Autonomous Database with Id
  **** is in the SCALE_IN_PROGRESS state.
  ```

- If you define a network Access Control List (ACL) on the source database, the currently set network ACL is cloned to the new database. If a database is cloned from a backup, the

current source database's ACL is applied (not the ACL that was valid at the time of the backup).

- If you create a clone and the source database has an access control list (ACL) and you specify the private endpoint network access option, **Virtual cloud network** for the target database, the ACL is not cloned to the new database. In this case, you must define security rules within your Network Security Group (or groups) to control traffic to and from your target database (instead of using the access control rules that were specified in the ACL on the clone source). See Configure Private Endpoints When You Provision or Clone an Instance for more information.
- Cloning an Autonomous Database instance copies your database files to a new instance. There is no downtime associated with cloning and the cloning operation has no impact on applications running on the source.
- For a **Metadata Clone**, the APEX Apps and the OML Projects and Notebooks are copied to the clone. For a **Metadata Clone**, the underlying database data of the APEX App or OML Notebook is not cloned.
- The Autonomous Database Details page for an Autonomous Database instance that was created by cloning includes the **Cloned From** field. This displays the name of the database where the clone was created.
- The value of the `DBA_SERVICES.FAILOVER_TYPE` column is not copied to a cloned Autonomous Database instance.

Notes for Cloning an Autonomous Database from a Backup

Provides Notes for cloning an Autonomous Database from a backup:

- If there is an ongoing clone from backup operation on a source database, you cannot initiate a new clone operation on the same backup being cloned until the ongoing operation completes. Thus, you cannot clone from backup twice concurrently from a specific backup (for example, a specific timestamp or a specific selected backup from the list of backups).
- For external tables, partitioned external tables, and the external partitions of hybrid partitioned tables a backup does not include the external files that reside on your Object Store. Thus, for the clone from backup operation, it is your responsibility to backup, and restore if necessary, the external files associated with external tables, external partitioned tables, or the external files for a hybrid partitioned table
- With clone from backup, the Oracle Machine Learning workspaces, projects, and notebooks of the source database are not cloned to the new database.
- After Autonomous Database has finished provisioning the clone, query the view `dba_pdbs` to see the `last_recover_time` and `last_recover_scn` columns. These columns provide the saved timestamp and the saved SCN of the data from which the cloned database was created.

  For example:

```
SELECT last_recover_time, last_recover_scn FROM dba_pdbs;
```

Notes for Cross Tenancy and Cross-Region Cloning

Provides information about cross tenancy and cross-region cloning.

- For cloning, when you choose a region other than the current region for your clone target, using either the Oracle Cloud Infrastructure CLI or Terraform you can only perform such a cross-region clone from the remote region. That is, call the create clone API from the

remote region to which you want to clone, with the source database OCID as that of the source that you want to clone from.

- The cross tenancy cloning option is only available using the CLI or the Autonomous Database REST APIs. Cross tenancy cloning is not available using the Oracle Cloud Infrastructure Console.

  See Cross Tenancy and Cross-Region Cloning for more information.

- The cross tenancy cloning option does not support cloning with customer managed keys on the source. See Manage Encryption Keys on Autonomous Database for more information on customer managed keys.

Resource Management Rules and Performance Data for a Cloned Database

Provides notes for resource management rules and performance data for a cloned database.

The following applies for resource management rules and performance data in a cloned database:

- During the provisioning for either a Full Clone or a Metadata Clone, any resource management rule changed by the user in the source database is carried over to the cloned database.
- For a cloned database, performance data for the time prior to the clone operation is not visible in the **Database Dashboard** card (under Monitor in Database Actions).

For more information on setting resource management rules, see Manage Runaway SQL Statements on Autonomous Database.

Optimizer Statistics for a Cloned Database

During the provisioning for either a Full Clone or a Metadata Clone, the optimizer statistics are copied from the source database to the cloned database.

The following applies for optimizer statistics for tables in a cloned database:

- Full Clone: loads into tables behave the same as loading into a table with statistics already in place.
- Metadata Clone: the first load into a table after the database is cloned clears the statistics for that table and updates the statistics with the new load.

For more information on Optimizer Statistics, see Optimizer Statistics Concepts.

Disable Oracle Scheduler Jobs for a Cloned Database

Oracle Scheduler jobs from a source database are copied to the clone when you perform a clone operation. You can assure that cloned Oracle Scheduler jobs are disabled on the clone.

To disable Oracle Scheduler jobs on Autonomous Database instances that are cloned, you can define a trigger using the `AFTER CLONE` clause. To use the `AFTER CLONE` clause, define the trigger on the source database before you perform the clone operation.

For example:

```
CREATE OR REPLACE TRIGGER after_clone_instance
     AFTER CLONE
     ON PLUGGABLE DATABASE
BEGIN
-- Disable specific jobs or use a cursor to grab all scheduled jobs and
```

```
disable them
     DBMS_SCHEDULER.DISABLE(name=> job_name);
END;
/
```

See CREATE TRIGGER Statement for more information.

Note the following:

- The value for the `DBA_SCHEDULER_WINDOWS.REPEAT_INTERVAL` column is not copied to a cloned Autonomous Database instance.

Move an Autonomous Database to a Different Compartment

Shows you the steps to move a database to a different Oracle Cloud Infrastructure compartment.

- To move a database you must have the right to manage Autonomous Databases in the database's current compartment and in the compartment you are moving it to.
- As soon as you move a database to a different compartment, the policies that govern the new compartment apply immediately and affect access to the database. Therefore, your access to the database may change, depending on the policies governing your Oracle Cloud user account's access to resources.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.

To move a database to a different compartment:

1. Choose your region. See Switching Regions for information on switching regions and working in multiple regions.
2. Choose your **Compartment**. See Compartments for information on using and managing compartments.
3. Select an Autonomous Database instance from the list in your compartment.
4. On the **Details** page, from the **More actions** drop-down list, select **Move resource**.
5. In the **Move Autonomous Database to a different compartment** page, select the new compartment.
6. Click **Move Autonomous Database**.

See Moving Database Resources to a Different Compartment for more information.

Upgrade an Autonomous Database Instance to Oracle Database 23ai

When you clone an Autonomous Database instance and the source database version is Oracle Database 19c, in regions where Oracle Database 23ai is available, you can choose either Oracle Database 19c or Oracle Database 23ai for the clone.

- Upgrade to Oracle Database 23ai by Creating a Full Clone
  If your Autonomous Database instance is using Oracle Database 19c, in regions where Oracle Database 23ai is available, you can upgrade to Oracle Database 23ai by creating a full clone.
- Upgrade to Oracle Database 23ai by Creating a Refreshable Clone
  You can create a refreshable clone where the source database version is Oracle Database 19c and the cloned database version is Oracle Database 23ai.

Upgrade to Oracle Database 23ai by Creating a Full Clone

If your Autonomous Database instance is using Oracle Database 19c, in regions where Oracle Database 23ai is available, you can upgrade to Oracle Database 23ai by creating a full clone.

To upgrade your Autonomous Database by creating a full clone, do the following:

1. Verify the Oracle Database version for your Autonomous Database instance is Oracle Database 19c.

   On the Oracle Cloud Infrastructure Console, the Autonomous Database details page includes the **Database version** field showing your database version.
2. During the cloning operation, in the Clone type area select **Full clone**, and in the Clone source area, select **Clone from database instance**.
3. During the cloning operation, in the cloning step where you configure the database, select Oracle Database 23ai in the **Choose database version** selection.

   Oracle Database 23ai only shows in regions where Oracle Database 23ai is available.

See Clone an Autonomous Database Instance for the detailed steps for this type of clone.

> **Note:**
>
> The clone has its own wallet and you cannot use the wallet of the source database to connect to the clone. All users must obtain a new wallet and use the new wallet to access the upgraded database. See Download Client Credentials (Wallets) for more information.

Upgrade to Oracle Database 23ai by Creating a Refreshable Clone

You can create a refreshable clone where the source database version is Oracle Database 19c and the cloned database version is Oracle Database 23ai.

Creating a refreshable clone with an Oracle Database 19c source and an Oracle Database 23ai refreshable clone allows you to test your workload with Oracle Database 23ai, or you can use the refreshable clone to upgrade to Oracle Database 23ai.

To create this type of refreshable clone, start with an Autonomous Database Database version 19c instance and follow the steps to create a refreshable clone, and select Database version 23ai for the clone. See Create a Refreshable Clone for an Autonomous Database Instance for more information.

> **Note:**
>
> An Oracle Database 19c source with an Oracle Database 23ai refreshable clone is especially useful when you want to use Capture-Replay to test your workloads running on Oracle Database 19c before you upgrade to Oracle Database 23ai. See Test 19c Workloads on a on 23ai Refreshable Clone with Capture-Replay for more information.

Note the following when the refreshable clone source Autonomous Database Database version is 19c and the cloned Autonomous Database Database version is 23ai:

- After you create the refreshable clone, the clone will be in the stopped state and is not available for read-only operations. The only operations available for this type of refreshable clone are refresh (either manual or automatic), and disconnect.

  See Refresh a Refreshable Clone on Autonomous Database and Disconnect a Refreshable Clone from the Source Database for more information.
- This type of refreshable clone only allows a refresh when the database is in the "Stopped" state. The stopped refresh applies only for refreshable clones that have a source Oracle Database 19c and the cloned database is Oracle Database 23ai.
- To use the refreshable clone you must disconnect the clone from the source database and start the database.

  See Start Autonomous Database and Disconnect a Refreshable Clone from the Source Database for more information.
- You can reconnect a disconnected refreshable within 24 hours. When the source Oracle Database 19c and the refreshable clone Oracle Database 23ai, the clone reverts to the stopped state when you reconnect and you may refresh data from the source again.

  See Reconnect a Refreshable Clone to the Source Database for more information.
- While the refreshable clone is connected to the source database, if no refresh is performed within 7 days, the refreshable clone can no longer be refreshed. In this case you have the following option:
  - You can disconnect the refreshable clone from the source database. After the disconnect operation completes you can use the database as a standalone database and the data reflects the state of the source database at the time of the last successful refresh.

    See Disconnect a Refreshable Clone from the Source Database for more information.
- Creating a cross-region refreshable clone with an Oracle Database 19c source and an Oracle Database 23ai refreshable clone is not supported.
- Creating a cross tenancy refreshable clone with an Oracle Database 19c source and an Oracle Database 23ai refreshable clone is not supported in either the same region or cross-region.

Use Refreshable Clones with Autonomous Database

Autonomous Database provides cloning where you can choose to create a full clone of the active instance, create a metadata clone, or create a refreshable clone. With a refreshable clone the system creates a clone that can be easily updated with changes from the source database.

For details on creating a clone that is a full clone or a metadata clone, see Cloning and Moving an Autonomous Database.

- About Refreshable Clones on Autonomous Database
  When you create a refreshable clone for an Autonomous Database instance the system clones the source database to a database that can be refreshed from the source.
- Prerequisites for Creating a Refreshable Clone
  Describes prerequisites for creating a refreshable clone.
- Create and Manage Refreshable Clones
  When you create a refreshable clone the system creates a clone that can be easily updated with changes from the source database.
- View Refreshable Clones for an Autonomous Database Instance
- Use the API
  Provides details for using the APIs to create, refresh, list, disconnect, and reconnect a refreshable clone.
- Refreshable Clone Notes
  Lists limitations and notes for Autonomous Database refreshable clones.

About Refreshable Clones on Autonomous Database

When you create a refreshable clone for an Autonomous Database instance the system clones the source database to a database that can be refreshed from the source.

Using Cloud Links you can specify that access to a data set from one or more databases is offloaded to a refreshable clone. When a consumer Autonomous Database is listed in a data set's offload list, access to the data set is directed to the refreshable clone. See Register a Data Set with Offload Targets for Data Set Access for more information.

You can create a Refreshable Clone with a Oracle Database 19c source where the Refreshable Clone is Oracle Database 23ai. This type of Refreshable Clone is especially useful when you want test workloads running on Oracle Database 19c and run the workload on Oracle Database 23ai before you upgrade to Oracle Database 23ai. See Test 19c Workloads on a on 23ai Refreshable Clone with Capture-Replay for more information.

- Refreshable Clone Features
  Describes refreshable clone features.
- About Refreshable Clone Billing
  Provides information about refreshable clone billing.
- Refreshable Clone Operations
  You can create a refreshable clone from an Autonomous Database instance. After you create a refreshable clone you can perform several operations on the refreshable clone, including: refresh, stop, start, restart, disconnect from source, and terminate.
- Refreshable Clone with Automatic Refresh Enabled
  By default you manually refresh a refreshable clone with changes from the source database. When you enable the automatic refresh option, Autonomous Database automatically refreshes a refreshable clone with data from the source database at given time intervals.
- Refreshable Clone Lifecycle States
  After you create a refreshable clone, the clone indicates its state on the Autonomous Database Information page in the Lifecycle state field. In addition, the Mode field indicates that a refreshable clone is **Read-Only**.

- Refreshable Clone Refresh Timing and Disconnecting from the Source Database
  A refreshable clone with **Automatic refresh** disabled has a one week refresh age limit and a banner on the Oracle Cloud Infrastructure console displays the date and time up to which you can refresh the refreshable clone. The banner also includes a **Refresh clone** button.
- Refreshable Clone Reconnect to the Source Database
  After you perform a **Disconnect refreshable clone** operation, a banner displays the date and time up to which you can reconnect the database to the source. The banner also includes a **Reconnect refreshable clone** button.
- Operations on an Autonomous Database with an Attached Refreshable Clone
  Describes details for using a source Autonomous Database instance that has one or more attached refreshable clones.

Refreshable Clone Features

Describes refreshable clone features.

A refreshable clone allows you to do the following:

- Maintain one or more copies of the source database for use as read-only databases. A clone database is available when you need it, and when you want to update the data you can refresh the clone from the source database.
- Create one or more clones in regions other than the region of your primary (source) database. Clones in remote regions can be refreshed from the source database.
- You can enable the refreshable clone automatic refresh option. With automatic refresh enabled Autonomous Database automatically refreshes the clone with data from the source. See Refreshable Clone with Automatic Refresh Enabled for more information.
- Share copies of a production database with multiple business units. For example, one business unit might use the source database for ongoing transactions and another business unit could at the same time use the refreshable clone database for read-only operations.

  This option also allows you to spread the cost of database usage across multiple business units. You can bill the different units separately, based on their usage of one or more refreshable clone databases.
- Use a refreshable clone as a test database. You can disconnect a refreshable clone from its source and perform DML operations or calculations as needed, in addition to querying data. This allows you to run DML and make changes while the database is disconnected. When you are done with your testing you can reconnect to the source database, which refreshes the clone to the point where it was when you disconnected.

  The reconnect operation is only available for 24 hours after the disconnect time. If you do not reconnect within the reconnect period, the clone is disconnected from the source database and refreshing and reconnecting are not possible.

> **Note:**
>
> Refreshable clones have a one week refresh age limit. If you do not perform a refresh within a week (7 days), the refreshable clone automatically disconnects from its source database. The refresh by date is shown in the Oracle Cloud Infrastructure Console banner. See Disconnect a Refreshable Clone from the Source Database for more information.
> See Upgrade to Oracle Database 23ai by Creating a Refreshable Clone for additional restrictions for the case when the source is Database version 19c and the refreshable clone is Database version 23ai.

About Refreshable Clone Billing

Provides information about refreshable clone billing.

See Oracle Autonomous Database Serverless Features Billing for information on refreshable clone billing.

Refreshable Clone Operations

You can create a refreshable clone from an Autonomous Database instance. After you create a refreshable clone you can perform several operations on the refreshable clone, including: refresh, stop, start, restart, disconnect from source, and terminate.

Operation	Description
Create	You can create a refreshable clone from an Autonomous Database instance. You can create more than one refreshable clone using the same Autonomous Database instance as a source. See Create a Refreshable Clone for an Autonomous Database Instance for the steps to create a refreshable clone.
View	You view a refreshable clone from the Oracle Cloud Infrastructure console Autonomous Database Details page. See View Refreshable Clones for an Autonomous Database Instance for more information.
Start or Restart	When a refreshable clone is stopped as indicated by the Lifecycle state **Stopped**, you can start the database. When a refreshable clone is available as indicated by the Lifecycle state **Available**, you can restart the database. Start and Restart are not allowed when a refreshable clone source is Database version 19c and the refreshable clone is Database version 23ai. See Upgrade to Oracle Database 23ai by Creating a Refreshable Clone for more information.
Refresh	For a refreshable clone, you can refresh the clone with data from the source database. See Refresh a Refreshable Clone on Autonomous Database for more information.
Edit Automatic Refresh Policy	When you enable the automatic refresh option a refreshable clone automatically refreshes from the source database at regular intervals. By default, automatic refresh is disabled and you must manually refresh at least once every 7 days. See Edit Automatic Refresh Policy for Refreshable Clone for more information.
Disconnect Clone from Source	You can disconnect a refreshable clone from the source database to make the clone a standard read/write database. See Disconnect a Refreshable Clone from the Source Database for more information.

Operation	Description
Reconnect Refreshable Clone	When a clone database is disconnected, you can use the clone as a standard read/write database. There is a 24 hour reconnect period where you can reconnect the clone to its source database. After the reconnect period the refreshable clone is disassociated from the source database and reconnecting to the source database is not allowed. See Reconnect a Refreshable Clone to the Source Database for more information.
Stop	When a refreshable clone is stopped, database operations are not available and charging for OCPU usage on the refreshable clone stops.
Terminate	If you want to terminate a refreshable clone, select **More actions** and **Terminate**. Terminating a refreshable clone disassociates the clone database from the source database.

Refreshable Clone with Automatic Refresh Enabled

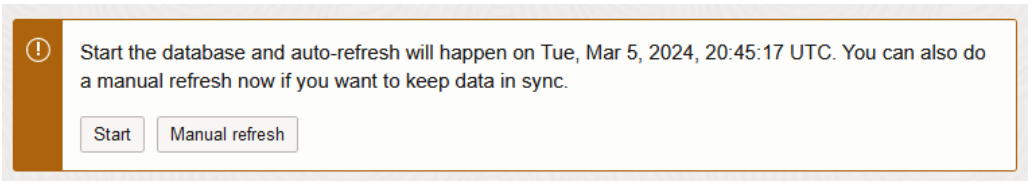
By default you manually refresh a refreshable clone with changes from the source database. When you enable the automatic refresh option, Autonomous Database automatically refreshes a refreshable clone with data from the source database at given time intervals.

> **Note:**
>
> The automatic refresh option is only available when the source database uses the ECPU compute model.

Note the following for a refreshable clone with automatic refresh enabled:

- You can manually refresh a refreshable clone between automatic refreshes. If you manually refresh to a refresh point later than the next refresh point specified for an automatic refresh, the automatic refresh may fail. Following refreshes, after the failed refresh, are attempted at the next scheduled automatic refresh.
- If a refreshable clone is stopped, disconnected, or already at a later timestamp than the specified automatic refresh point, the automatic refresh fails.
- If a scheduled refresh is missed due to the source database being stopped, or the refreshable clone being disconnected, or when the refreshable clone was manually refreshed to later point, the missed refresh is skipped. The refreshable clone refreshes at the next scheduled interval.
- When an automatic refresh is missed, Autonomous Database does the following:
  - Shows a banner on the Oracle Cloud Infrastructure Console.

    For example:

  - Generates the automatic refresh failed event. See Autonomous Database Lifecycle Events for more information.

- If no refresh is performed within 7 days, the refreshable clone automatically disconnects from its source database.

  See Disconnect a Refreshable Clone from the Source Database for more information.

- See Upgrade to Oracle Database 23ai by Creating a Refreshable Clone for additional restrictions for the case when the source is Database version 19c and the refreshable clone is Database version 23ai.

Refreshable Clone Lifecycle States

After you create a refreshable clone, the clone indicates its state on the Autonomous Database Information page in the Lifecycle state field. In addition, the Mode field indicates that a refreshable clone is **Read-Only**.

A refreshable clone indicates its state as follows:

- **Updating**: When a refreshable clone is refreshing or reconnecting, the Lifecycle state field shows **Updating**. While the database is refreshing connections and queries wait until the refresh completes. After the refresh completes the state is set to **Available**, and connections and queries resume.

  See Refresh a Refreshable Clone on Autonomous Database for more information.

- **Stopped**: When a refreshable clone is stopped, database operations are not available and charging for compute usage on the refreshable clone stops.

- **Available**: When the refreshable clone is available, database operations are available and you are charged for compute usage on the refreshable clone.

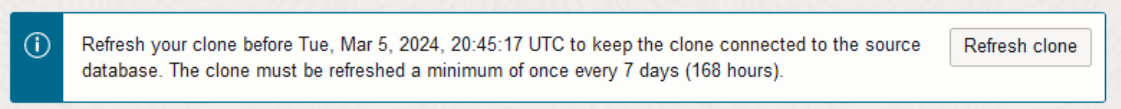
The Autonomous Database Information page Mode field indicates the database mode, as follows:

Read-Only: No data can be inserted into or updated in a refreshable clone as it is a read-only database. You can use a refreshable clone for read-only queries and for reporting.

See Disconnect a Refreshable Clone from the Source Database to change the database to Read/Write mode. In this case the refreshable clone is disconnected from the source database.

Refreshable Clone Refresh Timing and Disconnecting from the Source Database

A refreshable clone with **Automatic refresh** disabled has a one week refresh age limit and a banner on the Oracle Cloud Infrastructure console displays the date and time up to which you can refresh the refreshable clone. The banner also includes a **Refresh clone** button.



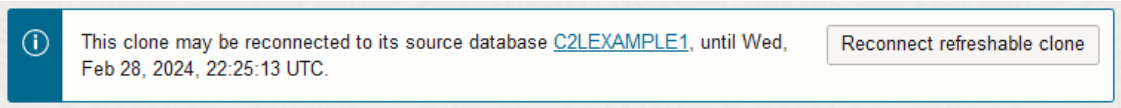
Refreshable clones have a one week refresh age limit. If you do not perform a refresh within a week (7 days), the refreshable clone automatically disconnects from its source database. The refresh by date is shown in the Oracle Cloud Infrastructure Console banner. See Disconnect a Refreshable Clone from the Source Database for more information.

When a refreshable clone has exceeded the maximum refresh time and is automatically disconnected, if you want to use a refreshable clone that can be refreshed from the source database, you must create a new refreshable clone.

See Upgrade to Oracle Database 23ai by Creating a Refreshable Clone for additional restrictions for the case when the source is Database version 19c and the refreshable clone is Database version 23ai.

Refreshable Clone Reconnect to the Source Database

After you perform a **Disconnect refreshable clone** operation, a banner displays the date and time up to which you can reconnect the database to the source. The banner also includes a **Reconnect refreshable clone** button.



When a disconnected database is not reconnected within 24 hours from the disconnect time, the Oracle Cloud Infrastructure Console removes the reconnect refreshable clone banner.

When a disconnected refreshable clone has exceeded the reconnect period, you have the following options:

- You can use the database as a standard Autonomous Database; there is no longer an option to reconnect the database to the source database.
- If you want to use a refreshable clone that can be refreshed from the source database, then you must create a new refreshable clone. If you create a new refreshable clone, then you might want to terminate the disconnected clone that is no longer able to refresh from the source database.

Operations on an Autonomous Database with an Attached Refreshable Clone

Describes details for using a source Autonomous Database instance that has one or more attached refreshable clones.

When you make certain changes on a source Autonomous Database instance that has one or more refreshable clones attached to it, the changes are applied to both the source database and to the refreshable clones as follows:

- Storage: The storage value you set on the source database applies to both the source database and to any attached refreshable clones.
- ADMIN password: The ADMIN password value you set on the source database applies to both the source database and to any attached refreshable clones.

To view the refreshable clones for a source database, on the Autonomous Database Details page, under **Resources**, click **Refreshable Clones**. The Autonomous Database resources area provides a link to each refreshable clone in the Display Name field, and includes the Last Refresh timestamp field and the Refresh Point timestamp field. The refresh point specifies the timestamp for the source database data to which the refreshable clone data is refreshed.

If you want to terminate a source database that has one or more attached refreshable clones, then before you terminate the source database you must do the following until there are no longer any attached refreshable clones. For each attached refreshable clone, do one of the following:

- Disconnect the refreshable clone from the source database. See Disconnect a Refreshable Clone from the Source Database for more information.
- Terminate the refreshable clone to disassociate the refreshable clone from the source database. You can terminate a refreshable clone by selecting **More actions** and **Terminate**.

Prerequisites for Creating a Refreshable Clone

Describes prerequisites for creating a refreshable clone.

To create a refreshable clone you must define the required access using OCI Identity and Access Management policy statements written by an administrator, whether you're using the Console, the REST API, the CLI, or another tool.

The following policies allow you to create a refreshable clone:

```
Allow group Group_Name to read autonomous-databases in compartment
Compartment_Name
      where target.id = 'oc1.autonomousdatabase.oc1..unique_ID'


Allow group Group_Name to manage autonomous-databases in compartment
Compartment_Name
```

The `where` clause is optional and provides a more fine grained way to grant access to a specific database.

You can limit cloning permissions so that the group can only clone Autonomous Databases but cannot create Autonomous Databases, or further limit permission to only create a particular type of clone: Full Clone, Metadata Clone, or Refreshable Clone. See IAM Permissions and API Operations for Autonomous Database for more information and examples.

This shows the policies for cloning within a tenancy. See Prerequisites for Cross Tenancy Cloning for cross tenancy cloning policies.

See IAM Policies for Autonomous Database and Getting Started with Policies for more information.

Create and Manage Refreshable Clones

When you create a refreshable clone the system creates a clone that can be easily updated with changes from the source database.

- Create a Refreshable Clone for an Autonomous Database Instance
  Shows you the steps to create an Autonomous Database refreshable clone from the Oracle Cloud Infrastructure Console.
- Edit Automatic Refresh Policy for Refreshable Clone
  When you enable the automatic refresh option a refreshable clone automatically refreshes from the source database at regular intervals. By default, automatic refresh is disabled and you must manually refresh at least once every 7 days.
- Refresh a Refreshable Clone on Autonomous Database
  Shows you the steps to refresh a refreshable clone from the Oracle Cloud Infrastructure Console.
- Disconnect a Refreshable Clone from the Source Database
- Reconnect a Refreshable Clone to the Source Database
- Create a Cross Tenancy or Cross-Region Refreshable Clone

Create a Refreshable Clone for an Autonomous Database Instance

Shows you the steps to create an Autonomous Database refreshable clone from the Oracle Cloud Infrastructure Console.

See Prerequisites for Creating a Refreshable Clone for required refreshable clone prerequisites.

Perform the following steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To create a refreshable clone do the following:

1. On the **Details** page, from the **More actions** drop-down list, select **Create clone**.
2. On the **Create Autonomous Database clone** page, choose the clone type **Refreshable clone** from the choices:
   - **Full clone**: creates a new database with the source database's data and metadata.
   - **Refreshable clone** To create a refreshable clone, select this clone type.
   - **Metadata clone**: creates a new database with the source database's metadata without the data.
3. Provide basic information for the Autonomous Database clone.
   - **Choose your preferred region**. Use the current region or select a different region. When you create a refreshable clone, the preferred region list only shows a remote region if your tenancy is subscribed to the remote region (you must be subscribed to the remote region where you create a refreshable clone).

     The list of available regions only shows a remote region if your tenancy is subscribed to the remote region (you must be subscribed to the paired remote region for the target region where you are cloning your database). See Autonomous Database Cross-Region Paired Regions for more information.
   - **Create in compartment**. See Compartments for information on using and managing compartments.
   - **Source database name** This field is read-only and shows the name of the source database.
   - **Display name** Specify a user-friendly description or other information that helps you easily identify the database.

     You can use the name provided, of the form: **Clone-of-***DBname* or change this to the name you want to use to identify the database. The supplied *DBname* is the name of the source database that you are cloning.
   - **Database name**: Specify the database name; it must consist of letters and numbers only. The maximum length is 30 characters. The same database name cannot be used for multiple Autonomous Databases in the same tenancy in the same region.

     The default database name is a generated 16-character string.

4. Configure the database (with ECPU compute model)

   - **Choose database version**: Select the database version. The available database versions are Oracle Database 23ai and Oracle Database 19c.

     Notes for selecting a database version:

     – If the source database version is Oracle Database 23ai, you can select Oracle Database 23ai for the clone.

     – If the source database version is Oracle Database 19c, in regions where Oracle Database 23ai is available, you can choose either Oracle Database 19c or Oracle Database 23ai for the clone.

     – Autonomous Database with Oracle Database 23ai in the Paid tier is available in all commercial public cloud regions.

     – An Oracle Database 19c source with an Oracle Database 23ai refreshable clone is especially useful when you want to use Capture-Replay to test your workloads running on Oracle Database 19c before you upgrade to Oracle Database 23ai. See Test 19c Workloads on a on 23ai Refreshable Clone with Capture-Replay for more information.

   - **ECPU count**: Specify the number of ECPUs for your database. The minimum value for the number of ECPUs is 2.

   - **Compute auto scaling**: By default compute auto scaling is enabled to allow the system to automatically use up to three times more CPU and IO resources to meet workload demand. If you do not want to use compute auto scaling then deselect this option.

     See Use Auto Scaling for more information.

   - **Show advanced options**: Click to show the compute model options or if you want to create or join an elastic pool:

     – **Enable elastic pool**:

       See Create an Elastic Pool While Provisioning or Cloning an Instance for more information.

     – **Compute model**: Shows the selected compute model.

       See Compute Models in Autonomous Database for more information.

     – **Bring your own license**: If you want to Bring Your Own License to the database, click **Enable** to show the **Update license and Oracle Database Edition** page.

       See Choose Bring Your Own License Option While Provisioning or Cloning for more information.

> **Note:**
>
> The storage for a refreshable clone is set to the same size as on the source database. To change the storage size for a refreshable clone, you must change the storage value on the source database.

5. Select **Enable automatic refresh** to specify that the refreshable clone automatically refreshes. By default, **Enable automatic refresh** is deselected and you must manually refresh at least once every 7 days. When you select **Enable automatic refresh**, the dialog shows two fields:

- **Frequency of refresh**: specifies the refresh frequency in hours or days. The minimum is one hour and the maximum is 7 days. The default **Frequency of refresh** value is 1 hour.
- **Data lag**: specifies the data lag in minutes, hours, or days, the minimum is 0 minutes and the maximum is 7 days. This is the value that specifies the amount of time behind the source that the data refresh is, where a value of 0 specifies that the refreshable clone refreshes to the latest available timestamp. The default **Data lag** value is 0.

6. Choose network access

> **Note:**
>
> After you clone your Autonomous Database you can change the network access option you select for the cloned instance.

- **Secure access from everywhere**

  By default all secure connections are allowed from everywhere.

- **Secure access from allowed IPs and VCNs only**

  This option restricts connections to the database according to the access control rules (ACLs) you specify. To add multiple ACLs for the Autonomous Database, select this option and click **Add access control rule**.
  See Configure Access Control Lists When You Provision or Clone an Instance for more information.

- **Private endpoint access only**

  This option assigns a private endpoint, private IP, and hostname to your database. Specifying this option allows traffic only from the VCN you specify; access to the database from all public IPs or VCNs is blocked. This allows you to define security rules, ingress/egress, at the Network Security Group (NSG) level and to control traffic to your Autonomous Database.

  See Configure Private Endpoints When You Provision or Clone an Instance for more information.

See IP Address Ranges for information about the public IP address ranges in Oracle Cloud Infrastructure. You must allow traffic to these CIDR blocks to ensure access to an Autonomous Database instance on a public endpoint.

7. (Optional) Provide contacts for operational notifications and announcements

   Click **Add contact** and in the **Contact email** field, enter a valid email address. If the database you are cloning has a customer contact list, the list is copied. To enter multiple **Contact email** addresses, repeat the process to add up to 10 customer contact emails.

   See View and Manage Customer Contacts for Operational Issues and Announcements for more information.

8. (Optional) Click Show advanced options to select advanced options.

   - **Encryption key**
     A refreshable clone uses the encryption key of the source Autonomous Database.

     See Use Customer-Managed Encryption Keys on Autonomous Database for more information.

   - **Maintenance**

Patch level By default the patch level is the patch level of the source database. Select **Early** to configure the instance with the early patch level. Select **Regular** to configure the instance with the regular patch level.

See Set the Patch Level for more information.

- **Management**
  Shows the character set and national character set for your database.

  See Choose a Character Set for Autonomous Database for more information.

- **Tools**
  Refreshable clones inherit their tool status for each built-in tool from the refreshable clone's source database. If the source of a refreshable clone changes its tool configuration, the change is reflected in the refreshable clone after the next refresh.

  See Configure Autonomous Database Built-in Tools for more information.

- **Tags**
  If you want to use Tags, enter the **Tag key** and **Value**. Tagging is a metadata system that allows you to organize and track resources within your tenancy. Tags are composed of keys and values which can be attached to resources.

  See Tagging Overview for more information.

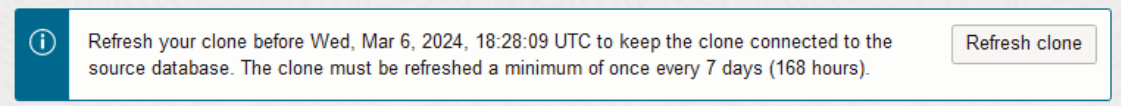
9. Click **Create Autonomous Database clone**.

On the Oracle Cloud Infrastructure console the **State** shows **Provisioning...** until the refreshable clone is available.

After the provisioning completes the Lifecycle state shows **Available** and the Mode is **Read-only**.

After you download the wallet for the refreshable clone and connect to the database, you can begin using the database to perform read-only operations such as running queries or building reports and notebooks.

When **automatic refresh** is disabled, after you create a refreshable clone, the Oracle Cloud Infrastructure console shows a banner similar to the following with a message indicating the date prior to which the next refresh must be completed (the banner shows the 7 day refresh limit).



See Refresh a Refreshable Clone on Autonomous Database for details on refreshing a refreshable clone.

Edit Automatic Refresh Policy for Refreshable Clone

When you enable the automatic refresh option a refreshable clone automatically refreshes from the source database at regular intervals. By default, automatic refresh is disabled and you must manually refresh at least once every 7 days.

The edit auto refresh settings dialog allows you to enable or disable automatic refresh for a refreshable clone. When automatic refresh is enabled you can change the automatic refresh options.

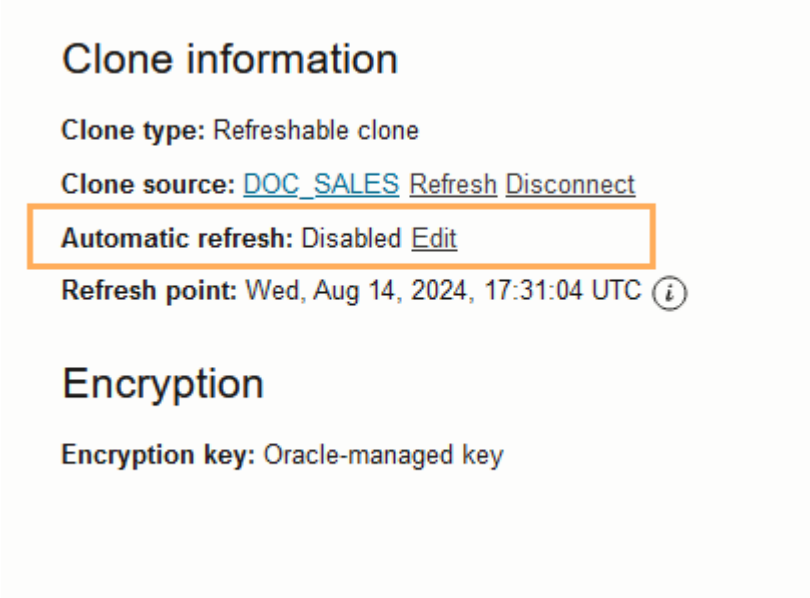
> **Note:**
>
> The automatic refresh option is only available when the source database uses the ECPU compute model.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To edit auto refresh options:

1. On the clone source, on the **Details** page, under **Resources**, select **Refreshable clones**.
2. In the Refreshable clones area, under **Display name**, select a refreshable clone.
3. On the refreshable clone, under the **Clone information** area, in the **Automatic refresh** field click **Edit**.



This shows the Edit automatic refresh policy page.

4. Select or deselect **Enable automatic refresh** to enable or disable automatic refresh for the refreshable clone. When **Enable automatic refresh** is selected, you can modify the options:
   - **Frequency of refresh**: specifies the refresh frequency in hours or days. The minimum is one hour and the maximum is 7 days. The default **Frequency of refresh** value is 1 hour.
   - **Data lag**: specifies the data lag in minutes, hours, or days, the minimum is 0 minutes and the maximum is 7 days. This is the value that specifies the amount of time behind the source that the data refresh is, where a value of 0 specifies that the refreshable clone refreshes to the latest available timestamp. The default **Data lag** value is 0.

5. Click **Save**.

Subscribe to the Information event `com.oraclecloud.databaseservice.autonomous.database.information` to be notified about the completion or failure of a scheduled Automatic Refresh. See Information Events on Autonomous Database for more information.

Refresh a Refreshable Clone on Autonomous Database

Shows you the steps to refresh a refreshable clone from the Oracle Cloud Infrastructure Console.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

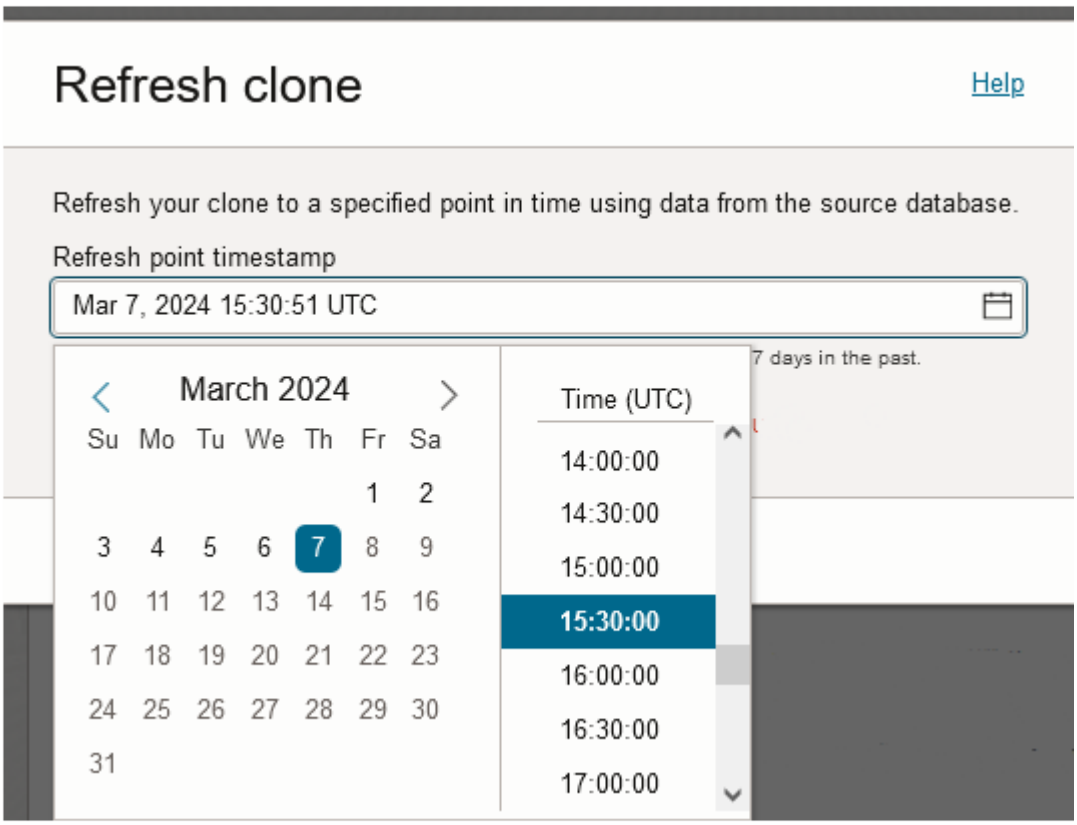
You can refresh a refreshable clone as follows:

1. On the Details page, under Clone information, in the **Clone source** field click the **Refresh** link.

   This shows the Refresh clone dialog.

2. In the Refresh clone dialog, specify the **Refresh point timestamp**.

   The refresh point specifies the timestamp of the source database data to which the refreshable clone is refreshed.



The timestamp can be a minimum of 1 minute in the past and a maximum of 7 days in the past. The timestamp must be later than, that is after, the last refresh point. Thus, if you create a refreshable clone, then you would need to wait some time before you refresh the clone.

3. Click **Refresh clone**.

On the Oracle Cloud Infrastructure console the Lifecycle state shows **Updating** until the refresh operation completes. While the database is updating connections and queries wait until the refresh completes.

You can see the status of the refresh operation under **Work requests**.

On the Oracle Cloud Infrastructure Console, under Clone Information the **Refresh point** shows the timestamp of the source database's data that the clone was refreshed to.

> **Note:**
>
> Refreshable clones have a one week refresh age limit. If you do not perform a refresh within a week (7 days), the refreshable clone automatically disconnects from its source database. The refresh by date is shown in the Oracle Cloud Infrastructure Console banner. See Disconnect a Refreshable Clone from the Source Database for more information.

Disconnect a Refreshable Clone from the Source Database

Shows you the steps to disconnect a refreshable clone from the source database.

When you disconnect a refreshable clone the refreshable clone is disassociated from the source database. This converts the database from a refreshable clone to a regular database.

Following the disconnect operation you are allowed to reconnect the disconnected database to the source database. The reconnect operation is limited to a 24 hour period.

There are several ways to find refreshable clones. See View Refreshable Clones for an Autonomous Database Instance for more information.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

Starting from the Source Database Console Disconnect a refreshable clone from the source database and convert the refreshable clone to a read/write database:

1. On the **Details** page, under **Resources** select **Refreshable clones**.

   This shows the list of refreshable clones.

2. In the row for the refreshable clone you want to disconnect, click ⋮ at the end of the row and select **Disconnect clone from source database**.
3. In the Disconnect refreshable clone dialog enter the source database name to confirm disconnecting the clone.
4. Click **Disconnect refreshable clone**.

Starting from the Clone Database Console Disconnect a refreshable clone from the source database and convert the refreshable clone to a read/write database:

1. On the Details page, under Clone information, in the **Clone source** field click the **Disconnect** link.
2. In the Disconnect refreshable clone dialog enter the source database name to confirm disconnecting the clone.
3. Click **Disconnect refreshable clone**.

The Autonomous Database Lifecycle state of the cloned database changes to **Updating**. When the disconnect operation completes the Lifecycle state changes to **Available** and the Mode shows **Read/Write**.

After you disconnect a refreshable clone, the Oracle Cloud Infrastructure Console updates with the following changes:

- The **Display name** on the Autonomous Databases page updates and does not show the **Refreshable clone** indicator for the cloned instance.
- On the disconnected instance, the Autonomous Database Details page updates and does not show the **Refreshable clone** indicator, and the **Clone information** area is removed.
- On the disconnected instance, the Autonomous Database details page displays a banner indicating the date and time up to which you can reconnect the database to the source. The banner also includes a **Reconnect refreshable clone** button. For example:



See Reconnect a Refreshable Clone to the Source Database for more information.

- On the source database, on the **Autonomous Database Details** page, when you click **Refreshable clones** under **Resources**, the list no longer shows an entry for the disconnected clone.

Notes for disconnecting a refreshable clone.

- Disconnecting a refreshable clone from the source may fail if the source database has scaled down its storage and the refreshable clone has a larger amount of data than the source. In this case, you have the following options:
  - You may scale up the storage on the source database temporarily before disconnecting, and then scale the source back down after disconnecting.
  - Refresh the clone to a point where the scale down of storage occurred.
- After you disconnect a refreshable clone you have 24 hours to reconnect. After the reconnect period is over, the reconnect operation is not available. If you do not reconnect a disconnected refreshable clone, the clone is a standard Autonomous Database and there is no longer an option to reconnect the database to the source database.
- A disconnected refreshable clone is no longer associated with the source database. To use the database or to initiate the reconnect operation, you must know the name of the refreshable clone database that was disconnected from the source database. You must initiate the reconnect operation from the disconnected clone database. You cannot reconnect a disconnected clone from the source database.

Reconnect a Refreshable Clone to the Source Database

Shows you the steps to reconnect a disconnected clone to the source database.

After you disconnect a refreshable clone, during the following 24 hour reconnect period you can reconnect to the source database. The reconnect operation restores all the data back to that time when the clone was disconnected.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

Reconnect a database and restore all the data back to that time when the refreshable clone was disconnected from the source database as follows:

1. On the **Details** page, from the **More actions** drop-down list, select **Reconnect refreshable clone**.

   This shows the Reconnect Refreshable Clone dialog.
2. In the Reconnect Refreshable Clone dialog, enter the source database name to confirm.

> **Note:**
>
> All data and metadata inserted, updated, or deleted from the database where you run the reconnect operation will be lost when the database is reconnected to the source database.

3. Click **Reconnect refreshable clone**.

The Autonomous Database Lifecycle state changes to **Updating**.

When the reconnect refreshable clone operation completes, the Oracle Cloud Infrastructure Console on the clone indicates the change as follows:

- The Mode shows **Read Only**.
- If the clone was available when you initiated the reconnect operation, the Lifecycle state changes to **Available**.
- If the clone was stopped when you initiated the reconnect operation, the Lifecycle state shows **Stopped**.
- On the Autonomous Database Details page, the Clone Information area is updated. The **Clone type** shows Refreshable Clone, the **Clone source** shows the source database link, the **Refresh** link, **Disconnect** link, and the **Refresh point** shows the last refresh timestamp.

Notes for reconnecting a disconnected clone to the source database:

- If you disconnect a refreshable clone, you have 24 hours to use the reconnect operation. After the reconnect period is over, the reconnect operation is not available. If you do not reconnect a disconnected refreshable clone, the Autonomous Database is a standard read/write database and there is no longer an option to reconnect the database to the source database.
- Network ACLs and Private Endpoint network configuration options on a refreshable clone can be changed when the clone is disconnected from the source database. When you reconnect a refreshable clone to its source database, the reconnect operation does not restore the network ACLs or Private Endpoint network configuration options on the refreshable clone.

Create a Cross Tenancy or Cross-Region Refreshable Clone

You can create a refreshable clone in a different tenancy (destination tenancy). When you create a cross tenancy refreshable clone, the refreshable clone in the destination tenancy can be in the same region or in a different region (cross-region).

The cross tenancy cloning option is only available using the CLI or the Autonomous Database REST APIs. This option is not available using the Oracle Cloud Infrastructure Console.

These steps cover creating a Refreshable clone. See Create a Cross Tenancy or Cross-Region Clone for details on creating a cross tenancy Full clone or Metadata clone.

To create a cross tenancy refreshable clone:

1. Perform the prerequisite steps to define the OCI Identity and Access Management policies to authorize cross tenancy cloning.

   See Prerequisites for Cross Tenancy Cloning for details.

> **Note:**
>
> If the polices you specify to allow cross tenancy cloning are subsequently revoked, cross-tenancy cloning is no longer allowed. In addition, for a refreshable clone, when the policies are revoked actions on the remote tenancy that require contact with the source tenancy, such as disconnect and refresh will fail.

2. On the tenancy where you want to create the refreshable clone, use the CLI or call the REST API.

   You run the CLI or call the REST API on the destination tenancy in the destination region where you want to create the refreshable clone. The source database resides in a different tenancy, and if specified in a different region. This creates either a cross-tenancy refreshable clone in the same region, or a cross-tenancy refreshable clone in a different region (cross-region).

   For example, with the CLI:

   ```
   oci db autonomous-database create-refreshable-clone
       --compartment-id ocid1.tenancy.oc1..aaaaaaaafcue47pqmrf4vigne_example
       --source-id ocid1.autonomousdatabase.oc1.iad.anuwcljs4bv3yyiae2_example
       --db-name adatabasedb1
       --data-storage-size-in-tbs 1
       --compute-model ECPU
       --compute-count 2
       --refreshable-mode MANUAL
   ```

   See create-refreshable-clone for more information.

   Use the `CreateAutonomousDatabase` API to create a cross tenancy refreshable clone. When you call this API, to create a refreshable clone you set the `source` value to `CLONE_TO_REFRESHABLE` and the `sourceID` to the OCID of the source database.

   See the following for information on the REST API:

   - CreateAutonomousDatabase
   - CreateRefreshableAutonomousDatabaseCloneDetails

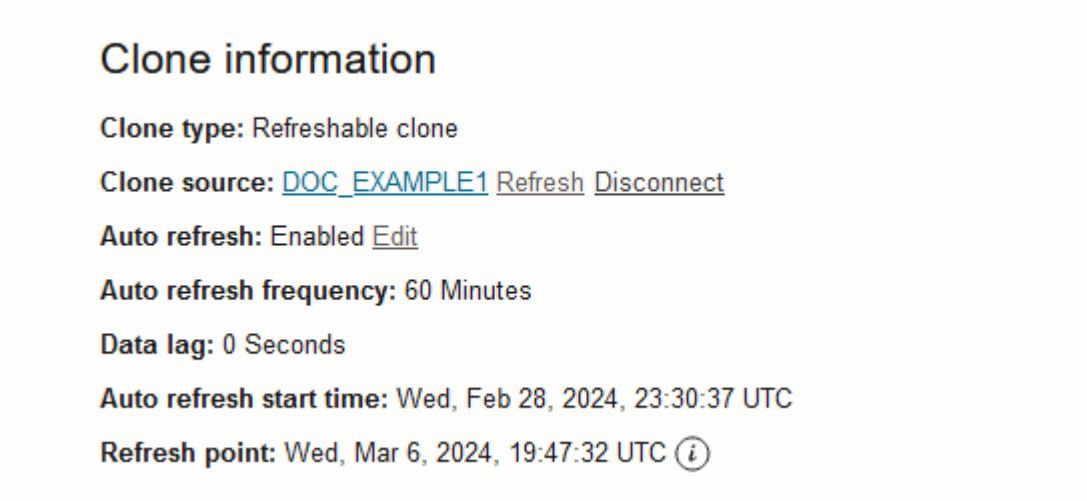
- For information about using the API and signing requests, see REST APIs and Security Credentials.
- For information about SDKs, see Software Development Kits and Command Line Interface.

View Refreshable Clones for an Autonomous Database Instance

If you know a refreshable clone's display name, you can view the instance by selecting from the list of Autonomous Databases on the Oracle Cloud Infrastructure Console.

On the Oracle Cloud Infrastructure Console, the details page for a refreshable clone shows a **Refreshable clone** tag to indicate the instance is a refreshable clone.

The **Clone information** area shows details for a refreshable clone. You can access the source database console by clicking the link in the **Clone source** field.



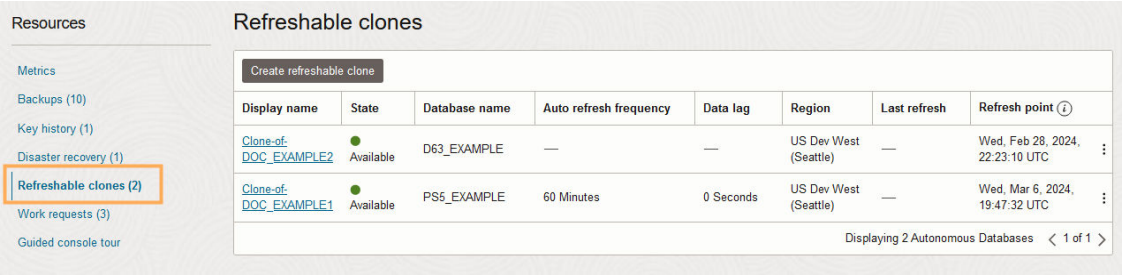
On the details under **Clone information** the console shows the **Refresh point**, indicating the timestamp of the source database's data that the clone was last refreshed with.

If an Autonomous Database has attached refreshable clones, you can view the refreshable clones as follows:

1. Choose your region. See Switching Regions for information on switching regions and working in multiple regions.
2. Choose your **Compartment**. See Compartments for information on using and managing compartments.
3. Select an Autonomous Database instance from the list in your compartment.
4. On the **Autonomous Database details** page, under **Resources**, click **Refreshable clones**.

This shows the list of refreshable clones for the Autonomous Database instance.

If you click  at the end of a row for a refreshable clone in the list, you can select an action to perform for the refreshable clone.



Use the API

Provides details for using the APIs to create, refresh, list, disconnect, and reconnect a refreshable clone.

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.

Use these API operations to manage a refreshable clone:

- To create a refreshable clone, use CreateAutonomousDatabase.
- To refresh a refreshable clone, use AutonomousDatabaseManualRefresh.
- To list the attached refreshable clones for a database, use ListAutonomousDatabaseClones.
- To disconnect a refreshable clone from the source database, use UpdateAutonomousDatabase.

Refreshable Clone Notes

Lists limitations and notes for Autonomous Database refreshable clones.

- See Upgrade to Oracle Database 23ai by Creating a Refreshable Clone for additional notes when you select a refreshable clone source with Database version 19c and a cloned Autonomous Database Database version 23ai.
- Always Free Autonomous Databases do not support refreshable clones.
- Autonomous Database does not support using customer-managed encryption keys with a refreshable clone. You cannot create a refreshable clone from a source database that uses customer-managed encryption keys. Additionally, you cannot switch to a customer-managed encryption key on a source database that has one or more refreshable clones.
- You cannot create a cascading series of refreshable clones. Thus, a refreshable clone cannot be created from another refreshable clone.
- You cannot backup or restore a refreshable clone.
- The ADMIN password for a refreshable clone is inherited from the source database. If you want to change the ADMIN password for a refreshable clone you must change the ADMIN password on the source database, and then refresh the clone for the ADMIN password on the clone to come into effect.
- Oracle APEX URLs do not work in a refreshable clone read-only database and the APEX URLs are disabled in the Oracle Cloud Infrastructure Console and in the Database Actions APEX and APEX Workspaces cards. APEX URLs are enabled for a read/write database when a refreshable clone is disconnected from the source.
- Oracle Machine Learning is disabled in a refreshable clone read-only database. The OML User Administration URLs are disabled in the Oracle Cloud Infrastructure Console and on the Database Actions Launchpad.
- Oracle Data Safe is not supported for a refreshable clone instance. The Data Safe data in the refreshable clone's source database, for example user registration, data masking, and so on are available and Data Safe can be enabled when a refreshable clone is disconnected from the source.
- When you scale up or scale down the storage on the source Autonomous Database instance for a refreshable clone, the change is immediately reflected in the database console and in the billing for the refreshable clone. When the refreshable clone is refreshed to a refresh point after the scale up or down operation, the system makes a corresponding change to the refreshable clone storage (scaling up or scaling down the refreshable clone storage to match the source database).
- You cannot use the rename operation on a refreshable clone instance or on a database that is the source for a refreshable clone.
- When the patch level of a source Autonomous Database instance is **Regular**, in regions that support the **Early** patch level you can set the patch level of a clone to **Early**. See Set the Patch Level for more information.

- Automatic Workload Repository (AWR) data and reports are not available for refreshable clones. In addition, the graphs that rely on AWR data are not available, including the following graphs:
  - The Running SQL statements graph on the Overview tab shown on the **Database Dashboard** card in Database Actions.
  - The SQL statement response time graph on the Overview tab shown on the **Database Dashboard** card in Database Actions.
  - The Time period graphs on the Monitor tab shown on the **Database Dashboard** card in Database Actions.
  - The Performance Hub graph data older than one hour is not available.
- Creating a cross-region refreshable clone with an Oracle Database 19c source and an Oracle Database 23ai refreshable clone is not supported.
- Creating a cross tenancy refreshable clone with an Oracle Database 19c source and an Oracle Database 23ai refreshable clone is not supported in either the same region or cross-region.
- For an Oracle Autonomous JSON Database (workload type **JSON Database**), note the following when reconnecting to the source database:
  - If, after you disconnect the refreshable clone, you promote both the clone and the source to Oracle Autonomous Transaction Processing (workload type **Transaction Processing**), you can reconnect the database to the source.
  - If after you disconnect the refreshable clone, you promote the source database to Oracle Autonomous Transaction Processing (workload type **Transaction Processing**) and do not promote the disconnected clone, the disconnected clone must also be promoted to Oracle Autonomous Transaction Processing (workload type **Transaction Processing**) before you perform the reconnect operation.
  - If after you disconnect the refreshable clone, you promote the disconnected database to Oracle Autonomous Transaction Processing (workload type **Transaction Processing**), you can still reconnect to the source but the reconnected database remains in the promoted state.

  See Promote to Autonomous Transaction Processing for more information.

Manage Autonomous Database Built-in Tools

Autonomous Database includes built-in tools that you can enable and disable when you provision or clone a database, or at any time for an existing database.

- About Autonomous Database Built-in Tools
  Provides a complete list of Autonomous Database built-in tools and a description of each tool.
- About Configuring Built-in Tools Compute Resources and Limits
- View Autonomous Database Built-in Tools Status
  Describes the steps to view the status of built-in tools on your Autonomous Database instance.
- Configure Autonomous Database Built-in Tools
  You can enable or disable Autonomous Database built-in tools for an existing Autonomous Database instance. In addition, with the ECPU compute model you can configure limits for the number of ECPUs and the idle time allowed for several of the built-in tools.

- Configure Autonomous Database Built-in Tools when you Provision or Clone an Instance
  During provisioning or when you clone an instance, you can enable or disable Autonomous Database built-in tools.
- Notes for Autonomous Database Built-in Tools
  Provides additional Autonomous Database built-in tools notes and lists differences that apply depending on the workload type of your Autonomous Database instance.

About Autonomous Database Built-in Tools

Provides a complete list of Autonomous Database built-in tools and a description of each tool.

Tool	Description
Oracle APEX	Oracle APEX is a low-code development platform that you can use to build scalable, secure enterprise applications that can be deployed anywhere. See Create Applications with Oracle APEX in Autonomous Database for more information.
Database Actions	Oracle Database Actions is a web-based interface that uses Oracle REST Data Services to provide development, data tools, administration and monitoring features for Oracle Autonomous Database. See Connect with Built-In Oracle Database Actions for more information.
Graph Studio	Graph Studio automates the creation of knowledge (RDF) and property graphs and includes interactive tooling for query, analysis, and visualization of these graphs in the Autonomous Database. You must log in as a graph-enabled user to access Graph Studio. Create this user in Database Actions. See Using Oracle Graph with Autonomous Database for more information.
Oracle ML User Interface	Oracle Machine Learning User Interface provides immediate access to the Oracle Machine Learning components and functionality on Autonomous Database, including OML Notebooks, OML AutoML UI, OML Models, and template example notebook. This includes: • OML Notebooks • OML AutoML UI • OML for Python • OML for R • OML Services
Data Transforms	Oracle Data Transforms allows you to design graphical data transformations in the form of data flows and workflows. The data flows define how the data is moved and transformed between different systems, while the workflows define the sequence in which the data flows run.
Web Access (ORDS)	Oracle REST Data Services (ORDS) provides HTTPS interfaces for working with the contents of your Oracle Database in one or more REST enabled schema See Develop with Oracle REST Data Services on Autonomous Database for more information.
MongoDB API	Oracle Database API for MongoDB enables MongoDB-compatible clients and drivers to connect directly to Autonomous Database. See Using Oracle Database API for MongoDB for more information.

Tool	Description
SODA Drivers	Simple Oracle Document Access (SODA) is a set of APIs that let you work with JSON documents managed by the Oracle Database without needing to use SQL. SODA drivers are available for REST, Java, Node.js, Python, PL/SQL, and C. See Use SODA for REST with Autonomous Database for more information.

About Configuring Built-in Tools Compute Resources and Limits

When your Autonomous Database instance is using the ECPU compute model, the built-in tools that use a VM have default compute resource limits. Optionally you can configure the compute resource limits for these tools.
For the built-in tools that use a VM, the **ECPU count** and **Maximum idle time** specify resources and configuration options for the VMs that run the associated built-in tool, and mean the following in terms of resource usage and billing:

- You do not pay for a tool's ECPU allocation if you do not use the tool.
- The VMs associated with a tool are provisioned when you start using a tool. For example, if Graph Studio is disabled, billing does not begin when you enable the tool. Billing begins when you start using Graph Studio.
- The **ECPU count** specifies the number of ECPUs that are dedicated to the tool. The built-in tools **ECPU count** value applies in addition to the database **ECPU count** you specify for the instance.
- After you start using a tool you are billed per ECPU hour reserved from the time the built-in Tool is launched. ECPU usage is measured each second, in units of whole ECPUs and averaged across an hour. If your database is running for less than an hour or you change the number of ECPUs for only part of an hour, your built-in Tool ECPU usage is billed per second for the average ECPU consumption during that hour. The minimum ECPU consumption is one minute.
- Billing stops for a built-in tool's allocated ECPUs when the built-in tool is disabled, the instance stops or is terminated, or when the built-in tool is idle for more than the specified **Maximum idle time**.

See Oracle Autonomous Database Serverless Billing for Database Tools for more information.

View Autonomous Database Built-in Tools Status

Describes the steps to view the status of built-in tools on your Autonomous Database instance.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

1. Navigate to the Autonomous Database details page.
2. Select the **Tool configuration** tab.

This displays the tools configuration details and status.

If your Autonomous Database instance is on a private endpoint and you enable the **Allow public access** option, the tools tab shows both a **Private access URL** and a **Public access URL**. The Private access URL provides the URL to access a tool from within the VCN where the database resides. The Public access URL provides the URL to access a tool from the public internet, to provide access from allowed public IPs or from allowed public VCNs. See Use a Private Endpoint with Public Access Allowed for more information.

Configure Autonomous Database Built-in Tools

You can enable or disable Autonomous Database built-in tools for an existing Autonomous Database instance. In addition, with the ECPU compute model you can configure limits for the number of ECPUs and the idle time allowed for several of the built-in tools.

- Configure Autonomous Database Built-in Tools (ECPU compute model)
- Configure Autonomous Database Built-in Tools (OCPU compute model)

Configure Autonomous Database Built-in Tools (ECPU compute model)

Describes how to enable or disable Autonomous Database built-in tools for an existing Autonomous Database instance.

Tool status is retained across starts and restarts for an Autonomous Database instance. For example, if a tool is disabled before you stop or restart an instance, the tool maintains the same status (disabled) after a start or restart.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

1. On the Autonomous Database details page, select the **Tool configuration** tab.
2. Select **Enable** to enable a tool.
3. (Optional) When enabled, Graph Studio, Oracle Machine Learning User Interface, and Data Transforms provide default **ECPU count** and **Max idle count** values and let you configure these values.

   Change the **ECPU count** value if you want to provide more or less resources for a built-in tool.

   Change the **Max idle time** value to specify the time, in minutes, after which an idle Graph Studio, Data Transforms, or Oracle Machine Learning VM is terminated.

   See About Configuring Built-in Tools Compute Resources and Limits for more information.
4. Click **Apply** to apply your changes.

The **Lifecycle state** changes to **Updating**. When the request completes, the **Lifecycle state** shows **Available**.

Configure Autonomous Database Built-in Tools (OCPU compute model)

Describes how to enable or disable Autonomous Database built-in tools for an existing Autonomous Database instance.

Tool status is retained across starts and restarts for an Autonomous Database instance. For example, if a tool is disabled before you stop or restart an instance, the tool maintains the same status (disabled) after a start or restart.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the  next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

1. On the Autonomous Database details page, select the **Tool configuration** tab.
2. Click **Edit tool configuration**.
3. Select or deselect in the **Enable tool** field to enable or disable a tool.
4. Click **Apply** to apply your changes.

Configure Tools

Tools	Description	Enable tool
Oracle APEX	Oracle APEX is a low-code development platform that you can use to build scalable, secure enterprise applications that can be deployed anywhere. Learn more.	Enabled
Database Actions	Oracle Database actions is a web-based interface that uses Oracle REST Data Services to provide development, data tools, administration and monitoring features for Oracle Autonomous Database. Learn more.	Enabled
Graph Studio	Graph Studio automates the creation of knowledge (RDF) and property graphs and includes interactive tooling for query, analysis, and visualization of these graphs in the Autonomous Database. You must log in as a graph-enabled user to access Graph Studio. Create this user in Database Actions. Learn more.	Enabled
Oracle Machine Learning User Interface	The Oracle Machine Learning User Interface provides immediate access to the Oracle Machine Learning components and functionality on Autonomous Database, including OML Notebooks, OML AutoML UI, OML Models, and template example notebooks. Learn more.	Enabled
Data Transforms	Oracle Data Transforms allows you to design graphical data transformations in the form of data flows and workflows. The data flows define how the data is moved and transformed between different systems, while the workflows define the sequence in which the data flows are executed. Learn more.	Enabled
Web Access (ORDS)	Oracle REST Data Services (ORDS) provides HTTPS interfaces for working with the contents of your Oracle Database in one or more REST enabled schemas. Learn more.	Enabled
MongoDB API	Oracle Database API for MongoDB enables MongoDB-compatible clients and drivers to connect directly to Autonomous Database. Learn more.	Disabled Use Access Control List or Private Endpoint as your network access type to be able to enable MongoDB API

Apply Cancel

Copyright © 2022, Oracle and/or its affiliates. All rights reserved.

The **Lifecycle state** changes to **Updating**. When the request completes, the **Lifecycle state** shows **Available**.

Configure Autonomous Database Built-in Tools when you Provision or Clone an Instance

During provisioning or when you clone an instance, you can enable or disable Autonomous Database built-in tools.

- Configure Autonomous Database Built-in Tools when you Provision or Clone an Instance (ECPU compute model)
- Configure Autonomous Database Built-in Tools when you Provision or Clone an Instance (OCPU compute model)

Configure Autonomous Database Built-in Tools when you Provision or Clone an Instance (ECPU compute model)

Describes how to enable or disable Autonomous Database built-in tools when you provision or clone an instance.

When you are cloning, note the following:

- When you are cloning and you select **Full Clone** or **Metadata Clone**, the built-in tools default settings are based on the defaults for the workload type of the clone. In this case, the defaults are not based on source database built-in tools settings.
- When you are cloning and you select **Refreshable Clone**, the refreshable clone inherits the built-in tool status from the source database. Editing the built-in tool status for refreshable clone is not be possible during or after provisioning. If the source database of refreshable clone changes its tool configuration, this is reflected in the refreshable clone after the next refresh.

Perform the steps to provision or clone an instance:

- Provision an Autonomous Database Instance
- Clone an Autonomous Database Instance

1. Click **Show advanced options** to select advanced options.
2. In advanced options, select the **Tools** tab.
3. Click **Edit tool configuration**.
4. Select or deselect a tool in the **Status** field to enable or disable the tool.
5. (Optional) When enabled, Graph Studio, Oracle Machine Learning User Interface, and Data Transforms provide default **ECPU count** and **Maximum idle count** values and let you configure these values.

   Change the **ECPU count** value if you want to provide more or less resources for a built-in tool.

   Change the **Maximum idle time value** to specify the time, in minutes, after which an idle Graph Studio, Data Transforms, or Oracle Machine Learning VM is terminated.

   See About Configuring Built-in Tools Compute Resources and Limits for more information.
6. (Optional) When enabled, Graph Studio, Oracle Machine Learning User Interface, and Data Transforms provide default **ECPU count** and **Maximum idle count** values and let you configure these values.

Change the **ECPU count** value if you want to provide more or less resources for a built-in tool.

Change the **Maximum idle time value** to specify the time, in minutes, after which an idle Graph Studio, Data Transforms, or Oracle Machine Learning VM is terminated.

See About Configuring Built-in Tools Compute Resources and Limits for more information.

7. Click **Apply** to apply the configuration changes and return to the provision or clone steps.

Configure Autonomous Database Built-in Tools when you Provision or Clone an Instance (OCPU compute model)

Describes how to enable or disable Autonomous Database built-in tools when you provision or clone an instance

When you are cloning, note the following:

- When you are cloning and you select **Full Clone** or **Metadata Clone**, the built-in tools default settings are based on the defaults for the workload type of the clone. In this case, the defaults are not based on source database built-in tools settings.
- When you are cloning and you select **Refreshable Clone**, the refreshable clone inherits the built-in tool status from the source database. Editing the built-in tool status for refreshable clone is not be possible during or after provisioning. If the source database of refreshable clone changes its tool configuration, this is reflected in the refreshable clone after the next refresh.

Perform the steps to provision or clone an instance:

- Provision an Autonomous Database Instance
- Clone an Autonomous Database Instance

1. Click **Show advanced options** to select advanced options.
2. In advanced options, select the **Tools** tab.
3. Click **Edit tool configuration**.
4. Select or deselect a tool in the **Status** field to enable or disable the tool.
5. Click **Apply** to apply the configuration changes and return to the provision or clone steps.

Notes for Autonomous Database Built-in Tools

Provides additional Autonomous Database built-in tools notes and lists differences that apply depending on the workload type of your Autonomous Database instance.

- Disabling ORDS affects built-in tools such as Database Actions, APEX, and MongoDB API. Therefore, if you disable ORDS, this also disables APEX, Database Actions, and MongoDB API.
- When ORDS is disabled, Database Actions, APEX, and MongoDB API are disabled. If you subsequently enable ORDS this does not automatically enable APEX, Database Actions, or the MongoDb APIs. You must manually enable each built-in tool that is dependent on ORDS. Enable ORDS first, then you can enable APEX, Database Actions, or the MongoDB APIs
- Always Free Autonomous Database does not provide configuration options for Autonomous Database tools and does not allow you to enable or disable Autonomous Database tools.

- By default when you provision an Autonomous Database with Autonomous Data Warehouse or Autonomous Transaction Processing workload types, the MongoDB API is disabled. If MongoDB API is disabled, you can enable it during provisioning or cloning, or in an existing database.

  You can enable MongoDB API only if you have an ACL or if you are on a private endpoint. See Configure Access for MongoDB for more information.
- For JSON Database, Graph Studio is not supported and enabling or disabling Graph Studio for JSON Database is not possible.

  For JSON Database, MongoDB API is enabled by default and you can disable it during the provisioning or cloning, or in an existing database.
- When you are cloning and you select **Full Clone** or **Metadata Clone**, the built-in tools default settings are based on the defaults for the workload type of the clone. In this case, the defaults are not based on source database built-in tools settings.
- When you are cloning and you select **Refreshable Clone**, the refreshable clone inherits the built-in tool status from the source database. Editing the built-in tool status for refreshable clone is not be possible during or after provisioning. If the source database of refreshable clone changes its tool configuration, this is reflected in the refreshable clone after the next refresh.
- When you provision or clone an Autonomous Database with the APEX workload type (APEX service), the only supported built-in tools are Oracle APEX, Web Access (ORDS), and Database Actions. In the APEX service, you can only disable Database Actions.
- All existing sessions end after you disable a built-in tool. In addition, disabling Oracle APEX or Web Access (ORDS) also affects the existing applications that are based on these tools. For example, a user application based on ORDS does not continue to work after you disable ORDS.

Use Autonomous Database Events

Oracle Cloud Infrastructure Events enable you to create automation based on state changes for your Autonomous Database.

You can use Oracle Cloud Infrastructure Events to subscribe to and be notified of Autonomous Database events. The Oracle Cloud Infrastructure Events allow you to create automation and to receive notifications based on state changes for resources.

- About Events Based Notification and Automation on Autonomous Database
- Critical Events on Autonomous Database
  Critical events on Autonomous Database are issues that cause disruption to the database.
- Information Events on Autonomous Database
  Information events provide important details about database lifecycle events, workload capture and reply events, and some other non-critical events on Autonomous Database.
- Autonomous Database Lifecycle Events
  Provides a list of Autonomous Database lifecycle events.
- Get Notified of Autonomous Database Events
  Using Oracle Cloud Infrastructure Events you can subscribe to Autonomous Database events individually or in two categories, Critical and Information events.

About Events Based Notification and Automation on Autonomous Database

You can subscribe to Oracle Cloud Infrastructure Events.

An event could be a resource life cycle state change or a system event impacting a resource. For example, an event is emitted on Autonomous Database when maintenance begins and ends. Subscribing to events allows you to receive notifications or to perform other types of automation for events.

You can subscribe to Autonomous Database events as follows:

- Critical Events: See Critical Events on Autonomous Database for details on the events in the Critical event category.
- Information Events: See Information Events on Autonomous Database for details on the events in the Information event category.
- Lifecycle Management Events. See Autonomous Database Lifecycle Events for the list of Autonomous Database events.
- Autonomous Data Guard Events. See Events and Notifications for a Standby Database for the list of Autonomous Data Guard events.

Critical events and Information events include an event category attribute (in the `eventCategory` field). Optionally you can subscribe to events and add an event category to receive notifications about events in a category, or to perform other types of automation for events by event type and event category.

The event categories are:

- **Incident**: Database incident events
- **Maintenance**: Database maintenance events
- **Performance**: Database performance events
- **Security**: Database security events
- **Other**: Other events

See Overview of Events for complete information on Oracle Cloud Infrastructure Events.

See Get Notified of Autonomous Database Events for information on creating event rules.

Critical Events on Autonomous Database

Critical events on Autonomous Database are issues that cause disruption to the database.

Event	Event Category	Description
`AdminPasswordWarning`	Security	Provides a message that the Autonomous Database ADMIN password is expiring within 30 days or is expired. • If the ADMIN password is within 30 days of being unusable, you receive a warning indicating the date when the ADMIN password can no longer be used. • If the ADMIN password expires and is no longer usable, Autonomous Database reports an event that specifies that the ADMIN user password has expired and must be reset.
`AutomaticFailoverBegin`	Incident	When an automatic failover begins for the Autonomous Database and the database is unavailable an `AutonomousDatabase-AutomaticFailoverBegin` event is generated. The event will only be triggered if you are using Autonomous Data Guard.

Event	Event Category	Description
`AutomaticFailoverEnd`	Incident	When an automatic failover that was triggered is complete and the database is now available an `AutonomousDatabase-AutomaticFailoverEnd` event is generated. The event will only be triggered if you are using Autonomous Data Guard.
`DatabaseDownBegin`	Incident	The Autonomous Database instance cannot be opened, or the services such as high, low, medium, tp, or tpurgent are not started or available. The following conditions do not trigger `DatabaseDownBegin`: • Operations performed during the maintenance window • Load balancer, network, or backup related issues • A user stopping the instance This event will not be triggered if you are using Autonomous Data Guard and the standby database is not available due to any of these conditions.
`DatabaseDownEnd`	Incident	The database is recovered from the down state, meaning the Autonomous Database instance is opened with its services, following a `DatabaseDownBegin` event. `DatabaseDownEnd` is triggered only if there was a preceding `DatabaseDownBegin` event. The following conditions do not trigger `DatabaseDownEnd`: • Operations performed during the maintenance window • A user starting the instance If you are using Autonomous Data Guard and the primary database goes down, this triggers a `DatabaseDownBegin` event. If the system fails over to the standby database, this triggers a `DatabaseDownEnd` event.
`DatabaseInaccessibleBegin`	Incident	The Autonomous Database instance is using customer-managed keys and the database becomes inaccessible (state shows **Inaccessible**). The following conditions trigger `DatabaseInaccessibleBegin`: • Oracle Cloud Infrastructure Vault master encryption key is deleted. • Oracle Cloud Infrastructure Vault master encryption key is disabled. • Autonomous Database instance is not able to reach the Oracle Cloud Infrastructure Vault.
`DatabaseInaccessibleEnd`	Incident	If the Autonomous Database instance recovers from the **Inaccessible** state that generated a `DatabaseInaccessibleBegin` event, when the database state changes to **Available**, the database triggers a `DatabaseInaccessibleEnd` event.

Event	Event Category	Description
`DatabaseUserPasswordWarning`	Security	Provides a message that the Autonomous Database user password is expiring within the specified grace period (default 30 days) or has already expired. • If an Autonomous Database user password is within the specified grace period of being unusable, you receive a warning indicating the date when the password can no longer be used. • If an Autonomous Database user password expires and is no longer usable, Autonomous Database reports an event that specifies that the password has expired and must be reset. The `DatabaseUserPasswordWarning` event is generated once per day.
`FailedLoginWarning`	Security	If the number of failed login attempts reaches 3 * *number of total users* in the last three (3) hours, the database triggers a `FailedLoginWarning` event. Failed login attempts could be due to a user entering an invalid username or password, or due to a connection timeout. The event is triggered when the failed login attempts threshold is exceeded within the monitored time period (the last three hours).
`InstanceRelocateBegin`	Incident	If an Autonomous Database is relocated to another Exadata infrastructure due to server maintenance, hardware refresh, a hardware issue, or as part of a resource scale up, the `InstanceRelocateBegin` event is triggered at the start of the relocation process.
`InstanceRelocateEnd`	Incident	If an Autonomous Database is relocated to another Exadata infrastructure due to server maintenance, hardware refresh, a hardware issue, or as part of a resource scale up, the `InstanceRelocateEnd` event is triggered at the end of the relocation process.
`WalletExpirationWarning`	Security	This event is generated when Autonomous Database determines that a wallet is due to expire in less than six (6) weeks. This event is reported at most once per week. This event is triggered when there is a connection that uses the wallet that is due to expire.

Use the event type **Autonomous Database - Critical** to specify critical events in Event rules.

Information Events on Autonomous Database

Information events provide important details about database lifecycle events, workload capture and reply events, and some other non-critical events on Autonomous Database.

The times shown with information events are in UTC.

Event	Event Category	Description
AJDNonJsonStorageExceeded	Performance	This event is generated when an Autonomous JSON Database has exceeded the maximum storage limit of 20GB of data stored outside of SODA collections. This limit does not apply to data stored in SODA collections or objects associated with SODA collections, such as indexes or materialized views. In addition to this event, an email is sent to the account owner. You must either reduce your usage of non-SODA-related data to below the 20GB limit or promote the Autonomous JSON Database to Autonomous Transaction Processing. See Promote to Autonomous Transaction Processing for more information.
APEXUpgradeAvailable	Incident	This event is generated when you are using Oracle APEX and a new release becomes available.
APEXUpgradeBegin	Incident	This event is generated when you are using Oracle APEX and your Autonomous Database instance begins an upgrade to a new Oracle APEX release or begins to apply an Oracle APEX Patch Set Bundle.
APEXUpgradeEnd	Incident	This event is generated when you are using Oracle APEX and your Autonomous Database instance completes an upgrade to a new Oracle APEX release or completes the installation of an Oracle APEX Patch Set Bundle.
AutomaticRefreshEnd	Incident	This event is generated when the scheduled Automatic Refresh completes. See Edit Automatic Refresh Policy for Refreshable Clone for more information.
AutomaticRefreshFailed	Incident	This event is generated when the scheduled Automatic Refresh fails. See Edit Automatic Refresh Policy for Refreshable Clone for more information.
DatabaseConnection	Security	This event is generated if a connection is made to database from a new IP address (the connection has not been made from the specified IP address in the last 30 days).

Event	Event Category	Description
`InactiveConnectionsDetected`	Security	This event is generated when the number of inactive database connections detected is above a certain ratio compared to all database connections for the Autonomous Database instance. Subscribing to this event can help you keep track of unused connections. This event is generated once per day, when the following is true: • The number of inactive connections, where the elapsed time for the inactive connection is greater than 24 hours, is great than 10% of the total number of connections in any state. For this calculation, inactive connections are connections where the `status` is `"INACTIVE"`. Use the following query to report detailed information on sessions that have been inactive for more than 24 hours: `SELECT sid, serial#, last_call_et FROM v$session WHERE status = 'INACTIVE' AND last_call_et > 60 * 60 *24;`
`LongPendingTransaction`	Incident	This event is generated when a transaction is open for longer than 24 hours. Only one event is generated per transaction. The generated event includes the transaction id, address, and session address.
`MaintenanceBegin`	Maintenance	This event is triggered when the maintenance starts and provides the start timestamp for the maintenance (this event does not provide the scheduled start time). This event is also triggered when the maintenance starts for a local Autonomous Data Guard standby.
`MaintenanceEnd`	Maintenance	This event is triggered when the maintenance ends and provides the end timestamp for the maintenance (this event does not provide the scheduled end time). This event is also triggered when the maintenance ends for a local Autonomous Data Guard standby.
`NewMaintenanceSchedule`	Maintenance	This event is generated when the maintenance date is updated and the new date is shown on the Oracle Cloud Infrastructure Console. This event is also triggered when the maintenance date is updated and the new date is shown for a local Autonomous Data Guard standby. When the Java version will be updated, the event `description` field indicates this with the following: `This maintenance involves Java version update, please expect down time of Java service during this maintenance window.`

Event	Event Category	Description
`NetworkUsageInbound`	Security	This event is triggered when the average incoming network traffic to your Autonomous Database for TCP or HTTP database connections in the last 24 hours has exceeded 50% of the average incoming traffic observed a day before. When the trigger conditions for this event are met, the `NetworkUsageInbound` event is generated once per day.
`NetworkUsageOutbound`	Security	This event is triggered when the average outgoing network traffic from your Autonomous Database for TCP or HTTP database connections in the last 24 hours has exceeded 50% of the average outgoing traffic observed a day before. When the trigger conditions for this event are met, the `NetworkUsageOutbound` event is generated once per day.
`OperatorAccess`	Other	This event is triggered when operator access is detected for the database. You can query the access details from the `DBA_OPERATOR_ACCESS` view with the request ID provided in the event description. The `OperatorAccess` event is generated once every 24 hours.
`QueryPerformanceDegradationDetected`	Performance	This event is triggered when the average elapsed time of a query has exceeded ten times its average elapsed time over the past seven (7) days. The `QueryPerformanceDegradationDetected` event `description` field indicates the potential cause of a query performance degradation. For example, with a message such as: `Query performance slow down detected for SQL ID 9x25vyx9ujnn8. This could potentially be due to a recent compute scale down performed on this Autonomous Database. You might need to scale up your instance's compute allocation to address this performance degradation.` When the trigger conditions for this event are met, the `QueryPerformanceDegradationDetected` event is generated at most once per day.

Event	Event Category	Description
`ScheduledMaintenanceWarning`	Maintenance	This event is generated when the instance is 24 hours from a scheduled maintenance and again when the instance is 1 hour (60 minutes) from the scheduled maintenance. This event is also triggered when a local Autonomous Data Guard standby is 24 hours from a scheduled maintenance and again when the local Autonomous Data Guard standby is 1 hour (60 minutes) from the scheduled maintenance. When the Java version will be updated, the event `description` field indicates this with the following message: `This maintenance involves Java version update, please expect down time of Java service during this maintenance window.`
`SQLParseErrorWarning`	Incident	If a SQL parse error occurs several times in a certain time window, for example 200 times or more in an hour, this generates "`Too Many Parse Errors`" entries in the alert log. The `SQLParseErrorWarning` event is triggered when there is a new entry in the alert log, indicating that a SQL statement has too many parse errors. For example, when an invalid query runs 200 times in less than an hour, this triggers the `SQLParseErrorWarning` event. This event is generated at most once per day. This means there can be a delay of up to 24 hours before the event is generated.
`WorkloadCaptureBegin`	Incident	This event is triggered when a workload capture is initiated.
`WorkloadCaptureEnd`	Incident	This event is triggered when a workload capture completes successfully. This generates a pre-authenticated (PAR) URL to download the capture reports. This URL is contained in the `captureDownloadURL` field of the event and is valid for 7 days from the date of generation.
`WorkloadReplayBegin`	Incident	This event is triggered when a workload replay is initiated.
`WorkloadReplayEnd`	Incident	This event is triggered when a workload replay completes successfully. This generates a pre-authenticated (PAR) URL to download the replay reports. This URL is contained in the `replayDownloadURL` field of the event and is valid for 7 days from the date of generation.

Note:

When Autonomous Data Guard is enabled, any of these events that occur on the standby database do not trigger an Information event.

Use the event type **Autonomous Database - Information** to specify information events in Event rules.

See View Patch and Maintenance Window Information, Set the Patch Level for information on maintenance windows.

Autonomous Database Lifecycle Events

Provides a list of Autonomous Database lifecycle events.

Friendly Name	Event Type
Autonomous Database - Access Control Lists Update Begin	`com.oraclecloud.databaseservice.updateautonomousdatabaseacl.begin`
Autonomous Database - Access Control Lists Update End	`com.oraclecloud.databaseservice.updateautonomousdatabaseacl.end`
Autonomous Database - Auto Scaling Disabled	`com.oraclecloud.databaseservice.autonomousdatabaseautoscaledisabled`
Autonomous Database - Auto Scaling Enabled	`com.oraclecloud.databaseservice.autonomousdatabaseautoscaleenabled`
Autonomous Database - Automatic Backup Begin	`com.oraclecloud.databaseservice.automaticbackupautonomousdatabase.begin`
Autonomous Database - Automatic Backup End	`com.oraclecloud.databaseservice.automaticbackupautonomousdatabase.end`
Autonomous Database - Change Compartment Begin	`com.oraclecloud.databaseservice.changeautonomousdatabasecompartment.begin`
Autonomous Database - Change Compartment End	`com.oraclecloud.databaseservice.changeautonomousdatabasecompartment.end`
Autonomous Database - Change Database Name Begin	`com.oraclecloud.databaseservice.changeautonomousdatabasename.begin`
Autonomous Database - Change Database Name End	`com.oraclecloud.databaseservice.changeautonomousdatabasename.end`
Autonomous Database - Create Backup Begin	`com.oraclecloud.databaseservice.autonomous.database.backup.begin`
Autonomous Database - Create Backup End	`com.oraclecloud.databaseservice.autonomous.database.backup.end`
Autonomous Database - Create Begin	`com.oraclecloud.databaseservice.autonomous.database.instance.create.begin`
Autonomous Database - Create End	`com.oraclecloud.databaseservice.autonomous.database.instance.create.end`
Autonomous Database - Critical See Critical Events on Autonomous Database for more information.	`com.oraclecloud.databaseservice.autonomous.database.critical`
Autonomous Database - Deregister Autonomous Database with Data Safe Begin	`com.oraclecloud.databaseservice.deregisterautonomousdatabasedatasafe.begin`
Autonomous Database - Deregister Autonomous Database with Data Safe End	`com.oraclecloud.databaseservice.deregisterautonomousdatabasedatasafe.end`
Autonomous Database - Disable Data Guard Begin	`com.oraclecloud.databaseservice.disableautonomousdataguard.begin`
Autonomous Database - Disable Data Guard End	`com.oraclecloud.databaseservice.disableautonomousdataguard.end`

Friendly Name	Event Type
Autonomous Database - Disconnect Refreshable Clone from Source Database Begin	`com.oraclecloud.databaseservice.disconnectrefreshableautonomousdatabaseclonefromsource.begin`
Autonomous Database - Disconnect Refreshable Clone from Source Database End	`com.oraclecloud.databaseservice.disconnectrefreshableautonomousdatabaseclonefromsource.end`
Autonomous Database - Enable Data Guard Begin	`com.oraclecloud.databaseservice.enableautonomousdataguard.begin`
Autonomous Database - Enable Data Guard End	`com.oraclecloud.databaseservice.enableautonomousdataguard.end`
Autonomous Database - Free Database Automatic Stop Reminder This event is emitted 48 hours prior to database stop.	`com.oraclecloud.databaseservice.freeautonomousdatabasestopreminder`
Autonomous Database - Free Database Automatically Stopped	`com.oraclecloud.databaseservice.freeautonomousdatabasestopped`
Autonomous Database - Free Database Automatic Termination Reminder. This event is emitted 48 hours prior to database termination.	`com.oraclecloud.databaseservice.freeautonomousdatabaseterminationreminder`
Autonomous Database - Free Database Automatically Terminated	`com.oraclecloud.databaseservice.freeautonomousdatabaseterminated`
Autonomous Database - Generate Wallet	`com.oraclecloud.databaseservice.generateautonomousdatabasewallet`
Autonomous Database - Information See Information Events on Autonomous Database for more information.	`com.oraclecloud.databaseservice.autonomous.database.information`
Autonomous Database - Long term backup ended This event is triggered when a long term backup completes.	
Autonomous Database - Long term backup schedule disabled This event is triggered when a long term backup schedule is disabled.	
Autonomous Database - Long term backup schedule enabled / updated This event is triggered when a long term backup schedule is either enabled or updated.	
Autonomous Database - Long term backup started This event is triggered when a long term backup starts.	
Autonomous Database - Manual Failover Begin	`com.oraclecloud.databaseservice.failoverautonomousdatabase.begin`
Autonomous Database - Manual Failover End This event is emitted after failover completes successfully or unsuccessfully. Unsuccessful manual failovers may result in data loss. Check the Autonomous Database details page in the Oracle Cloud Infrastructure Console for additional information in the event of an unsuccessful manual failover.)	`com.oraclecloud.databaseservice.failoverautonomousdatabase.end`

Friendly Name	Event Type
Autonomous Database - Manual Refresh Begin The "Manual Refresh" events are used for Autonomous Database refreshable clones.	com.oraclecloud.databaseservice.manualrefresh.begin
Autonomous Database - Manual Refresh End	com.oraclecloud.databaseservice.manualrefresh.end
Autonomous Database - Deregister Autonomous Database with Data Safe Begin	com.oraclecloud.databaseservice.deregisterautonomousdatabasedatasafe.begin
Autonomous Database - Deregister Autonomous Database with Data Safe End	com.oraclecloud.databaseservice.deregisterautonomousdatabasedatasafe.end
Autonomous Database - Register Autonomous Database with Data Safe Begin	com.oraclecloud.databaseservice.registerautonomousdatabasedatasafe.begin
Autonomous Database - Register Autonomous Database with Data Safe End	com.oraclecloud.databaseservice.registerautonomousdatabasedatasafe.end
Autonomous Database - Restart Begin	com.oraclecloud.databaseservice.restartautonomousdatabase.begin
Autonomous Database - Restart End	com.oraclecloud.databaseservice.restartautonomousdatabase.end
Autonomous Database - Restore Begin	com.oraclecloud.databaseservice.autonomous.database.restore.begin
Autonomous Database - Restore End	com.oraclecloud.databaseservice.autonomous.database.restore.end
Autonomous Database - Rotate Encryption Key Begin	com.oraclecloud.databaseservice.rotateautonomousdatabaseencryptionkey.begin
Autonomous Database - Rotate Encryption Key End	com.oraclecloud.databaseservice.rotateautonomousdatabaseencryptionkey.end
Autonomous Database - Start Begin	com.oraclecloud.databaseservice.startautonomousdatabase.begin
Autonomous Database - Start End	com.oraclecloud.databaseservice.startautonomousdatabase.end
Autonomous Database - Stop Begin	com.oraclecloud.databaseservice.stopautonomousdatabase.begin
Autonomous Database - Stop End	com.oraclecloud.databaseservice.stopautonomousdatabase.end
Autonomous Database - Switchover Begin	com.oraclecloud.databaseservice.switchoverautonomousdatabase.begin
Autonomous Database - Switchover End This event is emitted after switchover completes successfully or unsuccessfully. Unsuccessful switchovers may result in data loss. Check the Autonomous Database details page in the Oracle Cloud Infrastructure Console for additional information in the event of an unsuccessful switchover.	com.oraclecloud.databaseservice.switchoverautonomousdatabase.end
Autonomous Database - Terminate Begin	com.oraclecloud.databaseservice.deleteautonomousdatabase.begin
Autonomous Database - Terminate End	com.oraclecloud.databaseservice.deleteautonomousdatabase.end
Autonomous Database - Update Begin	com.oraclecloud.databaseservice.updateautonomousdatabase.begin

Friendly Name	Event Type
Autonomous Database - Update End	`com.oraclecloud.databaseservice.updateautonomousdatabase.end`
Autonomous Database - Update Open Mode Begin	`com.oraclecloud.databaseservice.updateautonomousdatabaseopenmode.begin`
Autonomous Database - Update Open Mode End	`com.oraclecloud.databaseservice.updateautonomousdatabaseopenmode.end`
Autonomous Database - Upgrade Database Version Begin	`com.oraclecloud.databaseservice.upgradeautonomousdatabasedbversion.begin`
Autonomous Database - Upgrade Database Version End	`com.oraclecloud.databaseservice.upgradeautonomousdatabasedbversion.end`

Get Notified of Autonomous Database Events

Using Oracle Cloud Infrastructure Events you can subscribe to Autonomous Database events individually or in two categories, Critical and Information events.

> **Note:**
>
> When you subscribe to an event category, either Critical or Information, you are notified when any events in the category occur. For example, to get notified when the database goes down, subscribe to the Autonomous Database Critical event type.

If you want to subscribe to Critical or Information events:

- To subscribe to Critical events, create an event rule using the **Autonomous Database - Critical** event type.
- To subscribe to Information events, create an event rule using the **Autonomous Database - Information** event type.

If you want to subscribe to one or more individual Critical or Information events, create an event rule and add conditions. For example, to create an event rule that only applies for the `AdminPasswordWarning` Critical event, create a rule as follows:

1. On the Oracle Cloud Infrastructure Console click **Observability & Management**.
2. Under **Events Service**, click **Rules**.
3. To add a rule, click **Create Rule**.
4. On the Create Rule page, enter a Display Name and a Description.
5. Under Rule Conditions, enter a rule condition
   a. Enter the Condition: **Event Type**.
   b. Enter the Service Name: **Database**.
   c. Enter the Event Type, one of: **Autonomous Database - Critical** or **Autonomous Database - Information**.
6. Click **+ Another Condition**.
7. Under Condition, select **Attribute**.
   a. Under Attribute Name, select **eventName**.

b. Under Attribute Values, enter the event name. For example, enter `AdminPasswordWarning`.

   See Critical Events on Autonomous Database and Information Events on Autonomous Database for the list of event names.

For example with these entries the Rule Logic area shows:

```
MATCH event WHERE (
   eventType EQUALS ANY OF (
      com.oraclecloud.databaseservice.autonomous.database.critical
   )
   AND (
         eventName MATCHES ANY OF (
             AdminPasswordWarning
         )
   )
)
```

8. In the **Actions** area, select the actions you want.
9. Click **Create Rule**.

Critical events and Information events include an event category attribute (in the `eventCategory` field). Optionally you can subscribe to events by event type and include an event category to receive notifications about events in that category, or to perform other types of automation for events.

The event categories are:

- **Incident**: Database incident events
- **Maintenance**: Database maintenance events
- **Performance**: Database performance events
- **Security**: Database security events
- **Other**: Other events

For example, if you want to subscribe Critical events with the event category `Security`, create an event rule as shown, and in Step 7, select the `eventCategory` attribute and enter a value.

The result would look like this:

```
MATCH event WHERE (
   eventType EQUALS ANY OF (
      com.oraclecloud.databaseservice.autonomous.database.critical
   )
   AND (
     eventCategory MATCHES ANY OF (
       Security
     )
   )
)
```

See Getting Started with Events for information on using Oracle Cloud Infrastructure Events, creating event rules, and for information on configuring actions for notifications.

Manage Time Zone File Updates on Autonomous Database

Autonomous Database provides several options to automatically update time zone files on an Autonomous Database database instance.

- About Time Zone File Update Options
  For time zone support Oracle Database uses time zone files that store the list of all time zones. The time zone files for Autonomous Database are periodically updated to reflect the latest time zone specific changes.
- Use AUTO_DST_UPGRADE Time Zone File Option
- Use AUTO_DST_UPGRADE_EXCL_DATA Time Zone File Option

About Time Zone File Update Options

For time zone support Oracle Database uses time zone files that store the list of all time zones. The time zone files for Autonomous Database are periodically updated to reflect the latest time zone specific changes.

Autonomous Database provides the following options for updating time zone files:

- `AUTO_DST_UPGRADE`: Automatically upgrades the time zone files and automatically updates the data on your database to use the latest time zone data. This option requires that you restart or stop and then start your Autonomous Database instance.
- `AUTO_DST_UPGRADE_EXCL_DATA`: Automatically upgrades the time zone files and does not automatically update the data on your database to use the latest time zone data. The time zone files are upgraded to the latest version when you enable this option and versions are kept up to date with the latest version. This option does not require that you restart your Autonomous Database instance.

> **Note:**
>
> By default, both `AUTO_DST_UPGRADE` and `AUTO_DST_UPGRADE_EXCL_DATA` are disabled. You can enable one or the other of these options, but not both.

With every Daylight Saving Time (DST) version release, there are DST file changes introduced to make existing data comply with the latest DST rules. Applying a change to the database time zone files not only affects the way new data is handled, but potentially alters data stored in `TIMESTAMP WITH TIME ZONE` columns.

When a load or import operation results in the following time zone related error, this indicates that the version of your time zone files is out of date and you need to update to the latest version available for your database:

```
ORA-39405: Oracle Data Pump does not support importing from a source database
with TSTZ version y
into a target database with TSTZ version n.
```

You can enable the `AUTO_DST_UPGRADE` feature that automatically upgrades an instance to use the latest available version of the time zone files, and automatically updates the rows on your database to use the latest time zone data. The latest version of time zone files are applied and values in `TIMESTAMP WITH TIME ZONE` columns are updated at the next database restart.

However, depending on your database usage, the required database restart might restrict you from using AUTO_DST_UPGRADE to upgrade to the latest time zone files.

You can enable AUTO_DST_UPGRADE_EXCL_DATA to update your time zone files. This option allows you to immediately update your database if you encounter import/export errors due to a version mismatch of time zone files (where you receive the ORA-3940 error).

Enabling AUTO_DST_UPGRADE_EXCL_DATA on your database can be beneficial in the following cases:

- Data on the instance is in UTC and the database is not impacted by DST time zone file updates.
- The instance does not have any data in TIMESTAMP WITH TIME ZONE columns.
- Data in the instance uses dates that are not altered with new time zone or daylight saving time policies.

In this case, when your Autonomous Database instance does not contain rows that are adversely impacted by the new time zone rules and you encounter the ORA-3940 error, you can enable the AUTO_DST_UPGRADE_EXCL_DATA option to update to the latest version of the time zone files. The AUTO_DST_UPGRADE_EXCL_DATA option prioritizes the success of Oracle Data Pump jobs over the data consistency issues due to changing DST.

See Oracle Support Document 406410.1 to help you determine whether time zone changes will affect your database.

> **Note:**
>
> Oracle recommends enabling the AUTO_DST_UPGRADE option when these limiting cases do not apply to your database.

In summary, the choice of enabling automatic time zone upgrades with AUTO_DST_UPGRADE or AUTO_DST_UPGRADE_EXCL_DATA involves the following considerations:

- Possible Data Inconsistency: Oracle recommends that you maintain your database on the latest time zone file version by enabling AUTO_DST_UPGRADE. This option automatically updates the rows on your database to use the latest time zone data.

  The alternative option, AUTO_DST_UPGRADE_EXCL_DATA allows you to maintain your database using the latest time zone version without requiring a database restart; however, data that might be affected by the updated time zone files is not updated, which could lead to possible data inconsistency.
- Database Restart: Oracle recommends that you maintain your database on the latest time zone file version by enabling AUTO_DST_UPGRADE. This option requires a restart to upgrade to the latest time zone files version and during the restart updates the rows for existing data of TIMESTAMP WITH TIME ZONE data type to use the latest version. Enabling AUTO_DST_UPGRADE can affect database restart time. A restart can require extra time compared to a restart without this option enabled when there are new time zone files and there is data that needs to be updated when the database restarts.

  When you enable AUTO_DST_UPGRADE_EXCL_DATA, at the time you enable this option the database is upgraded to the latest version the time zone files and during each maintenance window, if newer time zone files are available, Autonomous Database updates database to use the latest version. Enabling AUTO_DST_UPGRADE_EXCL_DATA does not require that you restart your Autonomous Database instance to keep up to date with the latest version of the time zone files.

See Datetime Data Types and Time Zone Support for more information.

Use AUTO_DST_UPGRADE Time Zone File Option

Autonomous Database provides the `AUTO_DST_UPGRADE` option to automatically update time zone files on an Autonomous Database database instance.

The `AUTO_DST_UPGRADE` feature automatically upgrades the time zone files and automatically upgrades the rows on your database to use the latest time zone data.

> **Note:**
>
> By default, both `AUTO_DST_UPGRADE` and `AUTO_DST_UPGRADE_EXCL_DATA` are disabled. You can enable one or the other of the automatic time zone file update options, but not both.

When you enable `AUTO_DST_UPGRADE` your Autonomous Database instance automatically applies updates for time zone files, depending on the state of the instance:

- **Stopped**: At the next start operation the update is automatically applied.
- **Available**: After a restart, or a stop and then start, the update is automatically applied.

When a load or import operation results in the following time zone related error, this indicates that your time zone files are out of date and you need to update to the latest version available for your database:

```
ORA-39405: Oracle Data Pump does not support importing from a source database
with TSTZ version n+1
into a target database with TSTZ version n.
```

To enable automatic time zone file updates with `AUTO_DST_UPGRADE`:

1. Enable the `AUTO_DST_UPGRADE` feature.

```
BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_FEATURE(
        feature_name => 'AUTO_DST_UPGRADE');
END;
/
```

2. Query `dba_cloud_config` to verify that `AUTO_DST_UPGRADE` is enabled.

```
SELECT param_name, param_value FROM dba_cloud_config WHERE
       LOWER(param_name) = 'auto_dst_upgrade';

PARAM_NAME              PARAM_VALUE
----------------------- ----------------------------------------
auto_dst_upgrade        enable
```

3. Query `dba_cloud_config` to check the time zone version.

```
SELECT param_name, param_value FROM dba_cloud_config
    WHERE LOWER(param_name) = 'latest_timezone_version';
```

```
PARAM_NAME              PARAM_VALUE
----------------------- -----------
latest_timezone_version 38
```

You can query the `DB_NOTIFICATIONS` view to see if the time zone version is the latest version:

```
SELECT type, time, description, expected_start_date FROM db_notifications
   WHERE TYPE='TIMEZONE VERSION';
```

When `AUTO DST UPGRADE` is enabled, Autonomous Database upgrades to the latest version of the time zone files (when you next restart, or stop and then start the database). The columns in your database with the datatype `TIMESTAMP WITH TIME ZONE` are converted to the new time zone version during the restart.

To disable `AUTO_DST_UPGRADE`:

1. Disable the `AUTO_DST_UPGRADE` feature:

```
BEGIN
   DBMS_CLOUD_ADMIN.DISABLE_FEATURE(
        feature_name => 'AUTO_DST_UPGRADE');
END;
/
```

2. Query `dba_cloud_config` to verify that `AUTO_DST_UPGRADE` is disabled.

```
SELECT param_name, param_value FROM dba_cloud_config WHERE
       LOWER(param_name) = 'auto_dst_upgrade';

0 rows selected.
```

See ENABLE_FEATURE Procedure and DISABLE_FEATURE Procedure for more information.

See TIMESTAMP WITH TIME ZONE Data Type for more information.

Use AUTO_DST_UPGRADE_EXCL_DATA Time Zone File Option

Autonomous Database provides the `AUTO_DST_UPGRADE_EXCL_DATA` option to automatically update time zone files on an Autonomous Database database instance.

The `AUTO_DST_UPGRADE_EXCL_DATA` automatically upgrades the time zone files and does not automatically upgrade the rows on your database to use the latest time zone data. When this option is enabled the database upgrades to the latest version of the time zone files and subsequently upgrades the database to use new versions of the time zone files during the Autonomous Database maintenance window (whenever a new version is available). This option does not require that you restart your Autonomous Database instance.

> **Note:**
>
> By default, both `AUTO_DST_UPGRADE` and `AUTO_DST_UPGRADE_EXCL_DATA` are disabled. You can enable one or the other of the automatic time zone file update options, but not both.

When a load or import operation results in the following time zone related error, this indicates that your time zone files are out of date and you need to update to the latest version available for your database:

```
ORA-39405: Oracle Data Pump does not support importing from a source database
with TSTZ version n+1
into a target database with TSTZ version n.
```

To enable automatic time zone file updates with `AUTO_DST_UPGRADE_EXCL_DATA`:

1. Enable the `AUTO_DST_UPGRADE_EXCL_DATA` feature.

   ```
   BEGIN
      DBMS_CLOUD_ADMIN.ENABLE_FEATURE(
           feature_name => 'AUTO_DST_UPGRADE_EXCL_DATA');
   END;
   /
   ```

2. Query `dba_cloud_config` to verify that `AUTO_DST_UPGRADE_EXCL_DATA` is enabled.

   ```
   SELECT param_name, param_value FROM dba_cloud_config WHERE
          LOWER(param_name) = 'auto_dst_upgrade_excl_data';

   PARAM_NAME                     PARAM_VALUE
   ------------------------------ ----------------------------------------
   auto_dst_upgrade_excl_data     enabled
   ```

3. Query `dba_cloud_config` to check the time zone version.

   ```
   SELECT param_name, param_value FROM dba_cloud_config
       WHERE LOWER(param_name) = 'latest_timezone_version';

   PARAM_NAME              PARAM_VALUE
   ----------------------- -----------
   latest_timezone_version 38
   ```

   You can query the `DB_NOTIFICATIONS` view to see if the time zone version is the latest version:

   ```
   SELECT type, time, description, expected_start_date FROM db_notifications
      WHERE TYPE='TIMEZONE VERSION';
   ```

When `AUTO_DST_UPGRADE_EXCL_DATA` is enabled, Autonomous Database upgrades to the latest version of the time zone files and checks and updates to new time zone file versions during each scheduled maintenance window. With `AUTO_DST_UPGRADE_EXCL_DATA` enabled, the

columns in your database with the datatype `TIMESTAMP WITH TIME ZONE` are not converted to the new time zone version.

To disable `AUTO_DST_UPGRADE_EXCL_DATA`:

1. Disable the `AUTO_DST_UPGRADE_EXCL_DATA` feature:

```
BEGIN
   DBMS_CLOUD_ADMIN.DISABLE_FEATURE(
        feature_name => 'AUTO_DST_UPGRADE_EXCL_DATA');
END;
/
```

2. Query `dba_cloud_config` to verify that `AUTO_DST_UPGRADE_EXCL_DATA` is disabled.

```
SELECT param_name, param_value FROM dba_cloud_config WHERE
       LOWER(param_name) = 'auto_dst_upgrade_excl_data';

0 rows selected.
```

See ENABLE_FEATURE Procedure and DISABLE_FEATURE Procedure for more information.

See TIMESTAMP WITH TIME ZONE Data Type for more information.

Monitor Autonomous Database Availability

Shows you the steps to view availability information for an Autonomous Database instance.

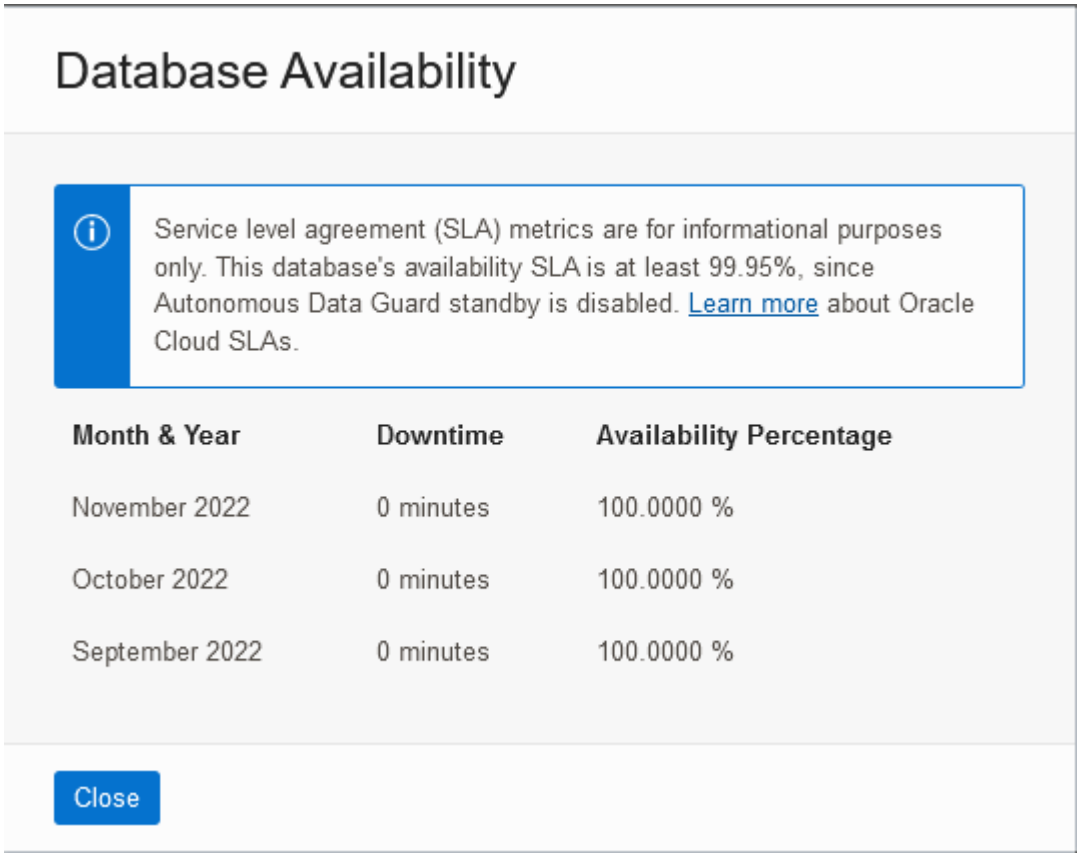
Database Availability is calculated based on the "Monthly Uptime Percentage" described in the Oracle PaaS and IaaS Public Cloud Services Pillar Document, document under **Delivery Policies** (see Autonomous Database Availability Service Level Agreement).

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To view Database Availability for an Autonomous Database instance:

1. On the **Details** page under General Information, click the Database Availability link in the **Lifecycle state** field.



2. Click **Close** to dismiss the page.

Note the following for the Autonomous Database instance availability:

- For databases with Cross-Region Autonomous Data Guard enabled, review both the primary database and the remote standby database availability.
- The Database Availability page shows information for the current month and for the previous two months.
- After you provision a new Autonomous Database instance, database availability information is only available for the current month.
- See View Metrics for an Autonomous Database Instance for more information on database availability. The Database Availability metric provides the Database Availability history for an Autonomous Database instance.
- Database availability information is not available prior to the introduction of this feature in April 2022.

See Monitor Regional Availability of Autonomous Databases for information on regional average availability metrics for Autonomous Database.

Monitor Regional Availability of Autonomous Databases

The regional availability is represented as a list of daily uptime percentages of Autonomous Databases across a set of services in each available region.

The regional availability data is updated daily on the following page:

Regional Average Availability Metrics for Oracle Autonomous Database

Notes for regional availability:

- The regional availability page shows information for the current month and for the previous two months.
- Regional availability information starts in June 2022 (availability data is not shown for dates prior to June 2022).

See View Metrics for an Autonomous Database Instance for more information regional availability.

See Monitor Autonomous Database Availability for information on Autonomous Database instance availability.

Manage Resource Availability with Compartment Quotas

Autonomous Database supports the compartment quotas feature of Oracle Cloud Infrastructure.

Compartment quotas give tenant and compartment administrators better control over how resources are consumed in Oracle Cloud Infrastructure, enabling administrators to easily allocate resources to compartments using the Console. Along with compartment budgets, compartment quotas create a powerful toolset to manage your spending in Oracle Cloud Infrastructure tenancies. You can start using compartment quotas from any compartment detail page in the Console.

Compartment quotas are similar to service limits. The biggest difference is that service limits are set by Oracle, and compartment quotas are set by administrators, using policies that allow them to allocate resources with a high level of flexibility.

Using compartment quotas, you can restrict, on a compartment by compartment basis, the availability of the following Autonomous Database resources in your tenancy:

- Number of Always Free Autonomous Databases. Tenancies can have a total of two Always Free Autonomous Databases, and these resources must be provisioned in the home region.
- Number of Autonomous Database for Developers databases.
- OCPUs available to Data Warehouse, Transaction Processing, JSON, or APEX databases.
- ECPUs available to Data Warehouse, Transaction Processing, JSON, or APEX databases.
- Total data storage available to Data Warehouse, Transaction Processing, JSON, or APEX databases.

> **Note:**
>
> Quotas are enforced only for provisioning databases and manually scaling their allocated ECPUs, OCPUs, or storage up or down.

You can manage the compartment quotas to limit the availability of a given resource in a given compartment using the `set`, `unset` and `zero` policy statements. These policy statements are written in a simple declarative language that is similar to the IAM policy language.

- **set:** Sets the maximum number of cloud resources that can be used for a compartment. For example, the following statement sets `atp-ecpu-count` quota to 10 ECPUs in the compartment `MyCompartment`:

```
set database quota atp-ecpu-count to 10 in compartment MyCompartment
```

- **unset:** Resets quotas back to the default service limits.
  For example, the following statement resets the quota for `atp-ecpu-count` in the entire tenancy:

  ```
  unset database quota atp-ecpu-count in tenancy
  ```

- **zero:** Removes access to a cloud resource for a compartment.
  For example, the following statement disallows use of ECPUs for APEX workload databases in the compartment `MyCompartment`:

  ```
  Zero apex-ecpu-count quotas in compartment MyCompartment
  ```

See Compartment Quotas for more information.

Compartment quotas are available on the OCI Console's Limits, Quotas and Usage page, which you can use to automatically generate quota statements. To access this page, on the **Governance & Administration** page, under **Tenancy Management** select **Limits, Quotas and Usage**. As a prerequisite, ensure you have the appropriate permissions to create quotas. See Policy Reference for the Quotas Service for more details.

Quotas for Oracle Autonomous Database

The following table lists the quotas for Oracle Autonomous Database

The family name associated with these quota names is: `database`

For example:

```
set database quota atp-ecpu-count to 10 in compartment MyCompartment
```

Name	Scope	Description
`adb-developer-count`	Regional	Number of Autonomous Database for Developers instances.
`adb-free-count`	Regional	Number of Always Free Autonomous Databases. Tenancies can have a total of two Always Free Autonomous Databases, and these resources must be provisioned in the home region. For each database, you can choose the workload type.
`adw-ecpu-count`	Regional	Number of Autonomous Databases for Analytics and Data Warehousing ECPUs for databases using Autonomous Database Serverless.
`adw-ocpu-count`	Regional	Number of Autonomous Databases for Analytics and Data Warehousing OCPUs for databases using Autonomous Database Serverless.
`adw-total-storage-tb`	Regional	Amount of storage (in TB) for Autonomous Database for Analytics and Data Warehousing databases using Autonomous Database Serverless.
`ajd-ecpu-count`	Regional	Number of JSON Database ECPUs with Autonomous Database Serverless.
`atp-ocpu-count`	Regional	Number of JSON Database OCPUs with Autonomous Database Serverless.
`ajd-total-storage-tb`	Regional	Total amount of storage (in TB) for JSON Database with Autonomous Database Serverless.

Name	Scope	Description
apex-ecpu-count	Regional	Number of APEX ECPUs with Autonomous Database Serverless.
apex-ocpu-count	Regional	Number of APEX OCPUs with Autonomous Database Serverless.
apex-total-storage-tb	Regional	Total amount of storage (in TB) for APEX with Autonomous Database Serverless.
atp-ecpu-count	Regional	Number of Autonomous Database for Autonomous Transaction Processing ECPUs for databases using Autonomous Database Serverless.
atp-ocpu-count	Regional	Number of Autonomous Database for Autonomous Transaction Processing OCPUs for databases using Autonomous Database Serverless.
atp-total-storage-tb	Regional	Amount of storage (in TB) for Autonomous Transaction Processing databases using Autonomous Database Serverless.
ex-cdb-count	Regional	External Container Database Count
ex-non-cdb-count	Regional	External Non Container Database Count
ex-pdb-count	Regional	External Pluggable Database Count

Managing and Viewing Maintenance Windows, Patching Options, Work Requests, and Customer Contacts

You can obtain information about system events, maintenance, and other important information for the operation of your Autonomous Database instance.

This includes the following:

- Oracle Cloud Infrastructure work requests allow you to monitor long-running operations such as cloning or backing up Autonomous Database.
- Oracle Cloud Infrastructure Console provides information on Autonomous Database maintenance windows and patching and provides a customer contacts area where you can view or manage the list of contacts to receive information on operational issues and service announcements.
- The DBA_OPERATOR_ACCESS view provides information on the top level SQL statements that Oracle Cloud Infrastructure cloud operations performs.

Topics

- About Work Requests
- View Patch and Maintenance Window Information, Set the Patch Level
  Autonomous Database uses predefined maintenance windows to automatically patch your database. You can view maintenance and patch information and see details for Autonomous Database maintenance history.
- View and Manage Customer Contacts for Operational Issues and Announcements
  You can view and manage the Autonomous Database customer contacts.

- View Oracle Cloud Infrastructure Operations Actions
  The `DBA_OPERATOR_ACCESS` view stores information on the actions that Oracle Cloud Infrastructure cloud operations performs on your Autonomous Database.

About Work Requests

You can view Oracle Cloud Infrastructure Work Requests on Autonomous Database.

A work request is an activity log that enables you to track each step in an operation's progress. Each work request has an Oracle Cloud Identifier (OCID) that allows you to interact with it programmatically.

Autonomous Database work requests are created for the following operations:

- Creating or terminating an Autonomous Database instance
- Starting or stopping an Autonomous Database instance
- Restoring an Autonomous Database instance
- Cloning an Autonomous Database instance
- Scaling database storage or CPU
- Updating the database license type
- Updating a database's network access control list (ACL)

Under **Resources** on the Oracle Cloud Infrastructure console click **Work Requests** to see recent work requests.

Click the link under **Operation** to see more details for a work request.

See Work Requests Integration for more information.

View Patch and Maintenance Window Information, Set the Patch Level

Autonomous Database uses predefined maintenance windows to automatically patch your database. You can view maintenance and patch information and see details for Autonomous Database maintenance history.

- About Scheduled Maintenance and Patching
  All Autonomous Database instances are automatically assigned to a maintenance window and different instances can have different maintenance windows.
- View Maintenance Event History
  You can view Autonomous Database maintenance event history for details about past maintenance events, such as the title, state, start time, and stop time.
- View Patch Details
  You can view Autonomous Database patch information, including a list of resolved issues and components.
- Set the Patch Level
  When you provision or clone an Autonomous Database instance you can select a patch level to apply to upcoming patches. You can also edit the patch level after an Autonomous Database instance is provisioned. There are two patch level options: **Regular** and **Early**.
- View Maintenance Status Notifications
  The `DB_NOTIFICATIONS` view stores information about maintenance status notifications for your Autonomous Database instance.

About Scheduled Maintenance and Patching

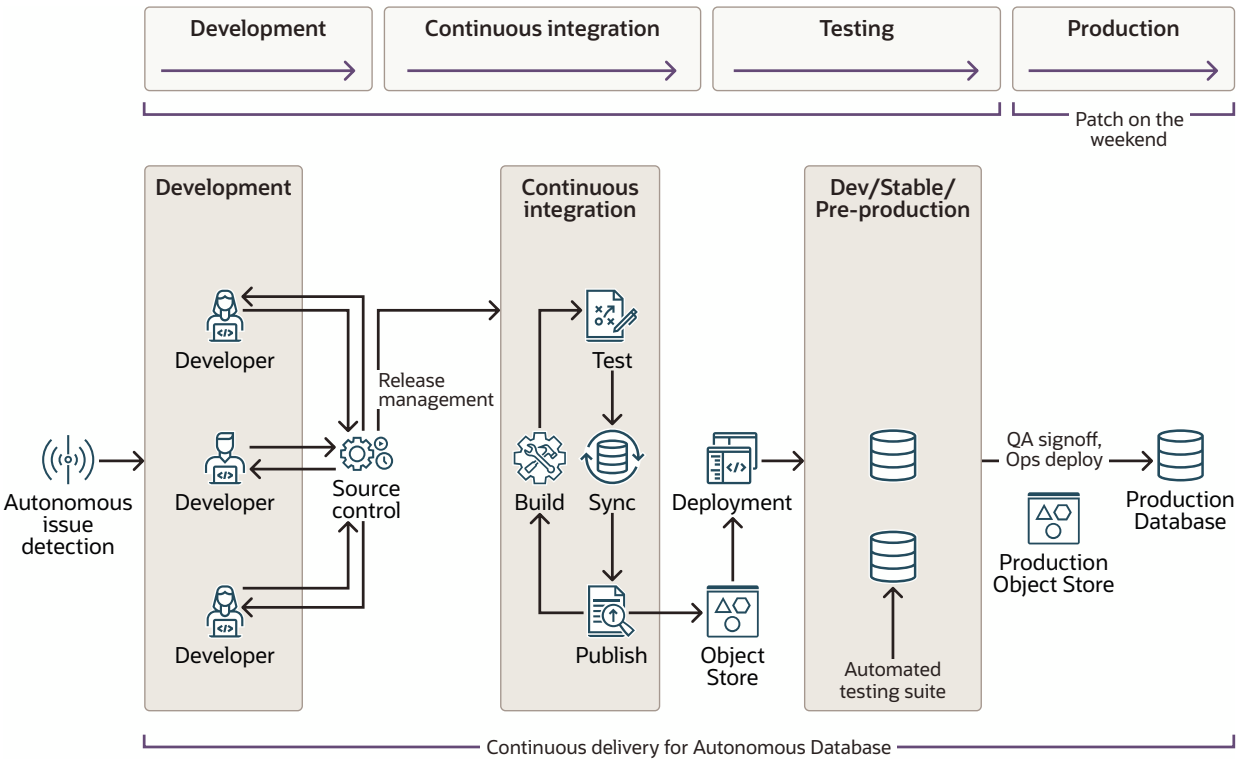
All Autonomous Database instances are automatically assigned to a maintenance window and different instances can have different maintenance windows.

Autonomous Database uses these maintenance windows to patch the entire stack used to run your database, including the database software, database dictionary, operating systems, Exadata storage, firmware, and more.

Patches include bug fixes, security fixes, and new features. Critical security fixes are always applied as soon as they are available. Patches are deployed uniformly across all databases, so you do not need to track one-off patches. After a fix for an issue is implemented, for example an issue seen in one database, the fix is deployed to all Autonomous Database instances.

All patches undergo a rigorous testing and validation process that is part of a continuous integration and development pipeline. Fixes are validated in multiple stages and environments before they are deployed to Early patch level and Always Free databases, followed by Regular patch level databases. This pipeline allows regressions to be caught and fixed before the patch is deployed to all databases. In the unlikely case that patching introduces a regression, Oracle has processes to mitigate the problem, including actions such as the following:

- Rolling back either a subset of the patch or the entire patch.
- Setting database parameters to disable the patch that introduced the regression.
- Applying an online patch to fix the problem without any downtime or connection disruption.



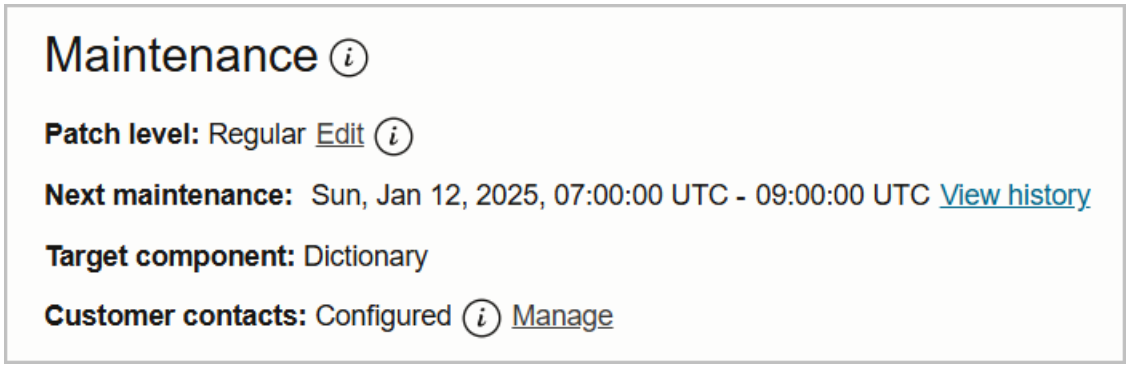
Automatic Regression Detection for Autonomous Database provides proactive handling of regressions and enables automated issue detection, diagnosis, and mitigation. This can reduce or eliminate the need for you to manually investigate issues and file service requests. Automatic Regression Detection monitors all databases, in both the Early and

Regular patch levels, but it is especially useful when you test a workload on an **Early** patch level database. If Automatic Regression Detection finds an issue and identifies a regression, it automatically reports the issue with detailed information to diagnose it. This automated reporting, as part of the continuous delivery patching cycle, allows Oracle to mitigate or fix the issue before changes reach production databases (**Regular** patch level databases). Automatic Regression Detection may not be able to find every issue; in cases where you see issues yourself, you can file a service request.

Automatic Regression Detection includes the following:

- Automatic Regression Detection monitors the database and, in the event of an incident, files a bug for the incident with detailed diagnostic information extracted from the Automatic Workload Repository.
- The incident reporting system produces notifications to Oracle Cloud Infrastructure Operations and Development teams with Oracle Automatic Incident Monitoring.
- Issues are mitigated by analyzing the Automatic Regression Detection alerts.
- Learning and improvements are made to Automatic Regression Detection on an ongoing basis.

The Autonomous Database Details page shows the **Patch level** field and the **Next maintenance** field that includes the date and time for the upcoming maintenance window; the date is updated automatically when the next maintenance window is scheduled. The **View history** link provides details on past maintenance. The **Target component** field shows the component to be updated in the next maintenance window.



When Autonomous Data Guard is enabled the console also shows maintenance information for a local standby database.

The Maintenance area includes the following information:

Maintenance Field	Description
Patch level	Shows the patch level for the instance. There are two patch level options: **Regular** and **Early**. Click **Edit** to manage patch level settings. See Set the Patch Level for more information.
Next maintenance	Specifies the time period for the next scheduled maintenance window. Click **View history** to see details on past maintenance. See View Maintenance Event History for more information.

Maintenance Field	Description
Target component	Lists the target components for the upcoming maintenance window. Possible values are: • Database: when the patch applies to the database home and executables. • Dictionary: when the patch applies to the data dictionary and to Oracle APEX. • Infrastructure: when the patch applies to the Exadata hardware or to the Grid Infrastructure.
Next maintenance (local peer)	Specifies the time period for the next scheduled maintenance window for a local Autonomous Data Guard standby. Click **View history** to see details on past maintenance.
Target component (local peer)	Lists the target components for the upcoming maintenance window on the Autonomous Data Guard. Possible values are: • Database: when the patch applies to the database home and executables. • Dictionary: when the patch applies to the data dictionary and to Oracle APEX. • Infrastructure: when the patch applies to the Exadata hardware or to the Grid Infrastructure.
Customer contacts	When customer contacts are set, Oracle sends notifications to the specified email addresses for Autonomous Database service-related issues. See View and Manage Customer Contacts for Operational Issues and Announcements for more information.

Notes for scheduled maintenance and patching:

- The Autonomous Database operations team never accesses your data unless you explicitly grant permission through a Service Request for a specified duration.
- If your database is in the stopped state during the maintenance window, the database changes from the patch are applied when you start your database.
- Your database remains available during the maintenance window. New connections to the database will always succeed. Your existing database connections may get disconnected briefly, depending on the component being patched; however, you can immediately reconnect and continue using your database:
  - For Database patches, existing connections may get disconnected if they are running for longer than the drain time after patching starts.
  - For Infrastructure patches, existing connections may get disconnected if they are running for longer than the drain time after patching starts.
  - For Dictionary patches, existing connections may get disconnected if they are holding locks on the dictionary objects being patched. Otherwise, the existing connections will not be impacted. For example, if your application is running a procedure in the `DBMS_CLOUD` package during patching and the package needs to be patched, the session using that package may get disconnected.
- You can use Oracle Cloud Infrastructure Events to be notified when maintenance begins and ends. See Information Events on Autonomous Database for more information.
- If you want to change the assigned maintenance window to a different 2-hour window on either Saturday or Sunday in the region's local time zone, file a Service Request at Oracle Cloud Support.

If you want a specific time period for your maintenance window on either Saturday or Sunday in the region's local time zone, you can request the time period with the same Service Request. If you request a specific time period for your maintenance window, the change can only be made if the time period you request is available for your database.

- If your database's allocated storage is 384 TB, you can choose a custom 2-hour window by filing a Service Request at Oracle Cloud Support (that is you can file a Service Request to request a specific day and time period on either Saturday or Sunday in the region's local time zone for your maintenance window).

See Test Workloads Against an Upcoming Patch for information on capturing a workload from a production database and replaying the workload on a target early patch level refreshable clone.

See Zero-Regression Service Level Objective for details on the zero-regression SLO for Autonomous Database.

View Maintenance Event History

You can view Autonomous Database maintenance event history for details about past maintenance events, such as the title, state, start time, and stop time.

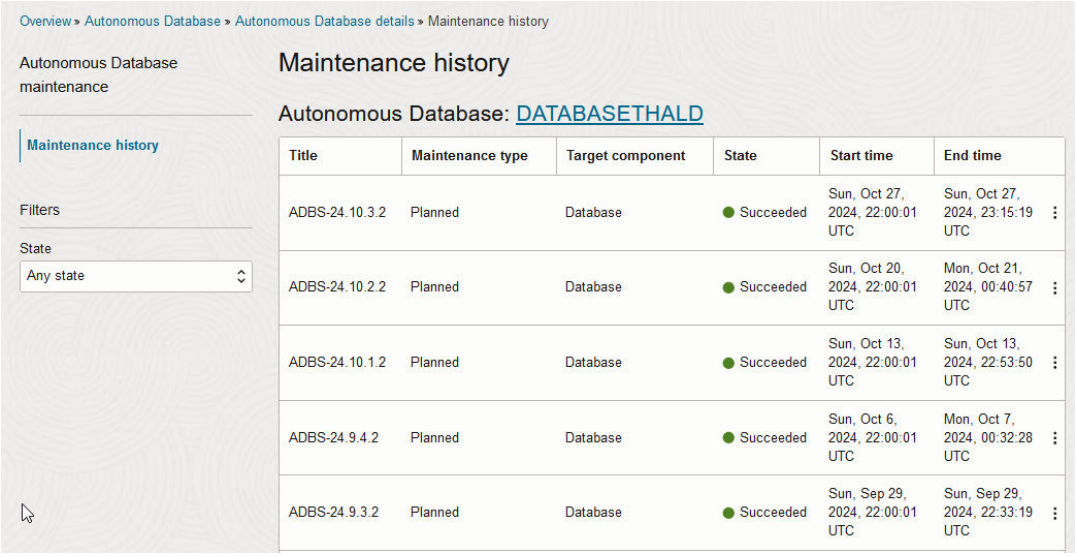
Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To view maintenance history, do the following:

1. On the Autonomous Database Details page, under **Maintenance**, click **View history**.
2. The Oracle Cloud Infrastructure Console displays the Maintenance History page.
3. (Optional) Use the **State** selector to filter events by state.

   For example, if you select **Succeeded**, the Maintenance History page shows only the maintenance events with the Succeeded state.



The Maintenance History page shows details for each maintenance event, including the following:

Field	Description
Title	The name of the maintenance event.
Maintenance type	Planned or Unplanned.
Target component	The type of the resource on which the maintenance event occurs: Database, Dictionary, or Infrastructure.
State	Succeeded, Failed, or In progress.
Start Time	Maintenance start time.
End Time	Maintenance end time.

> **Note:**
>
> Maintenance event history is available starting with maintenance events after February 2021.

View Patch Details

You can view Autonomous Database patch information, including a list of resolved issues and components.

The `DBA_CLOUD_PATCH_INFO` view provides patch information related to reported bugs (bugs reported by a customer). You can use this information to determine if a bug you reported is fixed and the patch version where the fix was applied to your Autonomous Database instance. If there were no customer bugs in a patch, `DBA_CLOUD_PATCH_INFO` does not include any rows for that patch.

To view patch information for a specific patch, do the following:

1. Select the Autonomous Database patch that you want to view. The Maintenance History page on the Oracle Cloud Infrastructure Console shows the list of patches under **Title**.
2. For the selected patch, query the `DBA_CLOUD_PATCH_INFO` view.

For example, for patch `ADBS-21.7.1.1`, use the following query:

```
SELECT * FROM DBA_CLOUD_PATCH_INFO WHERE PATCH_VERSION = 'ADBS-21.7.1.1';
```

3. For an issue of interest, query the view to obtain details for the issue:

```
SELECT * FROM DBA_CLOUD_PATCH_INFO WHERE PATCH_VERSION = 'ADBS-21.7.1.1'
and BUG_NUM = bug_number;
```

To view patch information for all available patches:

```
SELECT * FROM DBA_CLOUD_PATCH_INFO;
```

Notes for viewing patch information:

- The view `DBA_CLOUD_PATCH_INFO` is available to the ADMIN user.
- Patch information and details on resolved issues is available from `ADBS-21.7.1.1` onwards (starting in July 2021).
- The view `DBA_CLOUD_PATCH_INFO` has the following columns:

```
BUG_NUM, BUG_TITLE, COMPONENT_NAME, PATCH_VERSION
```

See View Maintenance Status Notifications for details about the patches applied during maintenance.

Set the Patch Level

When you provision or clone an Autonomous Database instance you can select a patch level to apply to upcoming patches. You can also edit the patch level after an Autonomous Database instance is provisioned. There are two patch level options: **Regular** and **Early**.

When you select patch level **Early**, patches are applied for the Autonomous Database instance one week before the **Regular** scheduled patch. The **Next Maintenance** field in the Oracle Cloud Infrastructure Console reflects a maintenance window date and time based on the patch level.

The default patch level for provisioning an Autonomous Database instance is **Regular**. The default patch level for cloning is the patch level specified for the source database.

Provisioning or cloning an instance and setting the patch level to **Early** allows you to use and to test upcoming patches before they are applied to all systems. Oracle recommends you select the **Early** patch level for your development and test databases if you want to test the upcoming patches before the patches reach production. You can also test your workloads using Oracle Real Application Testing to capture a workload on a production system and replay it with **Early** patch level. See Test Workloads with Oracle Real Application Testing for more information.

To set the patch level while provisioning or cloning an instance, do the following:

- If provisioning a new instance, follow the provisioning instructions and select the patch level, either **Regular** or **Early**. See Provision an Autonomous Database Instance for more information.
- If cloning an instance, follow the cloning instructions and select a patch level, either **Regular** or **Early**. See Clone an Autonomous Database Instance for more information.

To change the patch level for an existing Autonomous Database, do the following:

1. On the **Autonomous Database Details** page, under **Maintenance**, in the Patch level field click **Edit**.

> **Note:**
>
> The Edit button may be disabled under the following circumstances:
>
> - If early patch level is not available in your region for your database version.
> - If your database has Autonomous Data Guard enabled.
> - If your database is on early patch level and it is not feasible to move it to regular patch level. In this case, you should try again after the next maintenance window.

2. Select the patch level, **Regular** or **Early**, and click **Submit**.
The time it takes to change the patch level depends on the size of your database. You may see brief connection drops during this time.

Reporting Patch Issues to Oracle Support

When you report an issue for an **Early** patch level database, Oracle Support takes the necessary actions to prevent the problem from propagating to **Regular** patch level databases. A few examples of the possible actions are:

- The patch that caused the problem is removed before the regular patch level databases are patched.
- The patch that caused the problem is disabled using database parameters when it is being applied to the regular patch-level databases.
- Patching of regular patch level databases is paused until corrective action is taken.

If you have an issue to report, file a service request at Oracle Cloud Support or contact your support representative.

Oracle provides a service level objective of zero regressions in your production database. See Zero-Regression Service Level Objective for more information.

Notes for patching level:

- The option to set the patch level is not available in every region. In some regions all Autonomous Database instances are provisioned or cloned at the **Regular** patch level.
- Autonomous Data Guard is only available for instances with patch level **Regular**. When you configure an Autonomous Database instance with patch level **Early**, you cannot enable Autonomous Data Guard.
- Always Free Autonomous Database instances do not provide the **Early** patch level option.
- When the patch level of a source Autonomous Database instance is **Regular**, in regions that support the **Early** patch level you can set the patch level of a clone to **Early**.

View Maintenance Status Notifications

The `DB_NOTIFICATIONS` view stores information about maintenance status notifications for your Autonomous Database instance.

To show notification information:

1. Connect to your Autonomous Database instance.
2. Use the following query to view maintenance (patching) information.

```
SELECT * FROM DB_NOTIFICATIONS WHERE TYPE = 'MAINTENANCE';
```

The following provides details about the `DESCRIPTION` field values.

- **Maintenance run has ended**: Specifies maintenance has completed. The `MAINTENANCE_STATUS` shows the value `COMPLETED` with the start and end timestamps for the completed maintenance in `ACTUAL_START_DATE` and `ACTUAL_END_DATE`.
- **Maintenance run is scheduled for the instance**: Specifies a new maintenance has been scheduled. The `MAINTENANCE_STATUS` shows the value `SCHEDULED` with the expected start and end timestamps for the scheduled maintenance in `EXPECTED_START_DATE` and `EXPECTED_END_DATE`.
- **Maintenance run has begun**: Specifies the maintenance is in progress and provides the start timestamp for the active maintenance. The `MAINTENANCE_STATUS` shows the value `IN_PROGRESS` and `ACTUAL_START_DATE` stores the start timestamp.

The following table shows the `DB_NOTIFICATIONS` columns and datatypes.

Column	Datatype	Description
TYPE	VARCHAR2(128)	Specifies the type of the notification. Valid value is: MAINTENANCE.
TIME	TIMESTAMP(6) WITH TIME ZONE	Time when the notification entry was added.
EXPECTED_START_DATE	TIMESTAMP(6) WITH TIME ZONE	Scheduled maintenance start time.
EXPECTED_END_DATE	TIMESTAMP(6) WITH TIME ZONE	Scheduled maintenance end time.
ACTUAL_START_DATE	TIMESTAMP(6) WITH TIME ZONE	Actual maintenance start time.
ACTUAL_END_DATE	TIMESTAMP(6) WITH TIME ZONE	Actual maintenance end time.
MAINTENANCE_PRODUCT	VARCHAR2(128)	Product/component for which maintenance is scheduled/ongoing.
MAINTENANCE_STATUS	VARCHAR2(128)	Current status of the maintenance.
DESCRIPTION	VARCHAR2(128)	The notification message details.
PATCH_ID	VARCHAR2(128)	Patch version.

View and Manage Customer Contacts for Operational Issues and Announcements

You can view and manage the Autonomous Database customer contacts.

When customer contacts are set, Oracle sends notifications to the specified email addresses for Autonomous Database service-related issues. Contacts in the customer contacts list receive unplanned maintenance notices and other notices, including but not limited to notices for database upgrades and upcoming wallet expiration. When customer contacts are not set the notifications go to the tenancy admin email address associated with the account. Oracle recommends that you set the customer contacts so that the appropriate people receive service-related notifications.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the  next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To view or manage customer contact email addresses:

1. On the Autonomous Database Details page, under **Maintenance**, in the **Customer Contacts** field, click the **Manage** link.

   This shows the Manage Customer Contacts page. On the Manage Customer Contacts page you can view the email contact list.
2. To add a contact, click **Add Customer Contact**.
3. Enter the contact email address.

   Oracle recommends adding an administrator group's email address over an individual's email address when possible, so as to not miss important notifications or announcements.
4. To add additional contacts, click **Add Customer Contact** again.

The lifecycle state changes to Updating while the customer contact list is updated.

To remove customer contacts:

1. On the Manage Customer Contacts page select the email addresses to remove.
2. To remove all email addresses, select the top column next to **Email**.
3. Click **Remove**.
4. On the Confirm remove contact page, click **Remove** to confirm.

The lifecycle state changes to Updating while the customer contact list is updated.

To edit customer contacts:

1. On the Manage Customer Contacts page select the email addresses to edit.
2. To edit all email addresses, select the top column next to **Email**.
3. Click **Edit**.
4. On the Edit Contacts page edit the contacts as needed and click **Save**.

The lifecycle state changes to Updating while the customer contact list is updated.

Notes for adding customer contacts for operational issues and announcements:

- Oracle recommends adding an administrator group's email address over an individual's email address when possible, so as to not miss important notifications or announcements.
- In addition to sending information for operational issues and announcements to customer contacts, you can subscribe to Oracle Cloud Infrastructure Events to be notified of Autonomous Database events such as when regular maintenance is scheduled or when regular maintenance starts and ends. Event notification requires a subscription for each email address, independent of the customer contacts list.

  See Use Autonomous Database Events for more information.

View Oracle Cloud Infrastructure Operations Actions

The DBA_OPERATOR_ACCESS view stores information on the actions that Oracle Cloud Infrastructure cloud operations performs on your Autonomous Database.

These steps use **Database Actions** to query the DBA_OPERATOR_ACCESS view. You can query this view from any SQL connection by connecting as a user with the proper privileges.

1. On the Autonomous Database Details page select **Database Actions** and from the quick actions list click **SQL**.
2. Within the SQL Worksheet, query the DBA_OPERATOR_ACCESS view.

   For example:

   ```
   SELECT * FROM DBA_OPERATOR_ACCESS;
   ```

Notes for the DBA_OPERATOR_ACCESS view:

- You can continuously move this audit data to the object store for purposes such as integration with Security Information and Event Management (SIEM) tools. See Create and Configure a Pipeline for Export with Timestamp Column for more information on using pipelines to continuously export query results to object store.
- The DBA_OPERATOR_ACCESS view provides information starting on October 4, 2022, the date this feature was introduced. The view does not provide information prior to October 4, 2022.
- The DBA_OPERATOR_ACCESS view may not show results if Oracle Cloud Infrastructure cloud operations hasn't performed any actions or run any statements in your Autonomous Database instance. In this case, when you query the view you will not see any rows returned.
- DBA_OPERATOR_ACCESS View
  The DBA_OPERATOR_ACCESS view describes top level statements executed in the Autonomous Database instance by Oracle Cloud Infrastructure operations.

DBA_OPERATOR_ACCESS View

The DBA_OPERATOR_ACCESS view describes top level statements executed in the Autonomous Database instance by Oracle Cloud Infrastructure operations.

Column	Datatype	Description
SQL_TEXT	VARCHAR2(4000)	SQL text of the statement executed by the operator.
EVENT_TIMESTAMP	DATE	Timestamp of the operator action in UTC (Coordinated Universal Time).
REQUEST_ID	VARCHAR2(64)	Request number related to the reason behind the operator action. This could be a bug number, an SR number, or a change ticket request number that provides information on the reason for the action.
REASON	VARCHAR2(64)	Reason for the operator action. This provides context for the reason behind the action and may have a value such as: MITIGATION, DIAGNOSTIC COLLECTION, or CUSTOMER REQUEST.

Monitor and Manage Performance

Describes the Database Dashboard card in Database Actions where you can find information about the performance of an Autonomous Database instance.

- Monitor the Performance of Autonomous Database with Database Actions
  The **Database Dashboard** card in Database Actions provides the **Overview** and **Monitor** tabs where you can find information about the performance of an Autonomous Database instance.
- Monitor Autonomous Database with Performance Hub
  Use Performance Hub to monitor your database for a defined time period and download statistical reports. Performance Hub also lets you view real-time and historical performance data for an Autonomous Database instance.
- Manage Concurrency and Priorities on Autonomous Database
  Concurrency and prioritization of user requests in Autonomous Database is determined by the database service the user is connected with, and whether compute auto scaling is enabled.
- Manage CPU/IO Shares on Autonomous Database
  Autonomous Database comes with predefined CPU/IO shares assigned to different consumer groups. You can modify these predefined CPU/IO shares if your workload requires different CPU/IO resource allocations.
- Manage Runaway SQL Statements on Autonomous Database
  Specifies how you configure Autonomous Database to terminate SQL statements automatically based on their query runtime or their IO usage.
- Use Fast Ingest on Autonomous Database
- Enable Database Management Service to Monitor and Manage Autonomous Database
  You can use the Diagnostics & Management features in the Database Management service for advanced monitoring and management of Autonomous Databases.
- Perform SQL Tracing on Autonomous Database
  Use SQL tracing to help you identify the source of an excessive database workload, such as a high load SQL statement in your application.
- Monitor Performance with Autonomous Database Metrics
  You can monitor the health, capacity, and performance of your databases with metrics, alarms, and notifications. You can use Oracle Cloud Infrastructure Console or Monitoring APIs to view metrics.
- Available Metrics: oci_autonomous_database
  This topic describes the metrics emitted by the Database service in the `oci_autonomous_database` namespace.

Monitor the Performance of Autonomous Database with Database Actions

The **Database Dashboard** card in Database Actions provides the **Overview** and **Monitor** tabs where you can find information about the performance of an Autonomous Database instance.

- Use Database Actions to Monitor Activity and Utilization
  Database Actions provides the **Database Dashboard** card, with **Overview** and **Monitor** tabs to provide real-time and historical information about the utilization of an Autonomous Database instance.

- Use Database Actions to Monitor Active Session History Analytics and SQL Statements
  In Database Actions the Performance Hub card provides information about Active Session History (ASH) analytics and current and past monitored SQL statements.

Use Database Actions to Monitor Activity and Utilization

Database Actions provides the **Database Dashboard** card, with **Overview** and **Monitor** tabs to provide real-time and historical information about the utilization of an Autonomous Database instance.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.

To view the **Overview** tab that shows general information about utilization, do the following:

1. Access Database Actions as the ADMIN user.

   See Access Database Actions as ADMIN for more information.
2. On the Database Actions Launchpad, under **Monitoring**, click **Database Dashboard**.

> **Note:**
>
> You can bookmark the Launchpad URL and go to that URL directly without logging in to the Oracle Cloud Infrastructure Console. If you logout and use the bookmark, then you need to enter the `ADMIN` username and *password*, and click **Sign in**. See Set the ADMIN Password in Autonomous Database if you need to change the password for the `ADMIN` user.

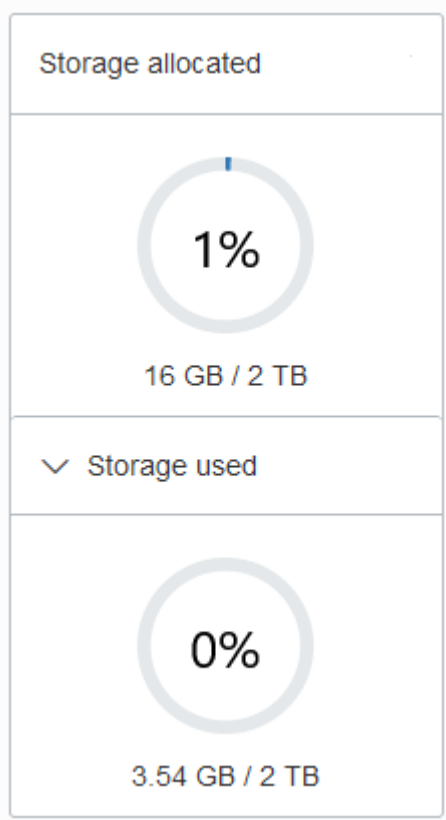
- Database Dashboard Overview
  The **Overview** tab shows real-time and historical information about the Autonomous Database utilization.
- Database Dashboard Activity
  The **Monitor** tab shows real-time and historical information about the Autonomous Database performance data, activity, and utilization.

Database Dashboard Overview

The **Overview** tab shows real-time and historical information about the Autonomous Database utilization.

The charts shown on this page include:

- **Storage:** This chart shows the provisioned, allocated, and used storage. The chart indicates what percentage of the space is currently in-use.



Provisioned storage is the amount of storage you select when you provision the instance or when you modify storage by scaling storage.

Storage allocated is the amount of storage physically allocated to all data tablespaces and temporary tablespaces and includes the free space in these tablespaces. This does not include storage for the sample schemas.

Storage used is the amount of storage actually used in all data and temporary tablespaces. This does not include storage for the sample schemas. The storage used is the storage in the Autonomous Database as follows:

- Storage used by all database objects. Note: the chart does not include storage for the sample schemas as they do not count against your storage.
- Storage for files users put in the file system.
- Storage used by temporary tablespaces.
- Used storage excludes the free space in the data and temporary tablespaces.

By default the chart does not show the used storage. Select **Storage used** to expand the chart to see used storage (the values are calculated when you open the chart).

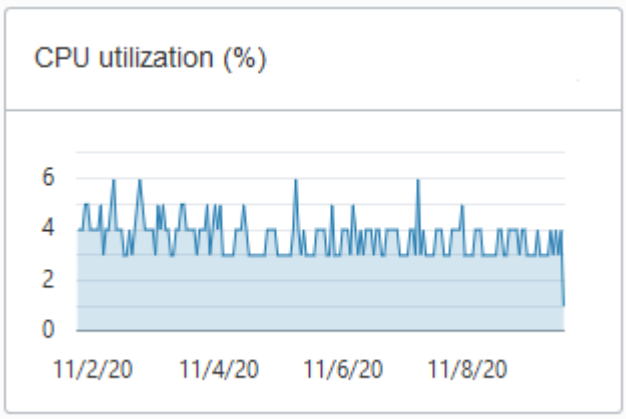
For an Autonomous JSON Database the chart shows an additional field showing the percentage of storage used that is not storing JSON documents.

> **Note:**
>
> If you drop an object, the space continues to be consumed until you empty the recycle bin. See Purging Objects in the Recycle Bin for more information.

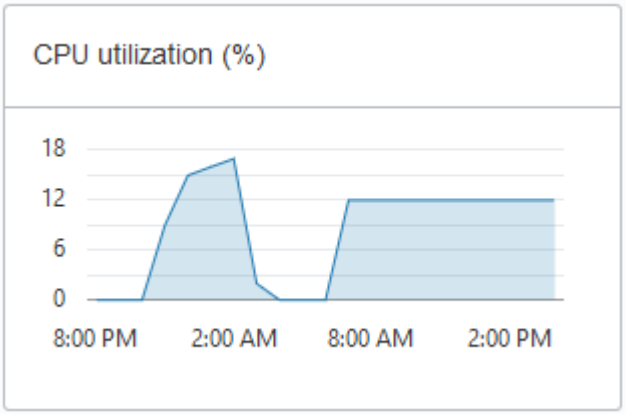
See Use Sample Data Sets in Autonomous Database for information on sample schemas SH and SSB.

- **CPU utilization (%) for ECPU Compute Model:** This chart shows the historical CPU utilization of the service:
  - Compute auto scaling disabled: this chart shows hourly data. A data point shows the average CPU utilization for that hour. For example, a data point at 10:00 shows the average CPU utilization for 9:00-10:00.

    The utilization percentage is reported with respect to the number of ECPUs the database is allowed to use. For example, if the database has four (4) ECPUs, the percentage in this graph is based on 4 ECPUs.

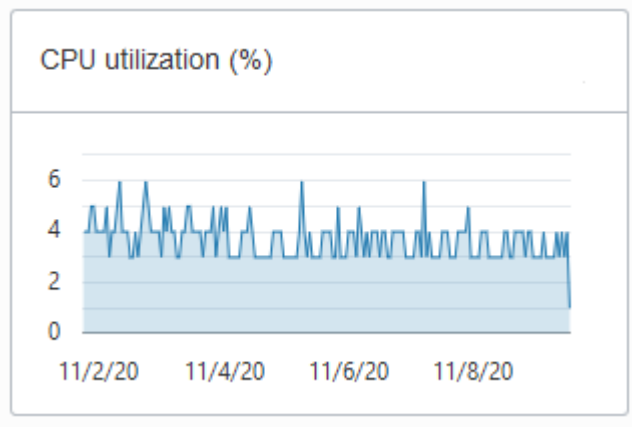


  - Compute auto scaling enabled: For databases with compute auto scaling enabled the utilization percentage is reported with respect to the maximum number of ECPUs the database is allowed to use, which is three times the number of ECPUs. For example, if the database has four (4) ECPUs with auto scaling enabled, the percentage in this graph is based on 12 ECPUs.

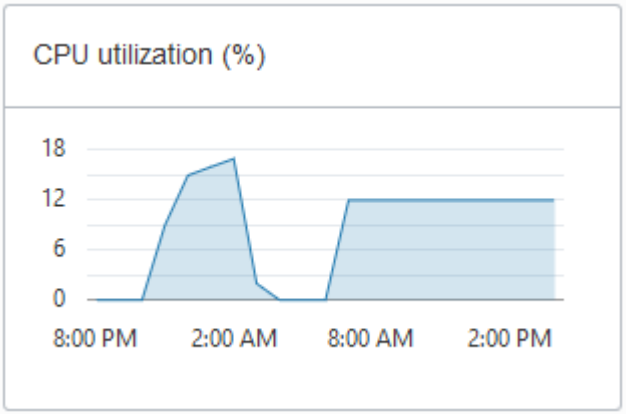


- **CPU utilization (%) for OCPU Compute Model:** This chart shows the historical CPU utilization of the service:
  - Compute auto scaling disabled: this chart shows hourly data. A data point shows the average CPU utilization for that hour. For example, a data point at 10:00 shows the average CPU utilization for 9:00-10:00.

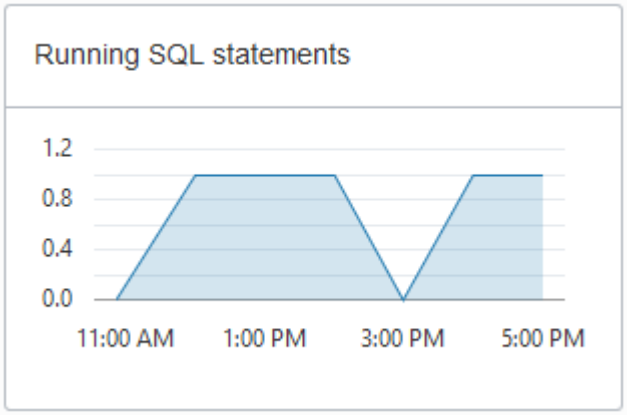
    The utilization percentage is reported with respect to the number of OCPUs the database is allowed to use. For example, if the database has four (4) OCPUs, the percentage in this graph is based on 4 OCPUs.



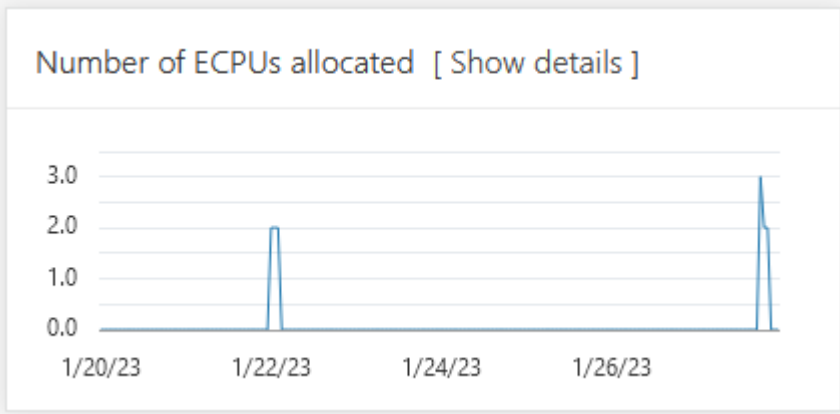
  - Compute auto scaling enabled: For databases with compute auto scaling enabled the utilization percentage is reported with respect to the maximum number of OCPUs the database is allowed to use. For example, if the database has four (4) OCPUs with auto scaling enabled, the percentage in this graph is based on 12 OCPUs.



- **Running SQL statements**: This chart shows the average number of running SQL statements historically. This chart shows hourly data. A data point shows the running SQL statements for that hour. For example, a data point at 10:00 shows the average number of running SQL statements for 9:00-10:00.



- **Number of ECPUs allocated (only shown for ECPU Compute Model):**



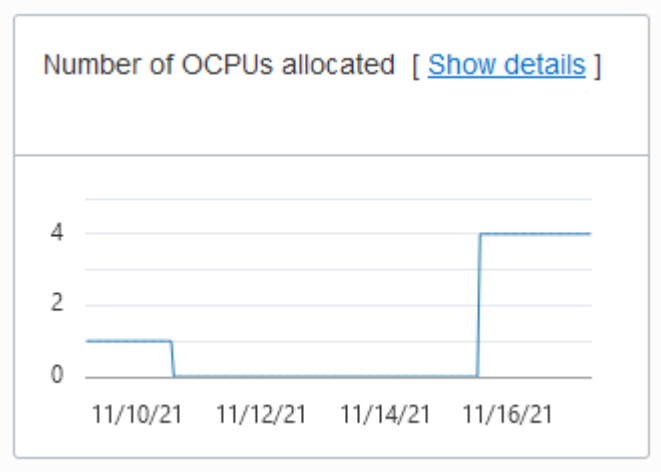
Notes for display results:

- Compute auto scaling disabled: For databases with compute auto scaling disabled, for each hour the chart shows the number of ECPUs allocated to the database if the database is open for at least some part of the hour.
- Compute auto scaling enabled: For databases with compute auto scaling enabled, for each hour the chart shows the average number of ECPUs used during that hour if that value is higher than the number of ECPUs provisioned. If the number of ECPUs used is not higher than the number of ECPUs provisioned, then the chart shows the number of ECPUs allocated for that hour.
- Stopped Database: If the database was stopped for the full hour the chart shows 0 ECPUs allocated for that hour.

Click **Show details** for more information, including the number of ECPUs allocated to the database and to external resources, and the total allocated ECPUs.

The Show details view includes separate values for database ECPU usage and external resource ECPU usage. External resources include: Cloud SQL, Graph, OML4PY, and others. The Total ECPUs are the total number of ECPUs in use on the Autonomous Database. The external ECPUs value shows how external ECPUs contribute to the total ECPU usage.

- **Number of OCPUs allocated (only shown for OCPU Compute Model):**



Notes for display results:

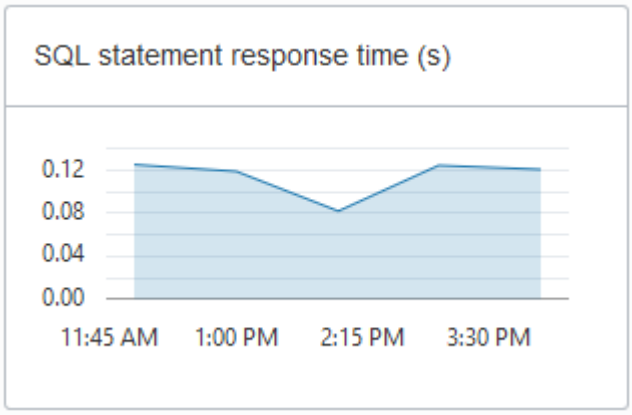
- Compute auto scaling disabled: For databases with compute auto scaling disabled, for each hour the chart shows the number of OCPUs allocated to the database if the database is open for at least some part of the hour.

  - Compute auto scaling enabled: For databases with compute auto scaling enabled, for each hour the chart shows the average number of OCPUs used during that hour if that value is higher than the number of OCPUs provisioned. If the number of OCPUs used is not higher than the number of OCPUs provisioned, then the chart shows the number of OCPUs allocated for that hour.
  - Stopped Database: If the database was stopped for the full hour the chart shows 0 OCPUs allocated for that hour.

  Click **Show details** for more information, including the number of OCPUs allocated to the database and to external resources, and the total allocated OCPUs.

  The Show details view includes separate values for database OCPU usage and external resource OCPU usage. External resources include: Cloud SQL, Graph, OML4PY, and others. The Total OCPUs are the total number of OCPUs in use on the Autonomous Database. The external OCPUs value shows how external OCPUs contribute to the total OCPU usage.

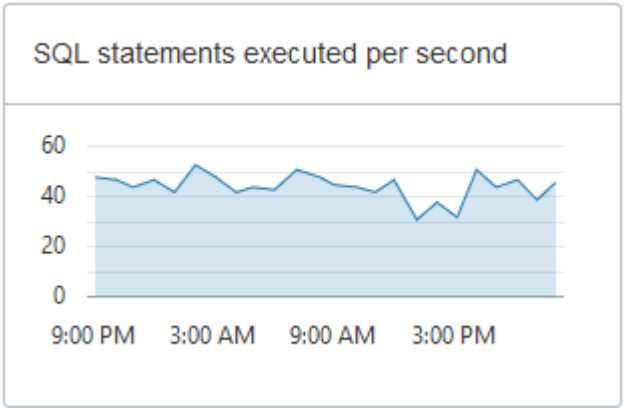
- **SQL statement response time (s)**: This chart shows the average response time, in seconds, of SQL statements historically. This chart shows hourly data. A data point shows the average SQL statement response time for that hour. For example, a data point at 10:00 shows the average SQL statement response time, in seconds, for the hour from 9:00-10:00.



- **SQL statements executed per second**

> **Note:**
>
> Database Dashboard does not show this chart when the Autonomous Database instance workload type is **Data Warehouse**.

.



The default retention period for performance data is thirty (30) days. The CPU utilization, running statements, and average SQL response time charts show data for the last eight (8) days by default.

You can change the retention period by modifying the Automatic Workload Repository retention setting with the PL/SQL procedure `DBMS_WORKLOAD_REPOSITORY.MODIFY_SNAPSHOT_SETTINGS()`. The maximum retention you can set is 30 days. See *Oracle Database PL/SQL Packages and Types Reference*.

If you need to store more performance data you can use the OCI Ops Insights AWR Hub. See Analyze Automatic Workload Repository (AWR) Performance Data for more information.

Database Dashboard Activity

The **Monitor** tab shows real-time and historical information about the Autonomous Database performance data, activity, and utilization.

> **Note:**
>
> The default view in the **Monitor** tab is real-time. This view shows performance data for the last hour.

The charts on this page are:

- **Database Activity**

  This chart shows the average number of sessions in the database using CPU or waiting on a wait event. See *Oracle Database Reference* for more information on wait events.

- **CPU Utilization** (with ECPU compute model)

  This chart shows the CPU utilization of each consumer group. The utilization percentage is reported with respect to the number of ECPUs the database is allowed to use. For example, if the database has four (4) ECPUs, the percentage in this graph is based on 4 ECPUs.

  For databases with compute auto scaling enabled the utilization percentage is reported with respect to the maximum number of ECPUs the database is allowed to use, which is three times the number of ECPUs. For example, if the database has four (4) ECPUs with auto scaling enabled, the percentage in this graph is based on 12 ECPUs.

See Manage Concurrency and Priorities on Autonomous Database for detailed information on consumer groups.

- **CPU Utilization** (with OCPU compute model)

  This chart shows the CPU utilization of each consumer group. The utilization percentage is reported with respect to the number of OCPUs the database is allowed to use. For example, if the database has four (4) OCPUs, the percentage in this graph is based on 4 OCPUs.

  For databases with compute auto scaling enabled the utilization percentage is reported with respect to the maximum number of OCPUs the database is allowed to use. For example, if the database has four (4) OCPUs with auto scaling enabled, the percentage in this graph is based on 12 OCPUs.

  See Manage Concurrency and Priorities on Autonomous Database for detailed information on consumer groups.

- **Running Statements**

  This chart shows the average number of running SQL statements in each consumer group.

  See Manage Concurrency and Priorities on Autonomous Database for detailed information on consumer groups.

- **Queued Statements**

  This chart shows the average number of queued SQL statements in each consumer group.

  See Manage Concurrency and Priorities on Autonomous Database for detailed information on consumer groups.

To see earlier data click **Time period**. The default retention period for performance data is thirty (30) days. By default in the Time Period view the charts show information for the last eight (8) days.

In the time period view you can use the calendar to look at a specific time period in the past 30 days. You can also use the time slider to change the period for which performance data is shown.

> **Note:**
>
> The retention time can be changed by changing the Automatic Workload Repository retention setting with the PL/SQL procedure `DBMS_WORKLOAD_REPOSITORY.MODIFY_SNAPSHOT_SETTINGS`. Be aware that increasing the retention time results in more storage usage for performance data. See *Oracle Database PL/SQL Packages and Types Reference*.

Use Database Actions to Monitor Active Session History Analytics and SQL Statements

In Database Actions the Performance Hub card provides information about Active Session History (ASH) analytics and current and past monitored SQL statements.

Perform the following prerequisite step as necessary:

- Access Database Actions as the ADMIN user. See Access Database Actions a ADMIN for more information.
- The Performance Hub Page

The Performance Hub Page

The Performance Hub page shows performance data for a time period you specify.
To navigate to the Performance Hub page, do either of the following:

- In the Launchpad page, select the **Monitoring** tab and click **Performance Hub**.
- Click **Selector** to display the navigation menu. Under Monitoring, select **Performance Hub**.

> **Note:**
>
> The Performance Hub page is available in the following user interface languages: French, Japanese, Korean, Traditional Chinese, and Simplified Chinese. If you change the language to German, Spanish, Italian, or Portuguese in Preferences, the Performance Hub page reverts to English.

The Performance Hub page consists of these parts:

- **Time Range Area**: Use the controls in time range area at the top of the page to specify the time period for which you want to view performance data.
- **ASH Analytics Tab**: Use this tab to explore ASH (Active Session History) information across a variety of different dimensions for the specified time period.
- **SQL Monitoring Tab**: Use this tab to view the top 100 SQL statement executions by different dimensions for the specified time period, and to view details of SQL statement executions you select.

See Performance Hub in the Database service documentation for more information.

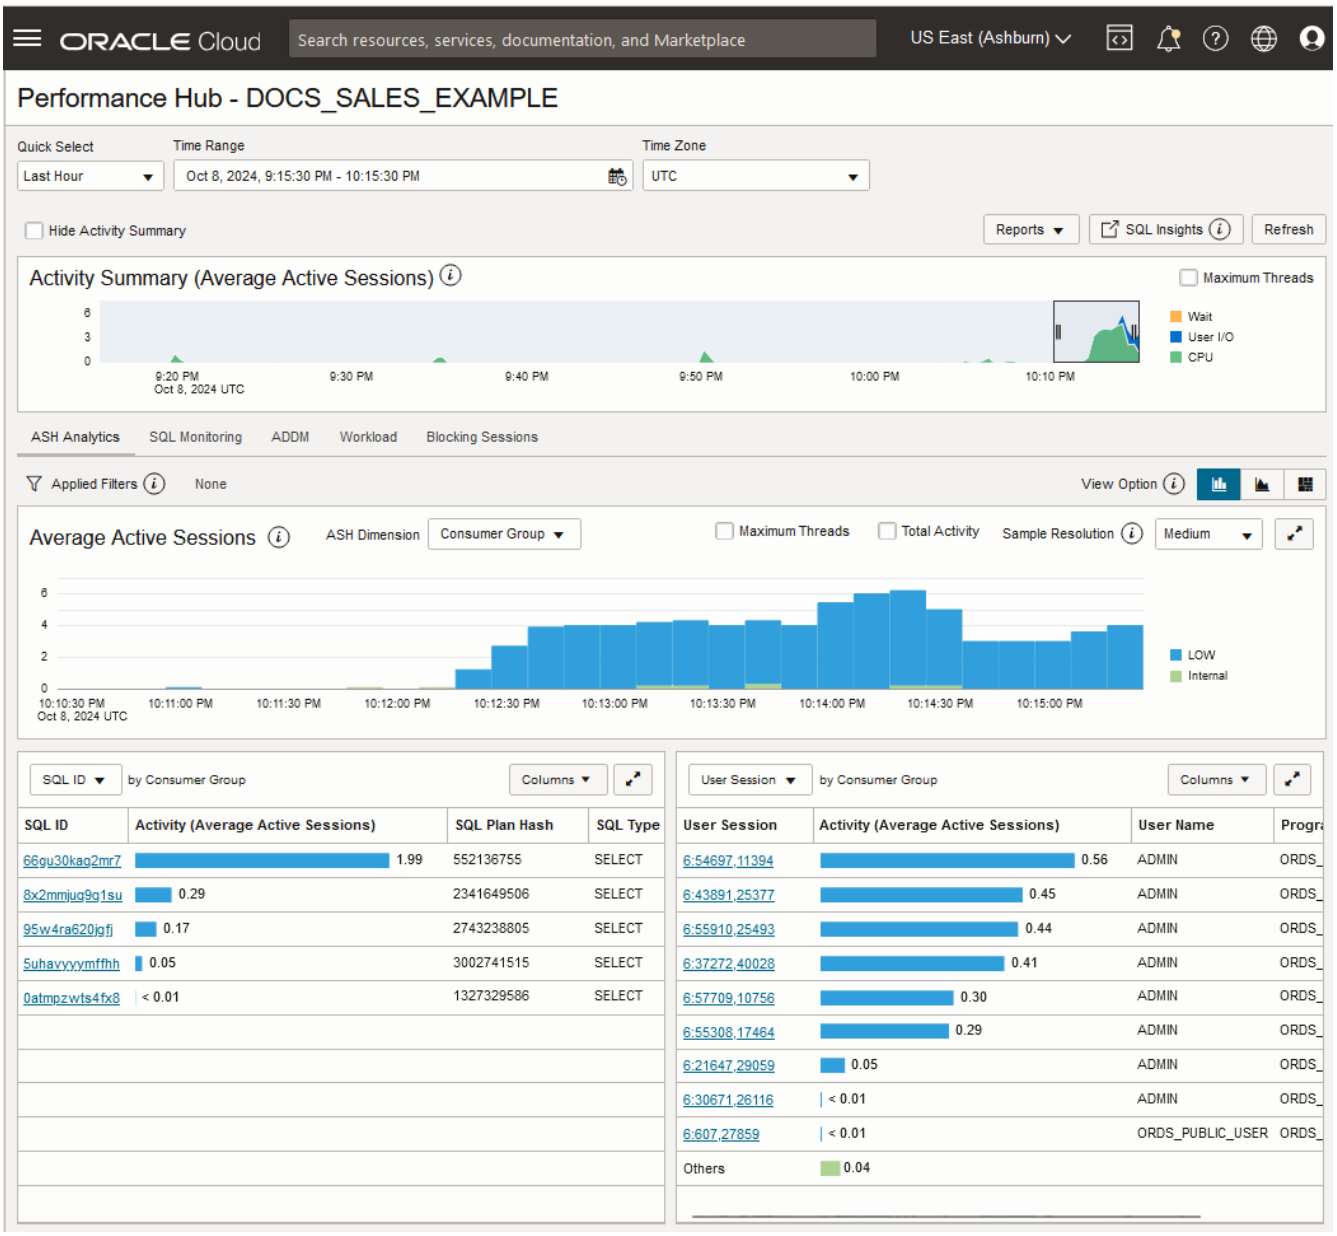
Monitor Autonomous Database with Performance Hub

Use Performance Hub to monitor your database for a defined time period and download statistical reports. Performance Hub also lets you view real-time and historical performance data for an Autonomous Database instance.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database**, and then click Autonomous Database.

1. On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.
2. From the Autonomous Database details page click **Performance Hub**.



See Using Performance Hub to Analyze Database Performance more information.

Component	Description
Top Area	The top of the Performance Hub page shows the following: • **Quick Select**: Use to quickly set the time range to one of Last Hour, Last 8 Hours, Last 24 Hours, or Last Week. • **Time Range**: Use to set the date and time range for Performance Hub to monitor. • **Time Zone**: Select an entry from this list to specify times based on Coordinated Universal Time (UTC), your local web browser time (Browser), or the database time zone setting (Database). • **Activity Summary (Average Active Sessions)**: shows active sessions during the selected time range. It displays the average number of active sessions broken down by CPU, User I/O, and Wait. It also shows the Max CPU usage. You can hide or show the activity summary by selecting **Hide Activity Summary**. • **Reports**: Allows you to generate and view different types of performance diagnostic reports, including: – **Automatic Workload Repository** (AWR): This report is available by default. The AWR report collects, processes, and maintains performance statistics for problem detection and self-tuning purposes. This data is both in memory and stored in the database. – **Active Session History** (ASH): This report requires that you enable Database Management. The ASH report contains data that helps you triage transient performance problems and perform scoped or targeted performance analysis by various dimensions such as time, session, module, action, or SQL identifier. See Enable Database Management Service to Monitor and Manage Autonomous Database for more information.
Active Session History (ASH) Analytics	**ASH Analytics** is displayed by default. This tab shows Active Session History (ASH) analytics charts to explore the Active Session History data. You can drill down into database performance across multiple dimensions such as Consumer Group, Wait Class, SQL ID, and User Name. Select an Average Active Sessions dimension and view the top activity for that dimension for the selected time period. For more information, see ASH Analytics.
SQL Monitoring	SQL statements are only monitored if they've been running for at least five seconds or if they're run in parallel. The table displays monitored SQL statement executions by dimensions including Last Active Time, CPU Time, and Database Time. The table displays currently running SQL statements and SQL statements that completed, failed, or were terminated. The columns in the table provide information for monitored SQL statements including Status, Duration, and SQL ID. For more information, see SQL Monitoring.
Automatic Database Diagnostic Monitor (ADDM)	**ADDM** tab provides access to analysis information gathered by the Automatic Database Diagnostic Monitor (ADDM) tool. ADDM analyzes AWR (Automatic Workload Repository) snapshots on a regular basis, locates root causes of any performance problems, provides recommendations for correcting the problems, and identifies non-problem areas of the system. AWR is a repository of historical performance data, therefore ADDM can analyze performance issues after the event, often saving time and resources in reproducing a problem. For more information, see Automatic Database Diagnostic Monitor (ADDM).

Component	Description
Workload	Use the **Workload** tab to visually monitor the database workload to identify spikes and bottlenecks. This tab shows four chart areas that show database workload in various ways: • **CPU Statistics**: Charts CPU usage. • **Wait Time Statistics**: Shows the wait time across the database's foreground sessions, divided by wait classes. • **Workload Profile**: Charts user (client) workload on the database. • **Sessions**: Shows the number of sessions, successful logons, and current logons. For more information, see Workload.
Blocking Sessions	**Blocking Sessions** hierarchically lists sessions that are waiting or are blocked by sessions that are waiting. You can set the minimum wait time required for sessions to be displayed in the list, and you can view a variety of information about a session to determine whether to let it continue or to end it. For more information, see Blocking Sessions.

See Using Performance Hub to Analyze Database Performance more information.

Manage Concurrency and Priorities on Autonomous Database

Concurrency and prioritization of user requests in Autonomous Database is determined by the database service the user is connected with, and whether compute auto scaling is enabled.

- Database Service Names for Autonomous Database
  Autonomous Database comes with predefined database services with different performance and concurrency characteristics. You must choose one of these services for each application or user that connects to the database.
- Idle Time Limits
  Autonomous Database has predefined idle time limits for sessions so that idle sessions do not hold system resources for a long time.
- Service Concurrency
  The consumer groups of the predefined service names provide different levels of performance and concurrency. The available service names are different depending on your workload: Data Warehouse, Transaction Processing, or JSON Database.
- Change MEDIUM Service Concurrency Limit (ECPU Compute Model)
  If your application requires customized concurrency, you can modify the concurrency limit for your Autonomous Database MEDIUM service.
- Change MEDIUM Service Concurrency Limit (OCPU Compute Model)
  If your application requires customized concurrency, you can modify the concurrency limit for your Autonomous Database MEDIUM service.
- Predefined Job Classes with Oracle Scheduler
  Autonomous Database includes predefined `job_class` values to use with Oracle Scheduler.

Database Service Names for Autonomous Database

Autonomous Database comes with predefined database services with different performance and concurrency characteristics. You must choose one of these services for each application or user that connects to the database.

The following table compares the database services. This table applies to databases with 2 OCPUs or more, or 4 ECPUs or more. You cannot use parallelism in databases that have fewer than 2 OCPUs or fewer than 4 ECPUs unless you use manual parallelism, as explained below.

Service Name	Degree of Parallelism Footnote 1	Resource Shares Footnote 2	Concurrent Statements Footnote 3	
			Without Compute Auto Scaling	**With Compute Auto Scaling**
`TPURGENT` Footnote 5	Set manually Footnote 4	12	Bounded by the sessions parameter. **OCPU model**: The sessions parameter is set to 300 times the number of base OCPUs. **ECPU model**: The sessions parameter is set to 75 times the number of base ECPUs.	Same values as without Auto Scaling
`TP` Footnote 5	No parallelism	8	Bounded by the sessions parameter. **OCPU model**: The sessions parameter is set to 300 times the number of base OCPUs. **ECPU model**: The sessions parameter is set to 75 times the number of base ECPUs.	Same values as without Auto Scaling
`HIGH`	Enabled **OCPU model**: The degree of parallelism will be equal to the number of base OCPUs **ECPU model**: The degree of parallelism will be half the number of base ECPUs rounded down to the nearest integer value.	4	3	9

Service Name	Degree of Parallelism Footnote 1	Resource Shares Footnote 2	Concurrent Statements Footnote 3	
			Without Compute Auto Scaling	**With Compute Auto Scaling**
MEDIUM	Enabled **OCPU model**: The degree of parallelism will be 2 with 2 OCPUs, 3 with 3 OCPUs, and 4 for more than 3 OCPUs. **ECPU model**: The degree of parallelism will be half the number of base ECPUs rounded down to the nearest integer value with 4 to 7 ECPUs. It will be 4 with 8 ECPUs or more.	2	**OCPU model**: 1.25 × number of base OCPUs **ECPU model**: 0.25125 × number of base ECPUs A decimal result is truncated.	**OCPU model**: 3.75 × number of base OCPUs **ECPU model:** 0.75375 × number of base ECPUs A decimal result is truncated.
LOW	No parallelism	1	Bounded by the sessions parameter. **OCPU model**: The sessions parameter is set to 300 times the number of base OCPUs. **ECPU model**: The sessions parameter is set to 75 times the number of base ECPUs.	Same values as without Auto Scaling.

Footnote 1

The HIGH and MEDIUM services have Parallel Query, DDL, and DML enabled by default.

If you are using Autonomous Database with Oracle Database 19c, parallel DML requires you to commit or rollback your transaction before you can access the same table again. If you try to access the same table without committing or rolling back the transaction, you will get the ORA-12839 error. This restriction has been lifted in Oracle Database 23ai, so you can access the same table again using the same session without a commit or rollback.

If needed, you can disable parallel DML in your session by running:

```
alter session disable parallel dml;
```

You can also create a logon trigger to run this statement for all your sessions. Note that the performance of your large DML statements may be impacted when you disable parallel DML.

Note that these degree of parallelism values may be doubled for simple queries like a query on a single table.

The parallelism hints in your SQL statements are ignored by default in Autonomous Database instances with the Data Warehouse workload type. The parallelism hints in your SQL statements are honored in Autonomous Database instances with Transaction Processing, JSON, or APEX workload types. If you want to control parallelism using hints in an Autonomous Database instance with the Data Warehouse workload, you can enable hints using the following command.

```
alter session set optimizer_ignore_parallel_hints=false;
```

See the following for more information:

- About Parallel DML Operations if you are using Autonomous Database with Oracle Database 23ai.
- Transaction Restrictions for Parallel DML if you are using Autonomous Database with Oracle Database 19c.

Footnote 2

The CPU shares assigned to each service determine how much CPU and IO the sessions using those services will get when CPU utilization in your database reaches 100%. In that case, each service will get CPU and IO resources proportional to its assigned CPU shares. You can modify these shares when needed to adjust them for your workload. The automated tasks that run on Autonomous Database, such as automatic optimizer statistics gathering, run with a resource share of 1. See Manage CPU/IO Shares on Autonomous Database for more information.

Footnote 3

Concurrency in this table means the number of active sessions running SQL statements concurrently. The number of sessions that do not use parallelism, for example sessions in the `LOW` service, is bounded by the `sessions` database parameter. This parameter determines the number of sessions you can have in the database independent of what service those sessions are using. After that limit is reached, new sessions will not be able to connect to the database. For example, if your database has the `sessions` parameter set to 300, you can have different numbers of sessions using `LOW` and `TP`, with the total being limited to 300.

The `HIGH` and `MEDIUM` services allow as many concurrent SQL statements to run as indicated in this table, assuming you are only using one of these services. You can still submit more SQL statements using those services, but subsequent statements will be queued after these numbers are reached. As the running statements finish, queued sessions will be picked from the queue. If you use parallelism in multiple services concurrently, or if you use parallelism hints to set the degree of parallelism for your SQL statements, the concurrency you get may be different than the documented numbers shown in this table.

To see the concurrency levels for your database, run the following command:

```
SELECT * FROM CS_RESOURCE_MANAGER.LIST_CURRENT_RULES();
```

You can also see the currency levels in the Set Resource Management Rules card in Database Actions. See Manage CPU/IO Shares on Autonomous Database for more information.

Footnote 4

Manual parallelism is the type of parallelism where you can control the parallelism for queries, DDL, and DML statements by setting parallelism degrees on your tables and indexes. See Manually Specifying the Degree of Parallelism for more information.

Footnote 5

The `TP` and `TPURGENT` services are not available for Data Warehouse workloads. If you have a workload that requires manual parallelism, which is only available with `TPURGENT`, file a Service Request at Oracle Cloud Support to request to enable the `TP` and `TPURGENT` services in your Autonomous Database instance.

Topics

- Where Do I Find Connection Descriptions for Database Services?
  The available database services are part of your database connect strings and you can view the services from the Oracle Cloud Infrastructure Console or view the database services names in the `tnsnames.ora` in the wallet file you can download.
- Which Database Service Should I Choose for My Connection, Application, or Tool?
  The service you choose for your connection, application, or tool depends on your workload characteristics.
- Database Service Names for Autonomous Data Warehouse
  Describes the available service names for connecting to Autonomous Database instance with a Data Warehouse workload type.
- Database Service Names for Autonomous Transaction Processing and Autonomous JSON Database
  Describes the available service names for connecting to Autonomous Database instance with a Transaction Processing, JSON, or APEX workload type.

Where Do I Find Connection Descriptions for Database Services?

The available database services are part of your database connect strings and you can view the services from the Oracle Cloud Infrastructure Console or view the database services names in the `tnsnames.ora` in the wallet file you can download.

See View TNS Names and Connection Strings for an Autonomous Database Instance for the steps to view the connection strings on the Oracle Cloud Infrastructure Console,

See Download Client Credentials (Wallets) for more information.

For example, for an Autonomous Database instance with the Transaction Processing workload type named sales, the TNS aliases will be `sales_tpurgent`, `sales_tp`, `sales_high`, `sales_medium`, and `sales_low`.

Which Database Service Should I Choose for My Connection, Application, or Tool?

The service you choose for your connection, application, or tool depends on your workload characteristics.

The following general guidelines can help you choose.a database service to use:

- `TP`, `TPURGENT`, or `LOW`: Use the `TP` or the `LOW` service for running OLTP applications or tools that are expected to run short-running queries and transactions. If you have different types of users who need to have different CPU and IO priorities, use `TP`, `TPURGENT`, and `LOW`, depending on their priorities, and adjust the CPU shares for these services based on your requirements.
- `MEDIUM`: Use the `MEDIUM` service for analytic workloads like long-running queries scanning large data sets, batch data load processes, adhoc operational queries, and so on. If the documented degree of parallelism and concurrency level for this service does not fit your

needs, you can set the concurrency level for it so that its degree of parallelism and concurrency is adjusted for your workload.

Use the `MEDIUM` service for running single-user benchmarks that measure the data scan, query and DML performance for analytic workloads and set its concurrency level to 1 to allow every SQL statement to use the maximum degree of parallelism possible.

See Change MEDIUM Service Concurrency Limit (ECPU Compute Model) for more information on changing the concurrency level.

- `HIGH`: Use the `HIGH` service for low-concurrency analytic workloads, such as a reporting database that is accessed by only a few users simultaneously.

If you need to switch to a different service in the same session to run SQL statements in that service, you can use the `CS_SESSION` package to do that. See CS_SESSION Package for more information.

Database Service Names for Autonomous Data Warehouse

Describes the available service names for connecting to Autonomous Database instance with a Data Warehouse workload type.

The service names for connections to an Autonomous Database instance with a Data Warehouse workload type are in the format:

- *databasename*_high
- *databasename*_medium
- *databasename*_low

These services map to the `LOW`, `MEDIUM`, and `HIGH` consumer groups.

For example, if you create an Autonomous Database with a Data Warehouse workload type and specify the database name as `DB2024`, your service names are:

- db2024_high
- db2024_medium
- db2024_low

If you connect using the `db2024_low` service, the connection uses the `LOW` consumer group.

The basic characteristics of these consumer groups are:

- HIGH: Highest resources, lowest concurrency. Queries run in parallel.
- MEDIUM: Less resources, higher concurrency. Queries run in parallel.

  Picking one of the predefined services provides concurrency values that work well for most applications. In cases where selecting one of the default services does not meet your application's performance needs, you can use the MEDIUM service and modify the concurrency limit. For example, when you run single-user benchmarks, you can set the concurrency limit of the MEDIUM service to 1 in order to obtain the highest degree of parallelism (DOP).

  Depending on your compute model, ECPU or OCPU, see the following for more information.

  – Change MEDIUM Service Concurrency Limit (ECPU Compute Model)
  – Change MEDIUM Service Concurrency Limit (OCPU Compute Model)
- LOW: Least resources, highest concurrency. Queries run serially.

> **Note:**
>
> After connecting to the database using one service, do not attempt to manually switch that connection to a different service by simply changing the consumer group of the connection. When you connect using a service, Autonomous Database performs more actions to configure the connection than just setting its consumer group. You can use the procedure `CS_SESSION.SWITCH_SERVICE` to switch to a different service.
> See SWITCH_SERVICE Procedure for more information.

Database Service Names for Autonomous Transaction Processing and Autonomous JSON Database

Describes the available service names for connecting to Autonomous Database instance with a Transaction Processing, JSON, or APEX workload type.

The available service names for connections to an Autonomous Database instance with one of: Transaction Processing, JSON, or APEX workload type are in the format:

- *databasename*_tpurgent
- *databasename*_tp
- *databasename*_high
- *databasename*_medium
- *databasename*_low

These services map to the `TPURGENT`, `TP`, `HIGH`, `MEDIUM` and `LOW` consumer groups.

For example, if you create an Autonomous Database with a Transaction Processing workload type and specify the database name as `DB2024`, your connection service names are:

- db2024_tpurgent
- db2024_tp
- db2024_high
- db2024_medium
- db2024_low

If you connect using the `db2024_tp` service, the connection uses the `TP` consumer group.

The basic characteristics of these consumer groups are:

- TPURGENT: The highest priority application connection service for time critical transaction processing operations. This connection service supports manual parallelism.
- TP: A typical application connection service for transaction processing operations. This connection service does not run with parallelism.
- HIGH: A high priority application connection service for reporting and batch operations. All operations run in parallel and are subject to queuing.
- MEDIUM: A typical application connection service for reporting and batch operations. All operations run in parallel and are subject to queuing.

  Picking one of the predefined services provides concurrency values that work well for most applications. In cases where selecting one of the default services does not meet your

application's performance needs, you can use the MEDIUM service and modify the concurrency limit. For example, when you run single-user benchmarks, you can set the concurrency limit of the MEDIUM service to 1 in order to obtain the highest degree of parallelism (DOP).

Depending on your compute model, ECPU or OCPU, see the following for more information.

- Change MEDIUM Service Concurrency Limit (ECPU Compute Model)
- Change MEDIUM Service Concurrency Limit (OCPU Compute Model)

- LOW: A lowest priority application connection service for reporting or batch processing operations. This connection service does not run with parallelism.

> **Note:**
>
> After connecting to the database using one service, do not attempt to manually switch that connection to a different service by simply changing the consumer group of the connection. When you connect using a service, Autonomous Database performs more actions to configure the connection than just setting its consumer group. You can use the procedure `CS_SESSION.SWITCH_SERVICE` to switch to a different service.
> See SWITCH_SERVICE Procedure for more information.

Idle Time Limits

Autonomous Database has predefined idle time limits for sessions so that idle sessions do not hold system resources for a long time.

A session may be terminated if it stays idle for more than five (5) minutes and the resources it consumes are needed by other users. This allows other active sessions to proceed without waiting for the idle session.

If you want sessions to be terminated after a certain amount of time, independent of the consumed resources needed by other users, then set the `MAX_IDLE_TIME` initialization parameter. The `MAX_IDLE_TIME` parameter specifies the maximum number of minutes that a session can be idle. After the specified amount of time, `MAX_IDLE_TIME` kills sessions.

> **Note:**
>
> Sessions that are idle for more than 24 hours are terminated whether they are holding resources or not.

See Initialization Parameters for more information.

Service Concurrency

The consumer groups of the predefined service names provide different levels of performance and concurrency. The available service names are different depending on your workload: Data Warehouse, Transaction Processing, or JSON Database.

Picking one of the predefined services provides concurrency values that work well for most applications. In cases where selecting one of the default services does not meet your

application's performance needs, you can use the MEDIUM service and modify the concurrency limit. For example, when you run single-user benchmarks, you can set the concurrency limit of the MEDIUM service to 1 in order to obtain the highest degree of parallelism (DOP).

In this topic, the "number of ECPUs" is the **ECPU count** shown in the Oracle Cloud Infrastructure Console. Likewise, if your database uses the OCPU compute model, the "number of OCPUs" is the **OCPU count**.

> **Note:**
>
> OCPU is a legacy billing metric and has been retired on Autonomous Database. ECPUs are the replacement billing metric for all new and existing Autonomous Database deployments. See Oracle Support Document 2998742.1 for more information.

- Service Concurrency Limits for Data Warehouse Workloads (ECPU Compute Model)
  The `tnsnames.ora` file provided with the credentials zip file contains three database service names identifiable as high, medium, and low for Autonomous Database with Data Warehouse workloads.
- Service Concurrency Limits for Transaction Processing Workloads (ECPU Compute Model)
  The `tnsnames.ora` file provided with the credentials zip file contains five database service names identifiable as tpurgent, tp, high, medium, and low for Autonomous Database with Transaction Processing or JSON workloads.
- Service Concurrency Limits for Data Warehouse Workloads (OCPU Compute Model)
  The `tnsnames.ora` file provided with the credentials zip file contains three database service names identifiable as high, medium and low for Autonomous Database with Data Warehouse workloads.
- Service Concurrency Limits for Transaction Processing and JSON Database Workloads (OCPU Compute Model)
  The `tnsnames.ora` file provided with the credentials zip file contains five database service names identifiable as tpurgent, tp, high, medium, and low for Autonomous Database with Transaction Processing or with JSON Database workloads.

Service Concurrency Limits for Data Warehouse Workloads (ECPU Compute Model)

The `tnsnames.ora` file provided with the credentials zip file contains three database service names identifiable as high, medium, and low for Autonomous Database with Data Warehouse workloads.

The following shows the details for the number of concurrent statements for each connection service for Data Warehouse workloads, with **Compute auto scaling** disabled, and with **Compute auto scaling** enabled.

> **Note:**
>
> The values in this table apply when the number of ECPUs is equal to or greater than 4.
> When the number of ECPUs is 2, all services use the concurrency limit 150. When the number of ECPUs is 3, all services use the concurrency limit 225 (this applies for **Compute auto scaling** enabled or disabled).

Database Service Name	Number of Concurrent Queries with Compute Auto Scaling Disabled	Number of Concurrent Queries with Compute Auto Scaling Enabled
`high`	3	9
`medium`	0.25125 × number of ECPUs A decimal result is truncated.	0.75375 × number of ECPUs A decimal result is truncated.
`low`	75 × number of ECPUs	75 × number of ECPUs

> **Note:**
>
> By default the HIGH consumer group can run at least 1 SQL statement when the MEDIUM consumer group is also running statements. If you change the concurrency for the MEDIUM service, this might not apply.

The following table shows sample concurrent connections values for a database with 32 ECPUs with **Compute auto scaling** disabled and with **Compute auto scaling** enabled.

Database Service Name	Number of Concurrent Queries with Compute Auto Scaling Disabled	Number of Concurrent Queries with Compute Auto Scaling Enabled
`high`	3	9
`medium`	8	24
`low`	Up to 2400	Up to 2400

See Database Service Names for Autonomous Database and Change MEDIUM Service Concurrency Limit (ECPU Compute Model) more information.

Service Concurrency Limits for Transaction Processing Workloads (ECPU Compute Model)

The `tnsnames.ora` file provided with the credentials zip file contains five database service names identifiable as tpurgent, tp, high, medium, and low for Autonomous Database with Transaction Processing or JSON workloads.

The following shows the details for the default number of concurrent statements for each connection service for Transaction Processing or JSON workloads.

> **Note:**
>
> The values in this table apply when the number of ECPUs is equal to or greater than 4.
> When the number of ECPUs is 2, all services use the concurrency limit 150. When the number of ECPUs is 3, all services use the concurrency limit 225 (this applies for **Compute auto scaling** enabled or disabled).

Database Service Name	Concurrent Statements with Compute Auto Scaling Disabled	Concurrent Statements with Compute Auto Scaling Enabled
`tpurgent`	75 × number of ECPUs	75 × number of ECPUs
`tp`	75 × number of ECPUs	75 × number of ECPUs
`high`	3	9
`medium`	0.25125 × number of ECPUs A decimal result is truncated.	0.75375 × number of ECPUs A decimal result is truncated.
`low`	75 × number of ECPUs	75 × number of ECPUs

See Database Service Names for Autonomous Database and Change MEDIUM Service Concurrency Limit (ECPU Compute Model) more information.

Service Concurrency Limits for Data Warehouse Workloads (OCPU Compute Model)

The `tnsnames.ora` file provided with the credentials zip file contains three database service names identifiable as high, medium and low for Autonomous Database with Data Warehouse workloads.

The following shows the details for the number of concurrent statements for each connection service for Data Warehouse workloads, with OCPU auto scaling disabled, and with OCPU auto scaling enabled.

> **Note:**
>
> The values in this table apply when the number of OCPUs is greater than 1. With 1 OCPU, the concurrency limit for each service is 300 and the DOP is 1.

OCPU is a legacy billing metric and has been retired on Autonomous Database. ECPUs are the replacement billing metric for all new and existing Autonomous Database deployments. See Oracle Support Document 2998742.1 for more information.

Database Service Name	Concurrent Statements with OCPU Auto Scaling Disabled	Concurrent Statements with OCPU Auto Scaling Enabled
`high`	3	9
`medium`	1.26 × number of OCPUs	3.78 × number of OCPUs
`low`	300 × number of OCPUs	300 × number of OCPUs

The following table shows sample concurrent connections values for a database with 16 OCPUs, showing the values with both OCPU auto scaling enabled and disabled.

Database Service Name	Number of Concurrent Queries with OCPU Auto Scaling Disabled	Number of Concurrent Queries with OCPU Auto Scaling Enabled
high	3	9
medium	20	60
low	Up to 4800	Up to 4800

See Database Service Names for Autonomous Database and Change MEDIUM Service Concurrency Limit (ECPU Compute Model) more information.

Service Concurrency Limits for Transaction Processing and JSON Database Workloads (OCPU Compute Model)

The tnsnames.ora file provided with the credentials zip file contains five database service names identifiable as tpurgent, tp, high, medium, and low for Autonomous Database with Transaction Processing or with JSON Database workloads.

The following shows the details for the default number of concurrent statements for each connection service for Transaction Processing or JSON Database workloads, with OCPU auto scaling disabled, and with OCPU auto scaling enabled.

> **Note:**
>
> The values in this table apply when the number of OCPUs is greater than 1. With 1 OCPU, the concurrency limit for each service is 300 and the DOP is 1.

OCPU is a legacy billing metric and has been retired on Autonomous Database. ECPUs are the replacement billing metric for all new and existing Autonomous Database deployments. See Oracle Support Document 2998742.1 for more information.

Database Service Name	Concurrent Statements with OCPU Auto Scaling Disabled	Concurrent Statements with OCPU Auto Scaling Enabled
tpurgent	300 × number of OCPUs	300 × number of OCPUs
tp	300 × number of OCPUs	300 × number of OCPUs
high	3	9
medium	1.26 × number of OCPUs	3.78 × number of OCPUs
low	300 × number of OCPUs	300 × number of OCPUs

See Database Service Names for Autonomous Database and Change MEDIUM Service Concurrency Limit (ECPU Compute Model) more information.

Change MEDIUM Service Concurrency Limit (ECPU Compute Model)

If your application requires customized concurrency, you can modify the concurrency limit for your Autonomous Database MEDIUM service.

Picking one of the predefined services provides concurrency values that work well for most applications. In cases where selecting one of the default services does not meet your application's performance needs, you can use the MEDIUM service and modify the

concurrency limit. For example, when you run single-user benchmarks, you can set the concurrency limit of the MEDIUM service to 1 in order to obtain the highest degree of parallelism (DOP).

> **Note:**
>
> Changing the concurrency limit is only allowed for an instance that has four (4) or more ECPUs.

For example, with **Compute auto scaling** disabled, if your instance is configured with 400 ECPUs, by default Autonomous Database provides a concurrency limit of 100 for the MEDIUM service:

`0.25125 x number of ECPUs` sessions (up to 100 concurrent queries). A decimal result is truncated.

In this example the MEDIUM service supports an application with up to 100 concurrent queries with DOP of 4. If you only need 50 concurrent queries and you want a higher DOP you can decrease the concurrency limit and the database increases the DOP. To do this, set the MEDIUM service concurrency limit to 50. When you change the concurrency limit the database calculates and sets the DOP based on the concurrency limit you select and the number of ECPUs. For this example, with the concurrency limit set to 50, the new DOP is 12.

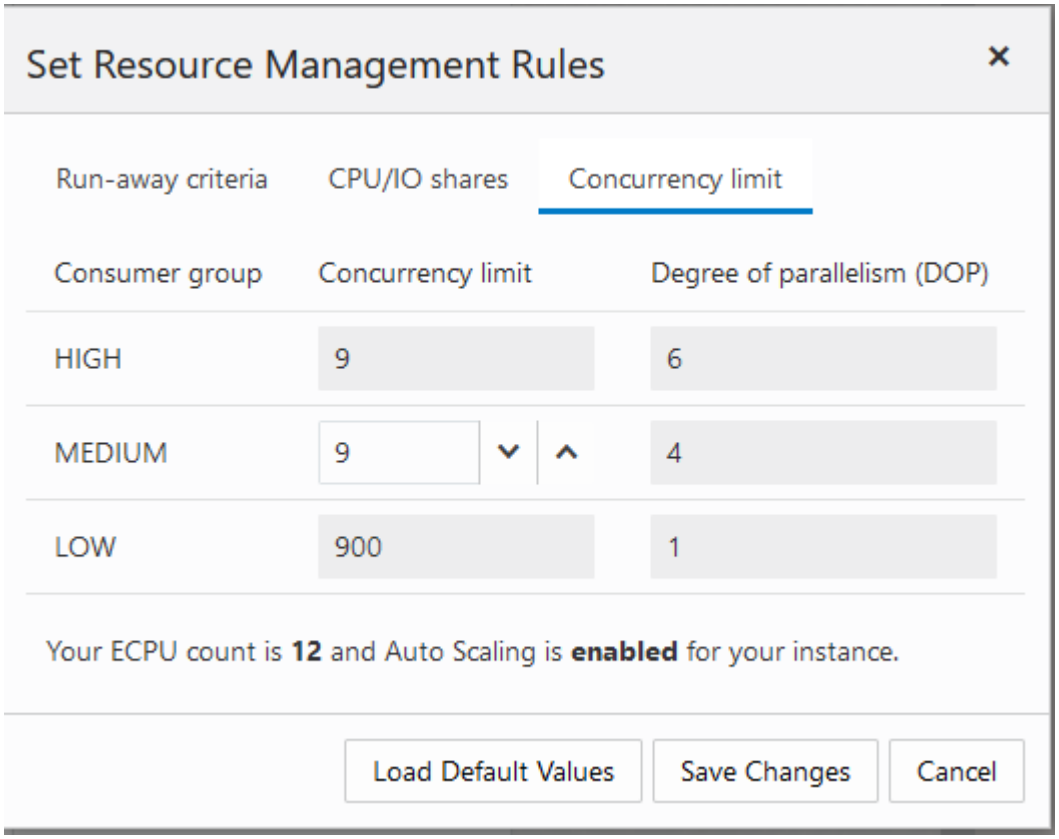
With **Compute auto scaling** enabled, the DOP is set to a value three times greater. In this example the DOP value would be 36.

You can change the concurrency limit for the MEDIUM service in Database Actions or using the PL/SQL package `CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE`.

Follow these steps to change the MEDIUM service concurrency limit in Database Actions:

1. Access Database Actions as the ADMIN user.

   See Access Database Actions as ADMIN for more information.
2. On the Database Actions Launchpad, under **Administration**, click **Set Resource Management Rules**.
3. On the Set Resource Management Rules page, select the **Concurrency limit** tab.



4. For the MEDIUM service, change the value to the desired concurrency limit by entering a value or by clicking the Decrement or Increment icons.

   If the concurrency limit you specify is not valid, based on the number of ECPUs, you will receive a message such as the following, listing the valid range of values for your instance:

   ```
   Enter a number between 1 and 12.
   ```

5. Click **Save Changes**.
6. Click **OK**.

To reset the concurrency limit for the MEDIUM service to its default value, click **Load Default Values** and click **Save Changes**.

- Change MEDIUM Service Concurrency Limit with PL/SQL Procedure UPDATE_PLAN_DIRECTIVE (ECPU Compute Model)
  As an alternative to using the **Set Resource Management Rules** card in Database Actions, you can use PL/SQL to change the concurrency limit for the MEDIUM service.
- Change MEDIUM Service Concurrency Limit Notes (ECPU Compute Model)

Related Topics

- Change MEDIUM Service Concurrency Limit (OCPU Compute Model)
  If your application requires customized concurrency, you can modify the concurrency limit for your Autonomous Database MEDIUM service.

Change MEDIUM Service Concurrency Limit with PL/SQL Procedure UPDATE_PLAN_DIRECTIVE (ECPU Compute Model)

As an alternative to using the **Set Resource Management Rules** card in Database Actions, you can use PL/SQL to change the concurrency limit for the MEDIUM service.

To change the MEDIUM service concurrency limit with `CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE`:

1. Call the PL/SQL procedure `CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE` to update the concurrency limit for the MEDIUM consumer group.

   For example, with 12 ECPUs, change the MEDIUM service's concurrency limit to 2, as follows:

   ```
   BEGIN
       CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(consumer_group => 'MEDIUM',
   concurrency_limit => 2);
   END;
   /
   ```

   If the `concurrency_limit` you specify is not valid, based on the number of ECPUs, the procedure shows an error message, listing the valid range of values for your instance. For example with 12 ECPUs:

   ```
   ORA-20000: Invalid or missing value. Concurrency limit must be between 1
   and 9 for the specified CPU count
   ```

   This error message example is from an instance with 12 ECPUs.

2. Use the PL/SQL function `CS_RESOURCE_MANAGER.LIST_CURRENT_RULES` to verify the updated MEDIUM service concurrency limit and degree of parallelism:

   ```
   SELECT * FROM CS_RESOURCE_MANAGER.LIST_CURRENT_RULES();
   ```

   This procedure returns the list of values for all consumer groups. After you modify the concurrency limit as specified in Step 1, check the MEDIUM service `CONCURRENCY_LIMIT` and `DEGREE_OF_PARALLELISM` values to verify your changes.

3. After you change the concurrency limit for the MEDIUM service, test your application by connecting with the MEDIUM service to verify that the customized concurrency limit meets your performance objectives.

When you want to go back to the default values, use the `CS_RESOURCE_MANAGER.REVERT_TO_DEFAULT_VALUES` PL/SQL procedure to revert to the default settings for the MEDIUM service.

For example:

```
BEGIN
    CS_RESOURCE_MANAGER.REVERT_TO_DEFAULT_VALUES(consumer_group => 'MEDIUM',
concurrency_limit => TRUE);
END;
/
```

See CS_RESOURCE_MANAGER Package for more information.

Change MEDIUM Service Concurrency Limit Notes (ECPU Compute Model)

- Changing the concurrency limit is only allowed for the MEDIUM service.
- Changing the concurrency limit is only allowed when the number of ECPUs is greater than or equal to 4.
- Changing the concurrency limit also changes the degree of parallelism (in some cases the value does not change, depending on the magnitude of the difference between the old concurrency limit and the new value you set).
- The concurrency limit you set must be in the range:
  - With **Compute auto scaling** disabled: between 1 and `.75 x the number of ECPUs`
  - With **Compute auto scaling** enabled: between 1 and `2.25 x the number of ECPUs`
- The MEDIUM service sets the following concurrency limit and DOP values by default:

MEDIUM Database Service	Default Value with Compute Auto Scaling Disabled	Default Value with Compute Auto Scaling Enabled
Concurrency Limit	`0.25125 × number of ECPUs` when the number of ECPUs 8 A decimal result is truncated 2 when the number of ECPUs is in the range `4 ≤ ECPUs < 8`	`0.75375 × number of ECPUs` when the number of ECPUs 8 A decimal result is truncated 6 when the number of ECPUs is in the range: `4 ≤ ECPUs < 8`
DOP	4 when the number of ECPUs 8 or `TRUNC (ECPU/2)`, when the number of ECPUs < 8	4 when the number of ECPUs 8 or `TRUNC (ECPU/2)`, when the number of ECPUs < 8

- By changing the value of the concurrency limit, the DOP of the MEDIUM service can go as low as 2 and as high as `.75 * number of ECPUs` (if **Compute auto scaling** is disabled) or `2.25 x number of ECPUs` (if **Compute auto scaling** is enabled).

  See Use Auto Scaling for information on **Compute auto scaling**.
- At any time you can return to the default values for the MEDIUM service concurrency limit and DOP.

Change MEDIUM Service Concurrency Limit (OCPU Compute Model)

If your application requires customized concurrency, you can modify the concurrency limit for your Autonomous Database MEDIUM service.

Picking one of the predefined services provides concurrency values that work well for most applications. In cases where selecting one of the default services does not meet your application's performance needs, you can use the MEDIUM service and modify the concurrency limit. For example, when you run single-user benchmarks, you can set the concurrency limit of the MEDIUM service to 1 in order to obtain the highest degree of parallelism (DOP).

> **Note:**
>
> Changing the concurrency limit is only allowed for an instance that has two (2) or more OCPUs.

For example, if your instance is configured with 100 OCPUs, by default Autonomous Database provides a concurrency limit of 126 for the MEDIUM service:

`1.26 x number of OCPUs` sessions (up to 126 concurrent queries)

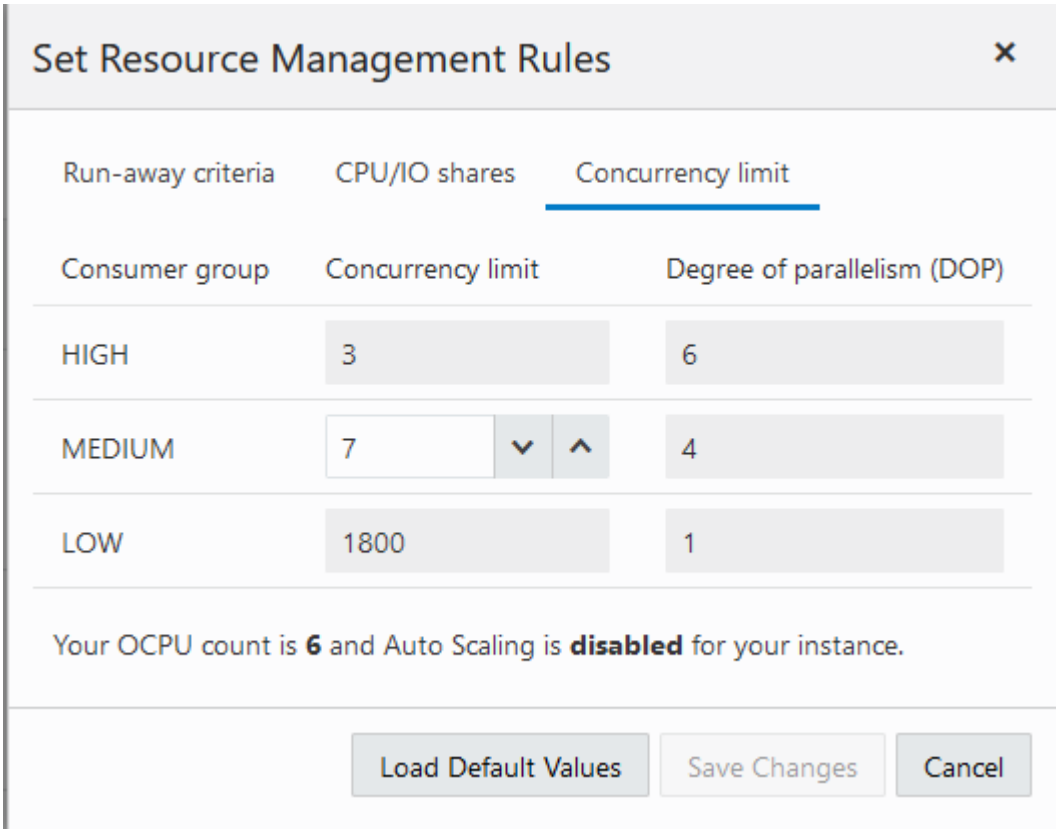
In this example using the MEDIUM service supports an application with up to 126 concurrent queries with DOP of 4. If you only need 50 concurrent queries and you want a higher DOP you can decrease the concurrency limit and thus increase the DOP. To do this, set the MEDIUM service concurrency limit to 50. When you change the concurrency limit the system calculates and sets the DOP based on the concurrency limit you select and the number of OCPUs. For this example, with the concurrency limit set to 50, the new DOP is 12. When **OCPU auto scaling** is enabled, the DOP is set to a value three times greater. In this example the DOP value would be 36.

You can change the concurrency limit for the MEDIUM service in Database Actions or using the PL/SQL package `CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE`.

Follow these steps to change the MEDIUM service concurrency limit in Database Actions:

1. Access Database Actions as the ADMIN user.

   See Access Database Actions as ADMIN for more information.
2. On the Database Actions Launchpad, under **Administration**, click **Set Resource Management Rules**.
3. On the Set Resource Management Rules page, select the **Concurrency limit** tab.



4. For the MEDIUM service, change the value to the desired concurrency limit by entering a value or by clicking the Decrement or Increment icons.

   If the concurrency limit you specify is not valid, based on the number of OCPUs, you will receive a message such as the following, listing the valid range of values for your instance:

   ```
   Please enter a concurrency limit between 1 and 300
   ```

   This error message example is from an instance with 100 OCPUs (the 300 maximum value shown is 3 x number of OCPUs).

5. Click **Save Changes**.
6. Click **OK**.

To reset the concurrency limit for the MEDIUM service to its default value, click **Load Default Values** and click **Save Changes**.

- Change MEDIUM Service Concurrency Limit with PL/SQL Procedure UPDATE_PLAN_DIRECTIVE (OCPU Compute Model)
  As an alternative to using the **Set Resource Management Rules** card in Database Actions, you can use PL/SQL to change the concurrency limit for the MEDIUM service.
- Change MEDIUM Service Concurrency Limit Notes (OCPU Compute Model)

Related Topics

- Change MEDIUM Service Concurrency Limit (ECPU Compute Model)
  If your application requires customized concurrency, you can modify the concurrency limit for your Autonomous Database MEDIUM service.

Change MEDIUM Service Concurrency Limit with PL/SQL Procedure UPDATE_PLAN_DIRECTIVE (OCPU Compute Model)

As an alternative to using the **Set Resource Management Rules** card in Database Actions, you can use PL/SQL to change the concurrency limit for the MEDIUM service.

To change the MEDIUM service concurrency limit with `CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE`:

1. Call the PL/SQL procedure `CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE` to update the concurrency limit for the MEDIUM consumer group.

   For example, with 3 OCPUs, change the MEDIUM service's concurrency limit to 2, as follows:

   ```
   BEGIN
       CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(consumer_group => 'MEDIUM',
   concurrency_limit => 2);
   END;
   /
   ```

   If the `concurrency_limit` you specify is not valid, based on the number of OCPUs, you will receive a message such as the following, listing the valid range of values for your instance:

   ```
   ORA-20000: Invalid or missing value. Concurrency limit must be between 1
   and 9 for the specified OCPU count
   ```

   This error message example is from an instance with 3 OCPUs.

2. Use the PL/SQL function `CS_RESOURCE_MANAGER.LIST_CURRENT_RULES` to verify the updated MEDIUM service concurrency limit and degree of parallelism:

   ```
   SELECT * FROM CS_RESOURCE_MANAGER.LIST_CURRENT_RULES();

   CONSUMER_GROUP ELAPSED_TIME_LIMIT IO_MEGABYTES_LIMIT SHARES
   CONCURRENCY_LIMIT DEGREE_OF_PARALLELISM
   -------------- ------------------ ------------------ ------
   ----------------- ---------------------
   HIGH
   4                 3                      3
   MEDIUM
   2                 2                      9
   LOW                                                        1
   900                     1
   ```

   This procedure returns the list of values for all consumer groups. After you modify the concurrency limit as specified in Step 1, check the MEDIUM service `CONCURRENCY_LIMIT` and `DEGREE_OF_PARALLELISM` values to verify your changes.

3. After you change the concurrency limit for the MEDIUM service, test your application by connecting with the MEDIUM service to verify that the customized concurrency limit meets your performance objectives.

When you want to go back to the default values, use the `CS_RESOURCE_MANAGER.REVERT_TO_DEFAULT_VALUES` PL/SQL procedure to revert to the default settings for the MEDIUM service.

For example:

```
BEGIN
    CS_RESOURCE_MANAGER.REVERT_TO_DEFAULT_VALUES(consumer_group => 'MEDIUM',
concurrency_limit => TRUE);
END;
/
```

See CS_RESOURCE_MANAGER Package for more information.

Change MEDIUM Service Concurrency Limit Notes (OCPU Compute Model)

- Changing the concurrency limit is only allowed for the MEDIUM service.
- Changing the concurrency limit is only allowed when the number of OCPUs is greater than 1.
- Changing the concurrency limit also changes the degree of parallelism (in most cases, depending on the magnitude of the difference between the old concurrency limit and the new value you set).
- The concurrency limit you set must be in the range:
  - With **OCPU auto scaling** disabled: between: `1` and `3 x the number of OCPUs`
  - With **OCPU auto scaling** enabled: between `1` and `9 x the number of OCPUs`
- The MEDIUM service sets the following concurrency limit and DOP values by default:

MEDIUM Database Service	Default Value with OCPU Auto Scaling Disabled	Default Value with OCPU Auto Scaling Enabled
Concurrency Limit	1.26 × number of OCPUs when the number of OCPUs 4 5 when the number of OCPUs < 4	3.78 × number of OCPUs when the number of OCPUs 4 15 when the number of OCPUs < 4
DOP	4 when the number of OCPUs 4 or The number of OCPUs, when the number of OCPUs < 4	4 when the number of OCPUs 4 or The number of OCPUs, when the number of OCPUs < 4

- By changing the value of the concurrency limit, the DOP of the MEDIUM service can go as low as 2 and as high as: `2 x number of OCPUs` (if compute auto scaling is disabled) or `6 x number of OCPUs` (if compute auto scaling is enabled).

  See Use Auto Scaling for information on compute auto scaling.
- At any time you can return to the default values for the MEDIUM service concurrency limit and DOP.

Predefined Job Classes with Oracle Scheduler

Autonomous Database includes predefined `job_class` values to use with Oracle Scheduler.

The predefined `job_class` values, `TPURGENT`, `TP`, `HIGH`, `MEDIUM` and `LOW` map to the corresponding consumer groups. These job classes allow you to specify the consumer group a job runs in with `DBMS_SCHEDULER.CREATE_JOB`.

The `DBMS_SCHEDULER.CREATE_JOB` procedure supports `PLSQL_BLOCK` and `STORED_PROCEDURE` job types for the `job_type` parameter in Autonomous Database.

For example: use the following to create a single regular job to run in `HIGH` consumer group:

```
BEGIN
   DBMS_SCHEDULER.CREATE_JOB (
     job_name => 'update_sales',
     job_type => 'STORED_PROCEDURE',
     job_action => 'OPS.SALES_PKG.UPDATE_SALES_SUMMARY',
     start_date => '28-APR-19 07.00.00 PM Australia/Sydney',
     repeat_interval => 'FREQ=DAILY;INTERVAL=2',
     end_date => '20-NOV-19 07.00.00 PM Australia/Sydney',
     auto_drop => FALSE,
     job_class => 'HIGH',
     comments => 'My new job');
END;
/
```

Notes for Oracle Scheduler:

- To use `DBMS_SCHEDULER.CREATE_JOB` additional grants for specific roles or privileges might be required. The `ADMIN` user and users with `DWROLE` have the required `CREATE SESSION` and `CREATE JOB` privileges. If a user does not have `DWROLE` then grants are required for `CREATE SESSION` and `CREATE JOB` privileges.
- The `instance_id` job attribute is ignored for Oracle Scheduler jobs running on Autonomous Database.

See Scheduling Jobs with Oracle Scheduler for more information on Oracle Scheduler and `DBMS_SCHEDULER.CREATE_JOB`.

See SET_ATTRIBUTE Procedure for information on job attributes.

Manage CPU/IO Shares on Autonomous Database

Autonomous Database comes with predefined CPU/IO shares assigned to different consumer groups. You can modify these predefined CPU/IO shares if your workload requires different CPU/IO resource allocations.

The CPU/IO shares assigned to the consumer groups determine the CPU/IO resources a consumer group can use with respect to the other consumer groups. The default CPU/IO shares depend on the Autonomous Database workload.

The CPU shares assigned to each service determine how much CPU and IO the sessions using those services will get when CPU utilization in your database reaches 100%. In that case, each service will get CPU and IO resources proportional to its assigned CPU shares. The Oracle internal processes that run as automated tasks on Autonomous Database, such as automatic optimizer statistics gathering, run with a resource share of 1.

Workload Type	Details
Data Warehouse	By default, the CPU/IO shares assigned to the consumer groups HIGH, MEDIUM, LOW are 4, 2, and 1, respectively. With the default settings the consumer group HIGH will be able to use 4 times more CPU/IO resources compared to LOW and 2 times more CPU/IO resources compared to MEDIUM, when needed. The consumer group MEDIUM will be able to use 2 times more CPU/IO resources compared to LOW, when needed.
Transaction Processing JSON Database	By default, the CPU/IO shares assigned to the consumer groups TPURGENT, TP, HIGH, MEDIUM, and LOW are 12, 8, 4, 2, and 1, respectively. With the default settings the consumer group TPURGENT will be able to use 12 times more CPU/IO resources compared to LOW, when needed. The consumer group TP will be able to use 4 times more CPU/IO resources compared to MEDIUM, when needed.

You can set CPU/IO shares in Database Actions or using the PL/SQL package `CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE`.

To use Database Actions to change the CPU/IO share values for consumer groups:

1. Access Database Actions as the ADMIN user.

   See Access Database Actions as ADMIN for more information.
2. On the Database Actions Launchpad, under **Administration**, click **Set Resource Management Rules**.
3. Select the **CPU/IO shares** tab to set CPU/IO share values for consumer groups.
4. Set the desired CPU/IO share value for a consumer group by entering a value or by clicking the Decrement or Increment icons.
5. Click **Save Changes**.
6. Click **OK**.

To reset CPU/IO shares values to the defaults, click **Load Default Values** and click **Save Changes** to apply the populated values.

As an alternative to using Database Actions, you can use the PL/SQL procedure `CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE` to change the CPU/IO share values for consumer groups.

For example, on an Autonomous Data Warehouse database, run the following script as the `ADMIN` user to set the CPU/IO shares to 8, 2, and 1 for consumer groups HIGH, MEDIUM, and LOW respectively. This allows the consumer group HIGH to use 4 times more CPU/IO resources compared to the consumer group MEDIUM and 8 times CPU/IO resources compared to the consumer group LOW:

```
BEGIN
   CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(consumer_group => 'HIGH', shares => 8);
   CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(consumer_group => 'MEDIUM', shares => 2);
   CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(consumer_group => 'LOW', shares => 1);
END;
/
```

For example, on an Autonomous JSON Database or on an Autonomous Transaction Processing database, run the following script as the ADMIN user to set CPU/IO shares to 12, 4, 2, 1, and 1 for the consumer groups TPURGENT, TP, HIGH, MEDIUM, and LOW respectively. This allows the consumer group TPURGENT to use 3 times more CPU/IO

resources compared to the consumer group TP and 12 times CPU/IO resources compared to the consumer group MEDIUM:

```
BEGIN
   CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(consumer_group => 'TPURGENT', shares => 12);
   CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(consumer_group => 'TP', shares => 4);
   CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(consumer_group => 'HIGH', shares => 2);
   CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(consumer_group => 'MEDIUM', shares => 1);
   CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(consumer_group => 'LOW', shares => 1);
END;
/
```

When you want to go back to the default shares values you can use the PL/SQL procedure `CS_RESOURCE_MANAGER.REVERT_TO_DEFAULT_VALUES` to revert to the default settings.

For example, on an Autonomous Data Warehouse database, run the following script as the `ADMIN` user to set the CPU/IO shares to default values for the consumer groups HIGH, MEDIUM, and LOW:

```
BEGIN
   CS_RESOURCE_MANAGER.REVERT_TO_DEFAULT_VALUES(consumer_group => 'HIGH', shares => TRUE);
   CS_RESOURCE_MANAGER.REVERT_TO_DEFAULT_VALUES(consumer_group => 'MEDIUM', shares => TRUE);
   CS_RESOURCE_MANAGER.REVERT_TO_DEFAULT_VALUES(consumer_group => 'LOW', shares => TRUE);
END;
/
```

For example, on an or on an Autonomous JSON Database or on an Autonomous Transaction Processing database, run the following script as the ADMIN user to set the default values for CPU/IO shares for the consumer groups TPURGENT, TP, HIGH, MEDIUM, and LOW:

```
BEGIN
   CS_RESOURCE_MANAGER.REVERT_TO_DEFAULT_VALUES(consumer_group => 'TPURGENT', shares => TRUE);
   CS_RESOURCE_MANAGER.REVERT_TO_DEFAULT_VALUES(consumer_group => 'TP', shares => TRUE);
   CS_RESOURCE_MANAGER.REVERT_TO_DEFAULT_VALUES(consumer_group => 'HIGH', shares => TRUE);
   CS_RESOURCE_MANAGER.REVERT_TO_DEFAULT_VALUES(consumer_group => 'MEDIUM', shares => TRUE);
   CS_RESOURCE_MANAGER.REVERT_TO_DEFAULT_VALUES(consumer_group => 'LOW', shares => TRUE);
END;
/
```

See CS_RESOURCE_MANAGER Package for more information.

Manage Runaway SQL Statements on Autonomous Database

Specifies how you configure Autonomous Database to terminate SQL statements automatically based on their query runtime or their IO usage.

You can set runtime run-away rules for query run time and IO usage in Database Actions or using the PL/SQL package `CS_RESOURCE_MANAGER`.

Follow these steps to use Database Actions to set runtime usage rules:

1. Access Database Actions as the ADMIN user.

   See Access Database Actions as ADMIN for more information.

2. On the Database Actions Launchpad, under **Administration**, click **Set Resource Management Rules**.
3. Select the **Run-away criteria** tab to set usage rules for a consumer group.
4. Select the **Consumer group**.
5. Set runaway criteria values:
   - **Query run time (seconds)**
   - **Amount of IO (MB)**
6. Click **Save Changes**.
7. Click **OK**.

When a SQL statement in the specified consumer group runs more than the specified runtime limit or does more IO than the specified amount, then the SQL statement will be terminated.

Click **Load Default Values** to load the default values; then click **Save Changes** to apply the populated values.

You can also use the procedure `CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE` to set these rules. For example, to set a runtime limit of 120 seconds and an IO limit of 1000MB for the HIGH consumer group run the following command when connected to the database as the ADMIN user:

```
BEGIN
     CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(consumer_group => 'HIGH',
io_megabytes_limit => 1000, elapsed_time_limit => 120);
END;
/
```

To reset the values and lift the limits, you can set the values to null:

```
BEGIN
     CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(consumer_group => 'HIGH',
io_megabytes_limit => null, elapsed_time_limit => null);
END;
/
```

See CS_RESOURCE_MANAGER Package for more information.

Use Fast Ingest on Autonomous Database

Fast ingest optimizes the processing of high-frequency, single-row data inserts into the database from applications, such as Internet of Things (IoT) applications, to improve data insert performance.

The intent of fast-ingest is to support applications that generate lots of informational data that has important value in the aggregate but that doesn't necessarily require full ACID guarantees. Many applications in the Internet of Things (IoT) have a rapid "fire and forget" type workload, such as sensor data, smart meter data or even traffic cameras. For these applications, data might be collected and written to the database in high volumes for later analysis.

Fast ingest is very different from normal Oracle Database transaction processing where data is logged and never lost once "written" to the database (that is, committed). In order to achieve the maximum ingest throughput, the normal Oracle transaction mechanisms are bypassed, and it is the responsibility of the application to check to see that all data was indeed written to

the database. Special APIs have been added that can be called to check if the data has been written to the database.

For information on fast ingest and the steps involved in using this feature, refer to Using Fast Ingest in *Database Performance Tuning Guide*.

In addition, the following are required to use Fast Ingest on Autonomous Database:

- **Enable the Optimizer to Use Hints**:
  To use fast ingest with Autonomous Database, you must enable the optimizer to use hints by setting the `optimizer_ignore_hints` parameter to `FALSE` at the session or system level, as appropriate.

  Depending on your Autonomous Database workload type, by default `optimizer_ignore_hints` may be set to `FALSE` at the system level. See Manage Optimizer Statistics on Autonomous Database for more information.
- **Create a Table for Fast Ingest**:
  Prerequisites for Fast Ingest Table contains the limitations for tables to be eligible for Fast Ingest (tables with the specified characteristics cannot use fast ingest).

Enable Database Management Service to Monitor and Manage Autonomous Database

You can use the Diagnostics & Management features in the Database Management service for advanced monitoring and management of Autonomous Databases.

> **Note:**
>
> Certain metrics and Performance Hub features are available free of cost for Autonomous Databases, however, you must create a private endpoint to enable communication and enable Database Management to use the advanced monitoring and management features, which are available at an additional cost.

Prerequisite Steps to Use Database Management Diagnostics & Management for Autonomous Databases:

Before you enable Database Management Diagnostics & Management for an Autonomous Database:

- Complete the general Oracle Cloud Infrastructure prerequisite tasks listed in Perform General Oracle Cloud Infrastructure Prerequisite Tasks.
- Obtain the permissions required to enable Diagnostics & Management for Autonomous Databases, as detailed in Permissions Required to Enable Diagnostics & Management for Autonomous Databases.
- Complete the Autonomous Database-specific prerequisite tasks listed in Autonomous Database-related Prerequisite Tasks.

Enable Database Management Diagnostics & Management for an Autonomous Database Instance:

To enable Database Management Diagnostics & Management for an Autonomous Database instance:

1. On the Autonomous Database details page, under Associated Services, in the **Database Management** field, click **Enable**.

2. On the **Enable Database Management** panel, enter the required information and click **Enable Database Management**. See Enable Diagnostics & Management for Autonomous Databases for more information.

After you enable Diagnostics & Management features in Database Management, you can use the full suite of available capabilities to monitor and manage a single Autonomous Database or a fleet of Autonomous Databases.

Here are some of the important performance monitoring and administration-related features:

- Database fleet overview to monitor multiple databases in one location
- Performance Hub for real-time monitoring of the database workload
- AWR Explorer to visualize AWR data over time
- SQL Monitoring to view top and current running SQL statements with detailed execution statistics
- SQL performance tuning to optimize database performance
- AWR and ADDM reports for in-depth performance analysis
- User administration to track configuration and details of database users
- Parameter configuration to update and monitor database configuration parameters
- Database jobs to streamline collections and run at scale utilizing Database Groups

See the following for more information on these features:

- Monitor the Health of Your Oracle Database Fleet
- Monitor and Manage a Specific Managed Database
- Diagnostics & Management Feature Support Matrix for Oracle Databases

Perform SQL Tracing on Autonomous Database

Use SQL tracing to help you identify the source of an excessive database workload, such as a high load SQL statement in your application.

- Configure SQL Tracing on Autonomous Database
  Shows the steps to configure SQL tracing on Autonomous Database.
- Enable SQL Tracing on Autonomous Database
  Shows the steps to enable SQL tracing for the database session.
- Disable SQL Tracing on Autonomous Database
  Shows the steps to disable SQL tracing on Autonomous Database.
- View Trace File Saved to Cloud Object Store on Autonomous Database
  Describes the output file naming for SQL trace files and shows the commands to use `TKPROF` to organize and view trace file data.
- View Trace Data in SESSION_CLOUD_TRACE View on Autonomous Database
  When you enable SQL Tracing, the same trace information that is saved to the trace file on Cloud Object Store is available in the `SESSION_CLOUD_TRACE` view in the session where the tracing was enabled.

Configure SQL Tracing on Autonomous Database

Shows the steps to configure SQL tracing on Autonomous Database.

> **Note:**
>
> If you enable SQL Tracing your application performance for the session may be degraded while the trace collection is enabled. This negative performance impact is expected due to the overhead of collecting and saving trace data.

To configure your database for SQL tracing, do the following:

1. Create a bucket to store trace files in your Cloud Object Storage.

   To save the SQL tracing files, the bucket can be in any Cloud Object Store that Autonomous Database supports.

   For example, to create a bucket in Oracle Cloud Infrastructure Object Storage, do the following

   a. Open the Oracle Cloud Infrastructure Console.

   b. Select **Storage** from the menu.

   c. Under Storage, select **Object Storage and Archive Storage**.

   d. Click **Create Bucket**.

   e. In the Create Bucket page, enter the **Bucket Name** and click **Create**.

   If you are using an Oracle Cloud Infrastructure Object Storage, note that SQL tracing files are only supported with buckets created in the standard storage tier, make sure you pick **Standard** as the storage tier when creating your bucket. See Overview of Object Storage for information on the Standard Object Storage Tier.

2. Create a credential for your Cloud Object Storage account using `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'DEF_CRED_NAME',
       username => 'adb_user@example.com',
       password => 'password'
   );
   END;
   /
   ```

   For Oracle Cloud Infrastructure the `username` is your Oracle Cloud Infrastructure user name. The `password` is your Oracle Cloud Infrastructure auth token. See Working with Auth Tokens for more information.

   See CREATE_CREDENTIAL Procedure for details on the parameters and their values, depending on the Cloud Object Storage.

3. Set initialization parameters to specify the Cloud Object Storage URL for a bucket for SQL trace files and to specify the credentials to access the Cloud Object Storage.

a. Set database property `DEFAULT_LOGGING_BUCKET` to specify the logging bucket on Cloud Object Storage.

   For example, if you create the bucket with Oracle Cloud Infrastructure Object Storage:

   ```
   SET DEFINE OFF;
   ALTER DATABASE PROPERTY SET
      DEFAULT_LOGGING_BUCKET = 'https://objectstorage.us-
   ashburn-1.oraclecloud.com/n/namespace-string/b/bucket_name/o/';
   ```

   Where *namespace-string* is the Oracle Cloud Infrastructure Object Storage namespace and *bucket_name* is the name of the bucket you previously created. See Understanding Object Storage Namespaces for more information.

   See Regions and Availability Domains for a list of regions.

   The Cloud Object Store you use for SQL Tracing files can be any Cloud Object Store that Autonomous Database supports.

b. Set the database property `DEFAULT_CREDENTIAL` to the credential you created in Step 2.

   Including the schema name with the credential is required. In this example the schema is "`ADMIN`".

   For example:

   ```
   ALTER DATABASE PROPERTY SET DEFAULT_CREDENTIAL = 'ADMIN.DEF_CRED_NAME';
   ```

Enable SQL Tracing on Autonomous Database

Shows the steps to enable SQL tracing for the database session.

> **Note:**
>
> If you enable SQL tracing your application performance for the session may be degraded while the trace collection is enabled. This negative performance impact is expected due to the overhead of collecting and saving trace data.

Before you enable SQL tracing you must configure the database to save SQL Trace files. See Configure SQL Tracing on Autonomous Database for more information.

To enable SQL tracing, do the following:

1. (Optional) Set a client identifier for the application. This step is optional but is recommended. SQL tracing uses the client identifier as a component of the trace file name when the trace file is written to Cloud Object Store.

   For example:

   ```
   BEGIN
     DBMS_SESSION.SET_IDENTIFIER('sqlt_test');
   END;
   /
   ```

2. (Optional) Set a module name for the application. This step is optional but is recommended. SQL tracing uses the module name as a component of the trace file name when the trace file is written to Cloud Object Store.

   For example:

   ```
   BEGIN
     DBMS_APPLICATION_INFO.SET_MODULE('modname', null);
   END;
   /
   ```

3. Enable the SQL Trace facility.

   ```
   ALTER SESSION SET SQL_TRACE = TRUE;
   ```

4. Run your workload.

   This step involves running the entire application or specific parts of the application. While you run your workload in the database session, SQL tracing data is collected.

5. Disable SQL Tracing.

   When you disable SQL tracing the collected data for the session is written to a table in your session and to a trace file in the bucket you configure when you set up SQL tracing. See Disable SQL Tracing on Autonomous Database for details.

Disable SQL Tracing on Autonomous Database

Shows the steps to disable SQL tracing on Autonomous Database.

To disable SQL tracing, do the following:

1. Disable the SQL Trace facility.

   ```
   ALTER SESSION SET SQL_TRACE = FALSE;
   ```

2. (Optional) as needed for your environment, you may want to reset the database property `DEFAULT_LOGGING_BUCKET` to clear the value for the logging bucket on Cloud Object Storage.

   For example:

   ```
   ALTER DATABASE PROPERTY SET DEFAULT_LOGGING_BUCKET = '';
   ```

When you disable SQL tracing, the tracing data collected while the session runs with tracing enabled is copied to a table and sent to a trace file on Cloud Object Store. You have two options to view trace data:

- View and analyze SQL Trace data in the trace file saved to Cloud Object Store. See View Trace File Saved to Cloud Object Store on Autonomous Database for more information.
- View and analyze SQL Trace data saved to the view `SESSION_CLOUD_TRACE`. See View Trace Data in SESSION_CLOUD_TRACE View on Autonomous Database for more information.

View Trace File Saved to Cloud Object Store on Autonomous Database

Describes the output file naming for SQL trace files and shows the commands to use `TKPROF` to organize and view trace file data.

You use SQL trace file data to analyze application performance on Autonomous Database. When you disable SQL trace in your database session, data is written to the Cloud Object Store bucket configured with `DEFAULT_LOGGING_BUCKET`.

The SQL Trace facility writes the trace data collected in the session to Cloud Object Store in the following format:

default_logging_bucket`/sqltrace/`*`clientID`*`/`*`moduleName`*`/sqltrace_`*`numID1`*`_`*`numID2`*`.trc`

The components of the file name are:

- *default_logging_bucket*: is the value of the `DEFAULT_LOGGING_BUCKET` database property. See Configure SQL Tracing on Autonomous Database for more information.
- *`clientID`*: is the client identifier. See Enable SQL Tracing on Autonomous Database for more information.
- *`moduleName`*: is the module name. See Enable SQL Tracing on Autonomous Database for more information.
- *`numID1`*`_`*`numID2`*: are two identifiers that the SQL Trace facility provides. The *`numID1`* and *`numID2`* numeric values uniquely distinguish each trace file name from other sessions using tracing and creating trace files in the same bucket in the Cloud Object Storage.

  When the database service supports parallelism and a session runs a parallel query, the SQL Trace facility can produce multiple trace files with different *`numID1`* and *`numID2`* values.

> **Note:**
>
> When SQL tracing is enabled and disabled multiple times within the same session, each trace iteration generates a separate trace file in Cloud Object Store. To avoid overwriting previous traces that were generated in the session, subsequently generated files follow the same naming convention and add a numeric suffix to the trace file name. This numeric suffix starts with the number 1 and is incremented by 1 for each tracing iteration thereafter.

For example, the following is a sample generated trace file name when you set the client identifier to "`sql_test`" and the module name to "`modname`":

```
sqltrace/sqlt_test/modname/sqltrace_5415_56432.trc
```

You can run `TKPROF` to translate the trace file into a readable output file.

1. Copy the trace file from Object Store to your local system.
2. Navigate to the directory in which the trace file is saved.
3. Run the `TKPROF` utility from the operating system prompt using the following syntax:

```
tkprof filename1 filename2 [waits=yes|no] [sort=option] [print=n]
    [aggregate=yes|no] [insert=filename3] [sys=yes|no] [table=schema.table]
    [explain=user/password] [record=filename4] [width=n]
```

The input and output files are the only required arguments.

4. To view online Help, invoke `TKPROF` without arguments.

See "Tools for End-to-End Application Tracing" in Oracle Database SQL Tuning Guide for information about using the `TKPROF` utility.

View Trace Data in SESSION_CLOUD_TRACE View on Autonomous Database

When you enable SQL Tracing, the same trace information that is saved to the trace file on Cloud Object Store is available in the `SESSION_CLOUD_TRACE` view in the session where the tracing was enabled.

While you are still in the database session you can view SQL tracing data in the `SESSION_CLOUD_TRACE` view. The `SESSION_CLOUD_TRACE` view includes two columns: `ROW_NUMBER` and `TRACE`:

```
DESC SESSION_CLOUD_TRACE

Name       Null? Type
---------- ----- ------------------------------
ROW_NUMBER       NUMBER
TRACE            VARCHAR2(32767)
```

The `ROW_NUMBER` specifies the ordering for trace data found in the `TRACE` column. Each line of trace output written to a trace file becomes a row in the table and is available in the `TRACE` column.

After you disable SQL tracing for the session, you can run queries on the `SESSION_CLOUD_TRACE` view.

For example:

```
SELECT trace FROM SESSION_CLOUD_TRACE ORDER BY row_number;
```

The data in `SESSION_CLOUD_TRACE` persists for the duration of the session. After you log out or close the session, the data is no longer available.

If SQL Trace is enabled and disabled multiple times within the same session, `SESSION_CLOUD_TRACE` shows the trace data for all the iterations cumulatively. Thus, re-enabling tracing in a session after previously disabling tracing does not remove the trace data produced by the earlier iteration.

Monitor Performance with Autonomous Database Metrics

You can monitor the health, capacity, and performance of your databases with metrics, alarms, and notifications. You can use Oracle Cloud Infrastructure Console or Monitoring APIs to view metrics.

- View Metrics for an Autonomous Database Instance
  Shows the steps to view the Autonomous Database metrics.
- View Metrics for Autonomous Databases in a Compartment
  Shows the steps to view metrics for Autonomous Databases in a compartment.

- Autonomous Database Metrics and Dimensions
  You can limit the instances where you see metrics with dimensions. The available dimensions include: workload type, instance display name, region, and the instance OCID.
- Use Custom Metrics on Autonomous Database
  Describes how to create and publish custom metrics on Autonomous Database.

View Metrics for an Autonomous Database Instance

Shows the steps to view the Autonomous Database metrics.

To view metrics you must have the required access as specified in an Oracle Cloud Infrastructure policy (whether you're using the Console, the REST API, or another tool). See Getting Started with Policies for information on policies.

Perform the following steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To view metrics for an Autonomous Database instance:

1. On the Details page, under **Resources**, click **Metrics**.
2. There is a chart for each metric. In each chart you can select the **Interval** and **Statistic** or use the default values.

> **Note:**
>
> The following table shows the default metrics shown on the Oracle Cloud Infrastructure Console.
> See Available Metrics: oci_autonomous_database for a list of all the database metrics and dimensions.

Metric Name	Description
CPU utilization	CPU utilization expressed as a percentage, aggregated across all consumer groups. The utilization percentage is reported with respect to the number of CPUs the database is allowed to use, which is the number of ECPUs. If your database uses OCPUs, the number of CPUs allowed is two times the number of OCPUs.
Storage utilization	The percentage of provisioned storage capacity currently in use. Represents the total allocated space for all tablespaces.
Sessions	The number of sessions in the database.
Execute count	The number of user and recursive calls that ran SQL statements during the selected interval.
Running statements	The number of running SQL statements, aggregated across all consumer groups, during the selected interval.

Metric Name	Description
Queued statements	The number of queued SQL statements, aggregated across all consumer groups, during the selected interval.
Database availability	The database is available for connections during the selected time interval (data for this metric lags by 5 minutes). Possible values for this metric: • 1 = Database is Available • 0 = Database is Unavailable You can set an alarm that is triggered if the database is not available (value 0). **Note:** Availability is calculated based on the "Monthly Uptime Percentage" described in the Oracle PaaS and IaaS Public Cloud Services Pillar Document document under **Delivery Policies** (see Autonomous Database Availability Service Level Agreement).
Failed connections	Shows the total number of failed connections to the database during the selected interval. A connection is counted as failed when a connection attempt reaches the database and logs any of the following errors: `ORA-12514`, `ORA-12519` or `ORA-12529`.

To create an alarm on a metric, in a metric chart or table, click **Options** and select **Create an Alarm on this Query**. See Managing Alarms for information on setting and using alarms.

For more information about metrics see Available Metrics: oci_autonomous_database.

You can also use the Monitoring API to view metrics. See Monitoring API for more information.

- View Logs and Audit Trails
  Shows the steps to view the Autonomous Database logs and audit trails.

View Logs and Audit Trails

Shows the steps to view the Autonomous Database logs and audit trails.

Note:

To view logs and audit trials you must have the required access as specified in an Oracle Cloud Infrastructure policy (whether you're using the Console, the REST API, or another tool). See Getting Started with Policies for information on policies.

To view audit trails and logs for an Autonomous Database instance:

1. On the Details page, under **Resources**, click **Metrics**.
2. In the Metrics area click **View audit and logs**.

3. In the Logging area, click **Logs** to view log information.
4. In the logging area, click **Audit** to view audit information.

See Audit Autonomous Database and Audit Logs for more information.

View Metrics for Autonomous Databases in a Compartment

Shows the steps to view metrics for Autonomous Databases in a compartment.

To view metrics you must have the required access as specified in an Oracle Cloud Infrastructure policy (whether you're using the Console, the REST API, or other tool). See Getting Started with Policies for information on policies.

- Open the Oracle Cloud Infrastructure Console by clicking the  next to Oracle Cloud.
- From the left navigation list click **Observability & Management**. Under **Monitoring**, click **Service Metrics**.

To use the metrics service to view Autonomous Database metrics:

1. On the Service Metrics page, under **Compartment** select your compartment.
2. On the Service Metrics page, under **Metric Namespace** select **oci_autonomous_database**.
3. If there are multiple Autonomous Databases in the compartment you can show metrics aggregated across the Autonomous Databases by selecting **Aggregate Metric Streams**.
4. If you want to limit the metrics you see, next to **Dimensions** click **Add** (click **Edit** if you have already added dimensions).
    a. In the **Dimension Name** field select a dimension.
    b. In the **Dimension Value** field select a value.
    c. Click **Done**.

    In the Edit dimensions dialog click **+Additional Dimension** to add an additional dimension. Click **x** to remove a dimension.

To create an alarm on a specific metric, click **Options** and select **Create an Alarm on this Query**. See Managing Alarms for information on setting and using alarms.

Autonomous Database Metrics and Dimensions

You can limit the instances where you see metrics with dimensions. The available dimensions include: workload type, instance display name, region, and the instance OCID.

Use dimensions by selecting values in the Oracle Cloud Infrastructure Console Service Metrics page or by setting dimension values with the API. See View Metrics for Autonomous Databases in a Compartment to view metrics and to select metric dimensions.

Use Custom Metrics on Autonomous Database

Describes how to create and publish custom metrics on Autonomous Database.

Autonomous Database provides many database related metrics that are available through the OCI Monitoring service. In addition, you can create custom metrics to collect, publish, and analyze your own metrics. A custom metric collects specified data from your Autonomous Database instance and publishes the data to the OCI Monitoring service using a REST API provided by the OCI SDK.

Prerequisites

1. Obtain Oracle Cloud Infrastructure access through Oracle Cloud Free Tier or a paid Cloud Account.

   See Get an Oracle Cloud Account for more information.

2. Create an Autonomous Database instance or have access to an existing Autonomous Database instance.

   See Provision an Autonomous Database Instance for more information.

3. Obtain ADMIN credentials for your Autonomous Database instance.

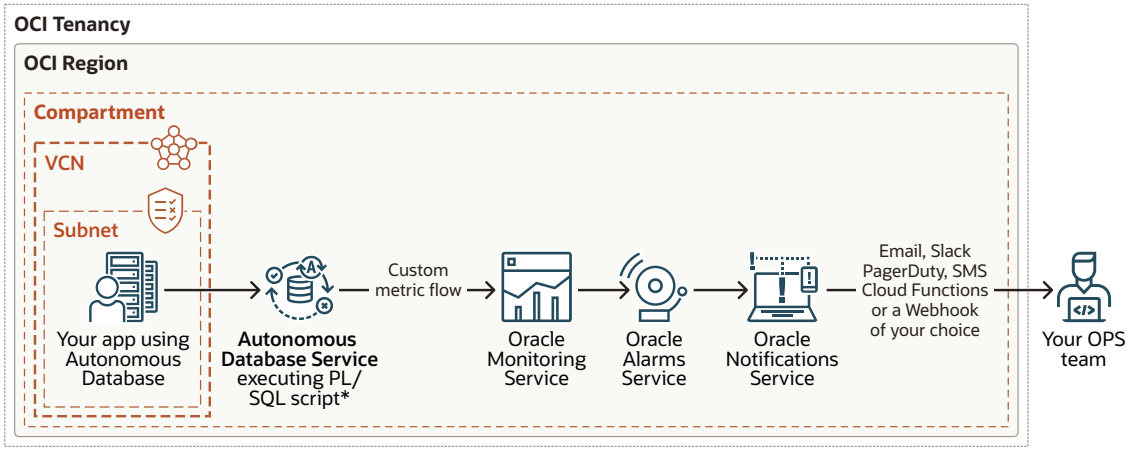
4. Use Database Actions or any of the Oracle Database clients, such as SQL Developer or SQL*Plus to connect to the database.

   See Connect to Autonomous Database for more information.

5. You need access to the OCI Monitoring service and OCI Identity and Access Management. See Publishing Custom Metrics Using the API for more information.

Overview of Custom Metrics with OCI Monitoring Service

The following figure shows the steps to create and publish custom metrics on Autonomous Database. This shows that you collect metric data with a script running on your Autonomous Database instance and you publish this data to the OCI Monitoring service, and how you can create alarms and notifications for metric values.



- Create and deploy a PL/SQL script on your Autonomous Database instance. This script runs periodically to compute, collect, and publish custom metrics to the OCI Monitoring service.
- The Autonomous Database instance can be on a public or on a private endpoint. The communication between the Autonomous Database instance and the OCI Monitoring service takes place on the Oracle Cloud network. This means, to publish metrics on the OCI Monitoring service you do not need to create a service gateway.

Create and Publish Custom Metrics

To create and publish custom metrics:

1. In OCI Identity and Access Management create a dynamic group for your Autonomous Database instance and create a policy to authorize the dynamic group to post metrics to the OCI Monitoring service.

   For example:

a. In the Oracle Cloud Infrastructure console click **Identity & Security**.

b. Under **Identity** click **Domains** and select an identity domain (or create a new identity domain).

c. Under **Identity domain**, click **Dynamic groups**.

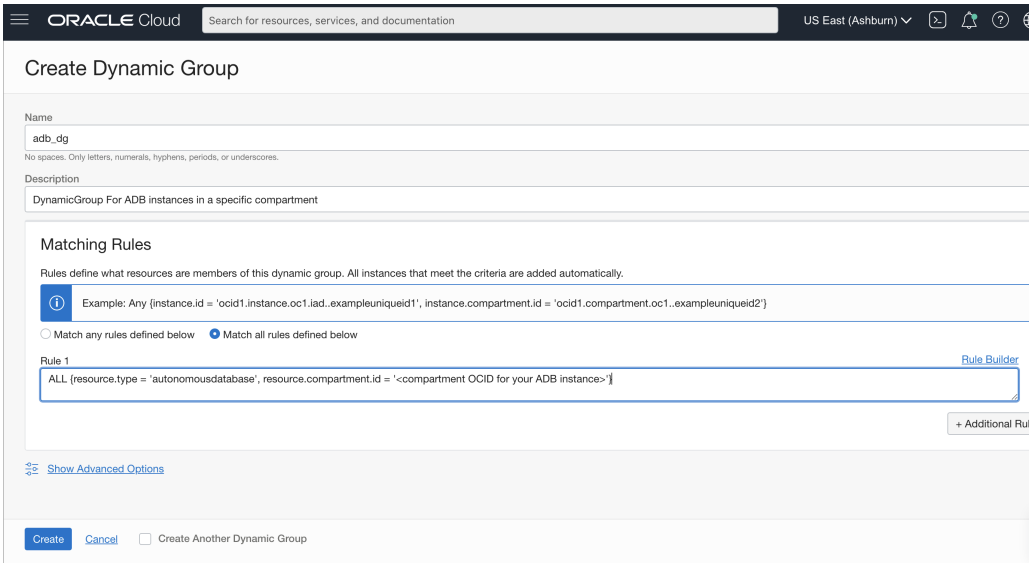
d. Click **Create dynamic group** and enter a **Name**, a **Description**, and a rule.

   For example, create a dynamic group named `adb_dg` for your Autonomous Database instance and create a rule:

   ```
   ALL {resource.type = 'autonomousdatabase', resource.compartment.id =
   'OCID_of_compartment'}
   ```

   Alternatively, you can select a single Autonomous Database instance instead of all the instances in the compartment:

   ```
   ALL {resource.type = 'autonomousdatabase', resource.id =
   'OCID_of_autonomousdatabase_instance'}
   ```

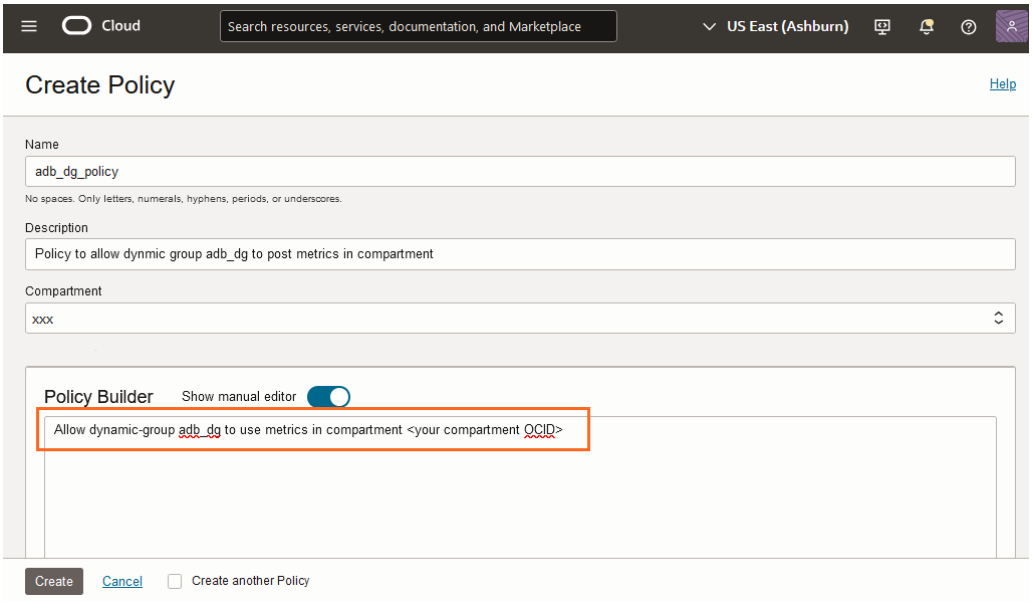


e. Click **Create**.

f. Create an OCI Identity and Access Management (IAM) policy to authorize the dynamic group.

   For example, create a policy to authorize the dynamic group `adb_dg` to post metrics to the OCI Monitoring service with the policy `adb_dg_policy`:

   ```
   Allow dynamic-group adb_dg to use metrics in compartment
   OCID_of_compartment
   ```



After you define the policy, the dynamic group with the definition `adb_dg` is authorized to post metrics in the compartment.

At this point, the PL/SQL routine running on Autonomous Database can't post any metrics to the OCI Monitoring service because there are no Autonomous Database users authorized to publish metrics to the OCI Monitoring service. You provide this capability when you enable the resource principal in the following step (step 2c).

2. Create a new database user or schema with the requisite privileges in your Autonomous Database instance or update an existing database user or schema with the requisite privileges.

   For example:

   a. Create a database user or use an existing user to publish metrics. .

      For example, create a new user `ECOMMERCE_USER` on the Autonomous Database instance:

      ```
      CREATE USER ECOMMERCE_USER IDENTIFIED BY "password";
      ```

      Now you have the new user `ECOMMERCE_USER`. The following steps are the same with this user or with the user you specify. However if you create a user with a different name you need to replace `ECOMMERCE_USER` with the alternative user name.

   b. Grant requisite Oracle Database related privileges to the user.

      ```
      GRANT CREATE TABLE, ALTER ANY INDEX, CREATE PROCEDURE,
            CREATE JOB, SELECT ANY TABLE, EXECUTE ANY PROCEDURE, UPDATE ANY
      TABLE,
            CREATE SESSION, UNLIMITED TABLESPACE, CONNECT, RESOURCE
            TO ECOMMERCE_USER;
      GRANT  SELECT ON "SYS"."V_$PDBS" TO ECOMMERCE_USER;
      GRANT  EXECUTE ON "C##CLOUD$SERVICE"."DBMS_CLOUD" to ECOMMERCE_USER;
      GRANT  SELECT ON SYS.DBA_JOBS_RUNNING TO ECOMMERCE_USER;
      ```

c. Enable Oracle Database credentials for the Oracle Cloud Resource Principal and give access to the user.

This step connects the dynamic group `adb_dg` you created in step 1 to the database user `ECOMMERCE_USER`, giving the user authorization to post metrics to the OCI Monitoring service.

For example:

```
EXEC DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL(username =>
'ECOMMERCE_USER');
```

See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

As an alternative to Resource Principal authentication, you can use OCI Native credentials for authentication. See CREATE_CREDENTIAL Procedure for more information.

d. (Optional) Verify the operations you completed in the previous step.

For example:

```
SELECT OWNER, CREDENTIAL_NAME FROM DBA_CREDENTIALS WHERE
CREDENTIAL_NAME = 'OCI$RESOURCE_PRINCIPAL'  AND OWNER =  'ADMIN';
```

To check if a user other than ADMIN has access to OCI Resource Principal, for example to check if the `ECOMMERCE_USER` has access to the database credential, use the `DBA_TAB_PRIVS` view:

```
SELECT * FROM DBA_TAB_PRIVS WHERE
DBA_TAB_PRIVS.GRANTEE='ECOMMERCE_USER';
```

3. Create a PL/SQL script that publishes metric data and schedule the script to run on your Autonomous Database instance.

a. For example, create a PL/SQL script with a metric that counts of the rows of the Star Schema Benchmark sample data and publishes to the OCI Monitoring service.

```
CREATE OR REPLACE PROCEDURE publish_lineorder_metric (p_sql_statement
IN VARCHAR2)
IS
l_result NUMBER;
l_compartment_ocid VARCHAR2(255);
l_db_name VARCHAR2(255);
l_region VARCHAR2(255);
l_json_payload CLOB;
l_cloud_identity CLOB;
BEGIN

-- 1. Execute the SQL statement and get the result
EXECUTE IMMEDIATE p_sql_statement INTO l_result;

 -- 2. Get compartment OCID, DB name, and region from cloud identity
 SELECT cloud_identity INTO l_cloud_identity FROM v$pdbs;

 l_compartment_ocid := JSON_VALUE(l_cloud_identity,
```

```
'$.COMPARTMENT_OCID');
 l_db_name := JSON_VALUE(l_cloud_identity, '$.DATABASE_NAME');
 l_region := JSON_VALUE(l_cloud_identity, '$.REGION');

DBMS_OUTPUT.PUT_LINE(l_cloud_identity);

 -- 3. Construct the JSON payload for the metric

 l_json_payload :=
 '{' ||
 ' "metricData": [' ||
 ' {' ||
 ' "namespace": "custom_metrics",' ||
 ' "compartmentId": "' || l_compartment_ocid || '",' ||
 ' "name": "lineorder_metric",' ||
 ' "dimensions": {' ||
 ' "dbName": "' || l_db_name || '"' ||
 ' },' ||
 ' "datapoints": [' ||
 ' {' ||
 ' "timestamp": "' || TO_CHAR(SYSTIMESTAMP AT TIME ZONE 'UTC', 'YYYY-MM-
DD"T"HH24:MI:SS.FF3"Z"') || '",' ||
 ' "value": ' || l_result ||
 ' }' ||
 ' ]' ||
 ' }' ||
 ' ]' ||
 '}';

 -- 4. Publish the metric using DBMS_CLOUD.SEND_REQUEST

 DECLARE
 l_response DBMS_CLOUD_TYPES.resp;
 BEGIN
 l_response := DBMS_CLOUD.SEND_REQUEST(
 credential_name => 'OCI$RESOURCE_PRINCIPAL',
 uri => 'https://telemetry-ingestion.' || l_region || '.oraclecloud.com/
20180401/metrics',
 method => 'POST',
 body => UTL_RAW.CAST_TO_RAW(l_json_payload)
 );

 -- 5. Check response status

 IF DBMS_CLOUD.GET_RESPONSE_STATUS_CODE(l_response) = 200 THEN
 DBMS_OUTPUT.PUT_LINE('Metric published successfully!');
 ELSE
 DBMS_OUTPUT.PUT_LINE('Error publishing metric: ' ||
DBMS_CLOUD.GET_RESPONSE_STATUS_CODE(l_response));
 DBMS_OUTPUT.PUT_LINE(DBMS_CLOUD.GET_RESPONSE_TEXT(l_response));
 END IF;

 EXCEPTION
 WHEN OTHERS THEN
 DBMS_OUTPUT.PUT_LINE('Error during SEND_REQUEST: ' || SQLERRM);
 END;
```

```
EXCEPTION
 WHEN OTHERS THEN
 DBMS_OUTPUT.PUT_LINE('Error: ' || SQLERRM);
END;
/
```

See PostMetricData for more information.

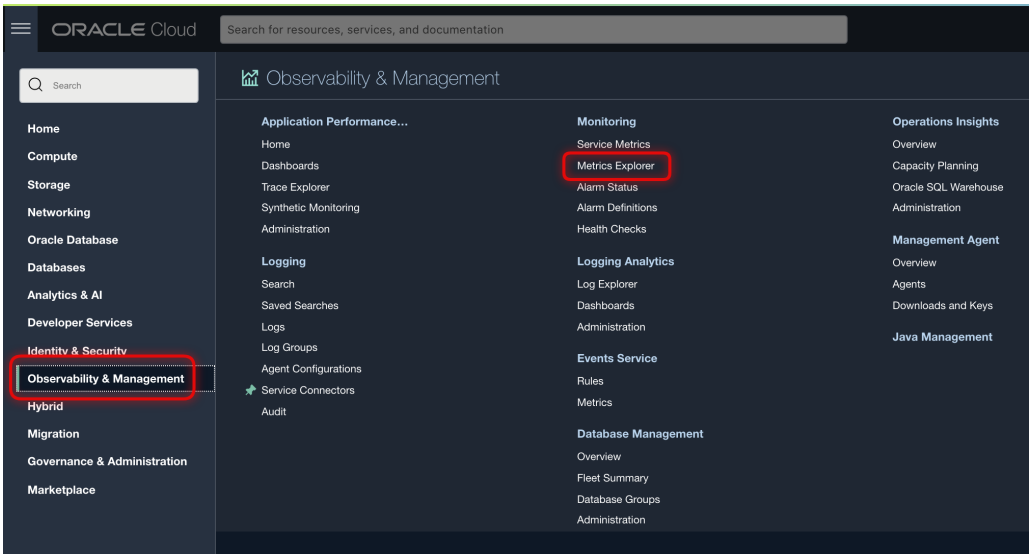
b. Schedule a job to call and publish the custom metric data point once per minute. In production the schedule would be configured to run according to your needs.

```
BEGIN
   DBMS_SCHEDULER.CREATE_JOB (
      job_name => 'publish_lineorder_count_job',
      job_type => 'STORED_PROCEDURE',
      job_action => 'ECOMMERCE_USER.publish_lineorder_metric',
      number_of_arguments => 1,
      start_date => SYSTIMESTAMP,
      repeat_interval => 'FREQ=MINUTELY;INTERVAL=1',
      enabled => FALSE,
      AUTO_DROP => FALSE,
      comments => 'Publishes the count of rows in SSB.LINEORDER every 1
minute'
);

  DBMS_SCHEDULER.SET_JOB_ARGUMENT_VALUE (
       job_name => 'publish_lineorder_count_job',
       argument_position => 1,
       argument_value => 'SELECT COUNT(*) FROM SSB.LINEORDER'
 );

   DBMS_SCHEDULER.ENABLE('ECOMMERCE_USER.PUBLISH_LINEORDER_COUNT_JOB');
END;
```

4. Explore the published custom metrics on the Metrics Explorer.
   a. From the Oracle Cloud Infrastructure Console navigation menu, click **Observability & Management**.
   b. Under **Observability & Management**, click **Metrics Explorer**.



c. In Metrics Explorer, select the namespace as `custom_metrics_from_adb`, `resourceGroup` as `ecommerece_adb`, and metric name as `lineorder_metric` (the name of the metric you created in the PL/SQL script).

   All the metadata and dimensions you set for custom metrics are available. You can construct Metrics Query Language (MQL) queries to analyze these metrics according to your needs and use case. You may also set up Oracle Cloud Alarms on the metric stream to alert your operational team. This automates the observability loop for the Autonomous Database metrics of your choice. Note that in the Metrics Explorer the metric screen by default shows graph view and you have the option to enable list view to highlight data points.

After you create custom metrics you can use the metrics just like any predefined metrics in OCI Monitoring service. This means you can analyze custom metrics with the Metrics Query Language (MQL) and set up alarms and notifications to notify you whenever an event of interest occurs.

See the following for more information:

- OCI Monitoring service
- PostMetricData API for publishing custom metrics
- dynamic groups and resource principles

Available Metrics: oci_autonomous_database

This topic describes the metrics emitted by the Database service in the `oci_autonomous_database` namespace.

To access a subset of the metrics from the Autonomous Database details page:

See View Metrics for an Autonomous Database Instance for details.

To access the complete metrics for the service:

See View Metrics for Autonomous Databases in a Compartment for details.

To access the metrics from the Metrics Explorer:

1. Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
2. From the Oracle Cloud Infrastructure left navigation menu click **Observability & Management**.
3. On the **Observability & Management** page, under **Monitoring**, select **Metrics Explorer**.
4. Use the **Metrics Explorer** to create a query and enter the **Metric namespace** value `oci_autonomous_database`.
5. Under Metric dimensions enter the **deploymentType** value **Shared**.
6. Enter a value in the **Metric name** field.

Database service metrics for Autonomous Databases include the following **dimensions**:

- **AUTONOMOUSDBTYPE**
  The type of Autonomous Database, Autonomous Data Warehouse (ADW) or Autonomous Transaction Processing (ATP).
- **deploymentType**
  The Exadata infrastructure type, shared or dedicated. When using the Console to view default metric charts for multiple Autonomous Databases, you must specify this dimension.

  This topic covers Oracle Autonomous Database Serverless metrics. See Monitor Databases with Autonomous Database Metrics for information on Oracle Autonomous Database on Dedicated Exadata Infrastructure metrics.
- **DISPLAYNAME**
  The friendly name of the Autonomous Database.
- **REGION**
  The **region** in which the Autonomous Database resides.
- **RESOURCEID**
  The **OCID** of the Autonomous Database.
- **RESOURCENAME**
  The name of the Autonomous Database.

The metrics listed in the following table are automatically available for any Autonomous Database that you create. You do not need to enable monitoring on the resource to get these metrics.

> **Note:**
>
> Valid alarm intervals are 5 minutes or greater due to the frequency at which these metrics are emitted. See To create an alarm for details on creating alarms.

In the following table, metrics that are marked with an asterisk (*) can be viewed only on the **Service Metrics** page of the Oracle Cloud Infrastructure console. All metrics can be filtered by the dimensions described in this topic.

Metric	Metric Display Name	Unit	Description	Collection Frequency
AverageActiveSessions	Average Active Sessions	integer	The average number of sessions that are actively executing or waiting for resources. Statistic: Mean Interval: 1 minute	1 minute
ConnectionLatency	Connection Latency	Milliseconds	The time taken to connect to an Oracle Autonomous Database Serverless instance in each region from a Compute service virtual machine in the same region. Statistic: Max Interval: 5 minutes	5 minutes
BlockingSessions	Blocking Sessions	count	The number of sessions that have been blocking other sessions for more than 60 seconds. Statistic: Max Interval: 1 minute	1 minute
ConnectionsDroppedByClient	Connections Dropped By Client	count	The number of connections dropped by the client. Statistic: Sum Interval: 1 minute	1 minute
ConnectionsDroppedByServer	Connections Dropped By Server	count	The number of connections dropped by the server. Statistic: Sum Interval: 1 minute	1 minute
CpuTime*	CPU Time	seconds per second	Average rate of accumulation of CPU time by foreground sessions in the database over the time interval. Statistic: Mean Interval: 1 minute	NA
CpuUtilization	CPU Utilization	percent	The CPU usage expressed as a percentage, aggregated across all consumer groups. The utilization percentage is reported with respect to the number of CPUs the database is allowed to use. Statistic: Mean Interval: 1 minute	NA
CurrentLogons*	Current Logons	count	The number of successful logons during the selected interval. Statistic: Count Interval: 1 minute	1 minute

Metric	Metric Display Name	Unit	Description	Collection Frequency
`DatabaseAvailability`	Database Availability	count	The database is available for connections during the selected time interval (data for this metric lags by 5 minutes). Possible values for this metric: • 1 = Database is Available • 0 = Database is Unavailable You can set an alarm that is triggered if the database is not available (value 0).	5 minutes

Note:

Availability is calculated based on the "Monthly Uptime Percentage" described in the Oracle PaaS and IaaS Public Cloud Services Pillar Document document under **Delivery Policies** (see Autonomous Database Availability Service Level Agreement).

Metric	Metric Display Name	Unit	Description	Collection Frequency
			Statistic: Mean Interval: 1 minute	
DBBlockChanges	DB Block Changes	count	The number of changes that were part of an update or delete operation that were made to all blocks in the SGA. Such changes generate redo log entries and become permanent changes to the database if the transaction is committed. This approximates total database work. This statistic indicates the rate at which buffers are being dirtied, during the selected time interval. Statistic: Sum Interval: 1 minute	1 minute
DBTime*	DB Time	seconds per second	The amount of time database user sessions spend executing database code (CPU Time + WaitTime). DB Time is used to infer database call latency, because DB Time increases in direct proportion to both database call latency (response time) and call volume. It is calculated as the average rate of accumulation of database time by foreground sessions in the database over the time interval. Statistic: Mean Interval: 1 minute	NA
ExecuteCount	Execute Count	count	The number of user and recursive calls that executed SQL statements during the selected interval. Statistic: Sum Interval: 1 minute	NA
FailedConnections*	Failed Connections	count	The number of failed database connections. Statistic: Sum Interval: 1 minute	1 minute
FailedLogons	Failed Logons	count	The number of log ons that failed because of an invalid user name and/or password, during the selected interval. The counter is reset to 0 by a successful login. Statistic: Mean Interval: 1 minute	1 minute

Metric	Metric Display Name	Unit	Description	Collection Frequency
HardParseCount	Parse Count (Hard)	count	The number of parse calls (real parses) during the selected time interval. A hard parse is an expensive operation in terms of memory use, because it requires Oracle to allocate a workheap and other memory structures and then build a parse tree. Statistic: Sum Interval: 1 minute	1 minute
LogicalReads	Session Logical Reads	count	The sum of "db block gets" plus "consistent gets", during the selected time interval. This includes logical reads of database blocks from either the buffer cache or process private memory.	1 minute
ParseCount*	Parse Count (Total)	count	The number of hard and soft parses during the selected interval. Statistic: Sum Interval: 1 minute	1 minute
ParseFailureCount	Parse Count (Failures)	count	The number of parse failures during the selected time interval. Statistic: Sum Interval: 1 minute	1 minute
PeerLag	Peer Lag	Seconds	The total time in seconds by which the Disaster Recovery peer lags behind its Primary database. For a local Disaster Recovery peer, the metric is displayed under the Primary database. For a cross-region Disaster Recovery peer, this metric is displayed under the Disaster Recovery peer. Statistic: Max Interval: 1 minute	1 minute
PhysicalReads	Physical Reads	count	The number of data blocks read from disk, during the selected time interval. Statistic: Sum Interval: 1 minute	1 minute
PhysicalReadTotalBytes	Physical Read Total Bytes	count	The size in bytes of disk reads by all database instance activity including application reads, backup and recovery, and other utilities, during the selected time interval. Statistic: Sum Interval: 1 minute	1 minute

Metric	Metric Display Name	Unit	Description	Collection Frequency
PhysicalWrites	Physical Writes	count	The number of data blocks written to disk, during the selected time interval. Statistic: Sum Interval: 1 minute	1 minute
PhysicalWriteTotalBytes	Physical Write Total Bytes	count	The size in bytes of all disk writes for the database instance including application activity, backup and recovery, and other utilities, during the selected time interval. Statistic: Sum Interval: 1 minute	1 minute
QueryLatency	Query Latency	Milliseconds	The time taken to display the results of a simple query on the user's screen. Statistic: Max Interval: 5 minutes	5 minutes
QueuedStatements	Queued Statements	count	The number of queued SQL statements, aggregated across all consumer groups, during the selected interval. Statistic: Sum Interval: 1 minute	NA
RedoGenerated	Redo Generated	count	Amount of redo generated in bytes, during the selected time interval. Statistic: Sum Interval: 1 minute	1 minute
RunningStatements	Running Statements	count	The number of running SQL statements, aggregated across all consumer groups, during the selected interval. Statistic: Mean Interval: 1 minute	1 minute
Sessions	Sessions	count	The number of sessions in the database. Statistic: Sum Interval: 1 minute	1 minute
SessionUtilization	Session Utilization	Percent	The maximum session utilization expressed as a percentage, aggregated across all consumer groups. Statistic: Mean Interval: 1 minute	1 minute
SQLNetBytesFromClient	Bytes Received via SQL*Net from Client	count	The number of bytes received from the client over Oracle Net Services, during the selected time interval. Statistic: Sum Interval: 1 minute	1 minute

Metric	Metric Display Name	Unit	Description	Collection Frequency
SQLNetBytesFromDBLink	Bytes Received via SQL*Net from DBLink	count	The number of bytes received from a database link over Oracle Net Services, during the selected time interval. Statistic: Sum Interval: 1 minute	1 minute
SQLNetBytesToClient	Bytes Sent via SQL*Net to Client	count	The number of bytes sent to the client from the foreground processes, during the selected time interval. Statistic: Sum Interval: 1 minute	1 minute
SQLNetBytesToDBLink	Bytes Sent via SQL*Net to DBLink	count	The number of bytes sent over a database link, during the selected time interval. Statistic: Sum Interval: 1 minute	1 minute
StorageAllocated*	Storage Space Allocated	GB	Amount of space allocated to the database for all tablespaces, during the interval. Statistic: Max Interval: 1 hour	1 hour
StorageMax	Maximum Storage Space	GB	Maximum amount of storage reserved for the database during the interval. Statistic: Max Interval: 1 hour	1 hour
StorageUsed*	Storage Space Used	GB	Maximum amount of space used during the interval. Statistic: Max Interval: 1 hour	1 hour
StorageUtilization	Storage Utilization	percent	The percentage of the reserved maximum storage currently allocated for all database tablespaces. Represents the total reserved space for all tablespaces. Statistic: Mean Interval: 1 hour	1 hour
TransactionCount*	Transaction Count	count	The combined number of user commits and user rollbacks during the selected interval. Statistic: Sum Interval: 1 minute	1 minute
UserCalls*	User Calls	count	The combined number of logons, parses, and execute calls during the selected interval. Statistic: Sum Interval: 1 minute	1 minute

Metric	Metric Display Name	Unit	Description	Collection Frequency
UserCommits	User Commits	count	The number of user commits during the selected time interval. When a user commits a transaction, the generated redo that reflects the changes made to database blocks must be written to disk. Commits often represent the closest thing to a user transaction rate. Statistic: Sum Interval: 1 minute	1 minute
UserRollbacks	User Rollbacks	count	Number of times users manually issue the ROLLBACK statement or an error occurs during a user's transactions, during the selected time interval. Statistic: Sum Interval: 1 minute	1 minute
WaitTime*	Wait Time	seconds per second	Average rate of accumulation of non-idle wait time by foreground sessions in the database over the time interval. Statistic: Mean Interval: 1 minute	1 minute

Use Sample Data Sets in Autonomous Database

For users who want to start using the service without creating their own tables, Autonomous Database provides the read-only Sales History and Star Schema Benchmark data sets.

These data sets are provided as Oracle Database schemas SH and SSB respectively. Any user can query these data sets without any manual configuration.

Notes for the Sample Data sets:

- Both the SH and SSB are provided as schema-only users, so you cannot unlock or drop those users or set a password.
- The storage size of the sample data sets is not counted towards or billed for in your database's storage.
- The sample data sets SH and SSB are not included in backups and do not impact the size or speed of either backup or restore operations.

Sales History (SH) Schema

The SH schema provides a small data set you can use to run the sample queries in the *Oracle Database Data Warehousing Guide*. Note that you need to prefix the table names with the schema name SH in your queries. For example, the following query shows you how the SQL function RANK() works:

```
SELECT channel_desc, TO_CHAR(SUM(amount_sold), '9,999,999,999') SALES$,
   RANK() OVER (ORDER BY SUM(amount_sold)) AS default_rank,
   RANK() OVER (ORDER BY SUM(amount_sold) DESC NULLS LAST) AS custom_rank
```

```
FROM sh.sales, sh.products, sh.customers, sh.times, sh.channels, sh.countries
WHERE sales.prod_id=products.prod_id AND sales.cust_id=customers.cust_id
  AND customers.country_id = countries.country_id AND
sales.time_id=times.time_id
  AND sales.channel_id=channels.channel_id
  AND times.calendar_month_desc IN ('2000-09', '2000-10')
  AND country_iso_code='US'
GROUP BY channel_desc;
```

For more information on the SH schema see *Sample Schemas* and *Schema Diagrams*.

Star Schema Benchmark (SSB) Schema

The SSB schema provides a well-known large sample data set. The SSB schema contains 1 TB of data (storage of the sample data sets does not count towards your database storage). You can use this schema to test the performance of your service. You can run the sample queries on this schema with different database services, HIGH, MEDIUM, LOW and with different number of ECPUs (OCPUs if your database uses OCPUs) to test the performance of Autonomous Database.

The SSB schema contains the tables: `lineorder`, `customer`, `supplier`, `part`, and `dwdate`. See Sample Star Schema Benchmark (SSB) Queries and Analytic Views for a list of sample queries you can use against the SSB schema. Note that you need to prefix the table names with the schema name `SSB` in your queries.

For more information on database services, see Database Service Names for Autonomous Database.

Use the Autonomous Database Free Container Image

Use the Autonomous Database Free Container Image to run Autonomous Database in a container in your own environment, without requiring access to the Oracle Cloud Infrastructure Console or to the internet.

- About the Free Container Image
  You can access the Autonomous Database Free Container Image from a repository and run it in your local environment.
- Container Registry Locations for Autonomous Database Free Container Image
  There are multiple locations where you can obtain Autonomous Database Free Container Image, including: Oracle Cloud Infrastructure Registry (Container Registry) and GitHub.
- Start Free Container Image
  After you download Autonomous Database Free Container Image, you can start the image in a container.
- Perform Database Operations using adb-cli
  The `adb-cli` command-line utility can be used to perform database operations after the container is up and running.
- Connect to an Free Container Image
  When the Free Container Image is running in a container, you can connect to an Autonomous Database instance.
- Migrate Data Between Autonomous Database Free Containers

About the Free Container Image

You can access the Autonomous Database Free Container Image from a repository and run it in your local environment.

Autonomous Database provides a fully managed Oracle Database that is available on Oracle Cloud Infrastructure. On Oracle Cloud Infrastructure, you perform lifecycle management operations and run Autonomous Database using the Oracle Cloud Infrastructure Console and you connect to your database through the public internet or through a private network that you set up (depending on your network configuration and security needs).

> **Note:**
>
> Autonomous Database supports both the Oracle Autonomous Database Free Container Image 19c and Oracle Autonomous Database Free Container Image 23ai versions.

The Autonomous Database Free Container Image provides an alternative to run Autonomous Database in a container in your own environment, without requiring access to Oracle Cloud Infrastructure Console or to the internet. When you run Autonomous Database in a container, the container provides a local, isolated environment with additional options for development, testing, and exploration of Oracle Autonomous Database features.

- Autonomous Database Free Container Image License
  Describes licensing for Autonomous Database Free Container Image.
- Autonomous Database Free Container Image Features
  The Autonomous Database Free Container Image provides many of the features available with Autonomous Database Serverless.
- Autonomous Database Free Container Image Recommendations and Restrictions
  Describes requirements and restrictions for Free Container Image.

Autonomous Database Free Container Image License

Describes licensing for Autonomous Database Free Container Image.

Autonomous Database Free Container Image is subject to licensing.

The container image you have selected and all of the software that it contains is licensed under Oracle Free Use Terms and Conditions that are provided in the container image. Your use of the container is subject to the terms of those licenses.

Autonomous Database Free Container Image Features

The Autonomous Database Free Container Image provides many of the features available with Autonomous Database Serverless.

- Each Free Container Image provides two Autonomous Database instances, one instance with Data Warehouse workload type and one instance with Transaction Processing workload type.
- The database will be started with either the Transaction Processing workload type or the Data Warehouse workload type based on the workload type you specify during startup.

- You can perform database operations using the `adb-cli` command-line utility.
- The Free Container Image resource allocation is 4 ECPUs and 20 GB of storage, and allows a maximum of 30 simultaneous database sessions.
- Each Free Container Image supports the Autonomous Database consumer groups:
  - Data Warehouse workload: You connect through HIGH, MEDIUM, or LOW services
  - Transaction Processing workload: You connect through HIGH, MEDIUM, LOW, TP, or TPURGENT services

  See About Autonomous Database Workload Types for more information.
- The Free Container Image provides support for the following built-in database tools:
  - Database Actions
  - Oracle REST Data Services (ORDS)
  - Oracle APEX
  - Oracle Database API for MongoDB

Autonomous Database Free Container Image Recommendations and Restrictions

Describes requirements and restrictions for Free Container Image.

Recommendations for Resource Allocation for the Free Container Image

The following is the recommended resource allocation for the Free Container Image:

- 4 CPUs
- 8 GB memory

Restrictions for Free Container Image

- There is no automatic patching or maintenance windows for the Free Container Image. The repository provides the latest version of the Free Container Image. Check the repository to find newer versions of Free Container Image.
- The following Autonomous Database built-in tools are not supported:
  - Graph
  - Oracle Machine Learning
  - Data Transforms
- When Autonomous Database runs in a container, the container provides a local Autonomous Database instance. The container image does not include the features that are only available through the Oracle Cloud Infrastructure Console or the APIs. Some features that are available in-database and also available through the Oracle Cloud Infrastructure Console can still be used through in-database commands, such as resetting the `ADMIN` password. The following lists some of the features that are not available:

Feature	Available or Not Available
Backup an Instance	Not available
Choose a Character Set	Not available
Clone an Instance	Not available
Create an Elastic Pool	Not available
Customer-managed keys	Not available

Feature	Available or Not Available
Database rename	Not available
Data Safe	Not available
Disable compute auto scaling	Not available
Disable Built-in Database Tools	Not available
Disable storage auto scaling	Not available
Disaster recovery options, including Autonomous Data Guard and Backup-Based Disaster Recovery.	Not available
Download wallet	Not available
Enable Built-in Database Tools	Not available
Enable compute auto scaling	Not available
Enable storage auto scaling	Not available
Join an Elastic Pool	Not available
Network ACLs	Not available
Oracle Cloud Infrastructure events	Not available
Performance hub	Not available
Private endpoints	Not available
Real Application Testing	Not available
Resource Principal based authentication	Not available
Restart Instance	Not available
Restore an Instance	Not available
Rotate wallet	Not available
Sample Schema	Not available
Scale down CPU and storage	Not available
Scale up CPU and storage	Not available
Selecting the Instance Patch Level	Not available
Start Instance	Not available
Stop Instance	Not available

> **Note:**
>
> When you run Free Container Image in a container, you can start an instance, stop an instance, or restart an instance by starting, stopping or restarting the container.

Container Registry Locations for Autonomous Database Free Container Image

There are multiple locations where you can obtain Autonomous Database Free Container Image, including: Oracle Cloud Infrastructure Registry (Container Registry) and GitHub.

You can obtain the Autonomous Database Free Container Image in multiple locations. The examples shown use `podman` commands (see Podman for more information).

1. Start a podman virtual machine.

For example:

```
podman machine init
podman machine set --cpus 4 --memory 8192
podman machine start
```

2. Obtain the Free Container Image.

   Oracle Cloud Infrastructure Registry:

   For example, use `podman` command to pull the latest Free Container Image image:

   ```
   podman pull container-registry.oracle.com/database/adb-free:latest-23ai
   ```

   > **Note:**
   >
   > Oracle Autonomous Database Serverless also supports 19c container images.
   >
   > - For 19c container image names, specify `latest` tag as your database version.
   > - In the above example, to pull the Autonomous Database Free Container Image 19c , use the following command in the image name tag:
   >
   >   ```
   >   podman pull container-registry.oracle.com/database/adb-free:latest
   >   ```
   >
   > Throughout the document, the image name tag refers to the `latest-23ai` version.

   For more details and additional information, search for "Oracle Autonomous Database Free" on Oracle Cloud Infrastructure Registry.

   GitHub Packages:

   For example, use `podman` command to pull the Free Container Image from: GitHub Packages:

   ```
   podman pull ghcr.io/oracle/adb-free:latest-23ai
   ```

3. Verify the image.

   For example:

   ```
   podman images container-registry.oracle.com/database/adb-free:latest-23ai
   ```

Start Free Container Image

After you download Autonomous Database Free Container Image, you can start the image in a container.

The database will be started with either the Transaction Processing workload type or the Data Warehouse workload type based on the workload type you specify.

- Start a container to run Free Container Image.

For example, with podman:

```
podman run -d \
-p 1521:1522 \
-p 1522:1522 \
-p 8443:8443 \
-p 27017:27017 \
-e WORKLOAD_TYPE='ATP' \
-e WALLET_PASSWORD=*** \
-e ADMIN_PASSWORD=*** \
--cap-add SYS_ADMIN \
--device /dev/fuse \
--name adb-free \
container-registry.oracle.com/database/adb-free:latest-23ai
```

Command notes:

- Autonomous Database Serverless also supports 19c container images.
  - For 19c container image names, specify `latest` tag as your database version.
  - In the above example, to start a container to run the Autonomous Database Free Container Image 19c, use the following command in the image name tag:

    ```
    container-registry.oracle.com/database/adb-free:latest
    ```

- Throughout the document, the image name tag refers to the `latest-23ai` version.
- The `WORKLOAD_TYPE` can be either `ATP` or `ADW`. The default value is `ATP`.
- By default, the database is named either `MYATP` or `MYADW`, depending on the passed `WORKLOAD_TYPE` value. Optionally, if you want a different database name than the default, you can set the `DATABASE_NAME` parameter. The database name can contain only alphanumeric characters.
- On container startup, the corresponding `MY<WORKLOAD_TYPE>.pdb` is downloaded and plugged in from a public Object Storage bucket.
- The wallet is generated using the provided `WALLET_PASSWORD`.
- It is necessary to change the `ADMIN_PASSWORD` upon initial login.
- Ensure you check the following requirements when you create or modify the `ADMIN_PASSWORD`:
  - The password must be between 12 and 30 characters long and must include at least one uppercase letter, one lowercase letter, and one numeric character.
  - The password cannot contain the username.
- Ensure you check the following requirements when you create or modify the `WALLET_PASSWORD`:
  - The password must be between 8 and 30 characters long and include alphabetic characters combined with numeric characters or special characters.
- For an OFS mount, the container starts with `SYS_ADMIN` capability. Also, virtual device `/dev/fuse` must be accessible.
- This `-p` options specify that the following ports are forwarded to the container process:

Port	Description
1521	TLS
1522	mTLS
8443	HTTPS port for ORDS / APEX and Database Actions
27017	Mongo API

If you are behind a corporate proxy, include `-e` options to specify environment variables for proxies. For example, with podman:

```
podman run -d \
-p 1521:1522 \
-p 1522:1522 \
-p 8443:8443 \
-p 27017:27017 \
-e WORKLOAD_TYPE='ATP' \
-e WALLET_PASSWORD=*** \
-e ADMIN_PASSWORD=*** \
-e http_proxy=http://example-corp-proxy.com:80/ \
-e https_proxy=http://example-corp-proxy.com:80/ \
-e no_proxy=localhost,127.0.0.1 \
-e HTTP_PROXY=http://example-corp-proxy.com:80/  \
-e HTTPS_PROXY=http://example-corp-proxy.com:80/  \
-e NO_PROXY=localhost,127.0.0.1 \
--cap-add SYS_ADMIN \
--device /dev/fuse \
--name adb-free \
container-registry.oracle.com/database/adb-free:latest-23ai
```

Perform Database Operations using `adb-cli`

The `adb-cli` command-line utility can be used to perform database operations after the container is up and running.

To use `adb-cli`, you can define the following alias for convenience:

```
alias adb-cli="podman exec <container_name> adb-cli"
```

Available Commands

You can view the list of available commands using the following command:

```
adb-cli --help
Usage: adb-cli [OPTIONS] COMMAND [ARGS]...
    ADB-S Command Line Interface (CLI) to perform container-runtime database
operations
Options:
    -v, --version Show the version and exit.
    --help Show this message and exit.
Commands:
    add-database
    change-password
```

Add Database

You can add a database using the following command:

```
adb-cli add-database --workload-type "ATP" --admin-password
"Welcome_MY_ATP_1234"
```

Change Password

You can change the admin password using the following command:

```
adb-cli change-password --database-name "MYATP"
    --old-password "Welcome_MY_ATP_1234" --new-password "Welcome_12345"
```

Connect to an Free Container Image

When the Free Container Image is running in a container, you can connect to an Autonomous Database instance.

To connect to an Autonomous Database instance running in a container, you can either setup the wallet and connect or use a TLS walletless connection.

- Connect to ORDS, APEX, or Database Actions from an Autonomous Database Container Image
  The container `hostname` is used to generate self-signed SSL certificates to serve HTTPS traffic on port 8443. Oracle APEX and Database Actions can be accessed using the container host (or simply `localhost`).
- Available TNS Aliases for Connecting to an Autonomous Database Container Image
  You can use any one of the following aliases to connect to an Autonomous Database Container Image.
- Connect to an Autonomous Database Container Image with a Wallet
  Describes the steps to connect to an Autonomous Database container image with a Wallet.
- Connect to an Autonomous Database Container Image without a Wallet (TLS)
  Describes the steps to connect to an Autonomous Database container image without a Wallet (TLS).

Connect to ORDS, APEX, or Database Actions from an Autonomous Database Container Image

The container `hostname` is used to generate self-signed SSL certificates to serve HTTPS traffic on port 8443. Oracle APEX and Database Actions can be accessed using the container host (or simply `localhost`).

Application	URL
Oracle APEX	https://localhost:8443/ords/apex
Database Actions	https://localhost:8443/ords/sql-developer

> **Note:**
>
> For additional databases plugged in using `adb-cli add-database` command, use the URL formats `https://localhost:8443/ords/{database_name}/apex` and `https://localhost:8443/ords/{database_name}/sql-developer` to access APEX and Database Actions respectively.

Available TNS Aliases for Connecting to an Autonomous Database Container Image

You can use any one of the following aliases to connect to an Autonomous Database Container Image.

Table 3-7 Available TNS Aliases

Protocol	Transaction Processing Workload	Data Warehouse Workload
mTLS	• myatp_medium • myatp_high • myatp_low • myatp_tp • myatp_tpurgent	• myadw_medium • myadw_high • myadw_low
TLS	• myatp_medium_tls • myatp_high_tls • myatp_low_tls • myatp_tp_tls • myatp_tpurgent_tls	• myadw_medium_tls • myadw_high_tls • myadw_low_tls

The TNS alias mappings for these connect strings are in `$TNS_ADMIN/tnsnames.ora`. See Manage Concurrency and Priorities on Autonomous Database for information on the service names in `tnsnames.ora`.

Connect to an Autonomous Database Container Image with a Wallet

Describes the steps to connect to an Autonomous Database container image with a Wallet.

Perform the following steps to setup a wallet and connect:

1. Copy the generated wallet to your host.

   When you start the container, Autonomous Database generates a wallet in `/u01/app/oracle/wallets/tls_wallet`.

   For example, copy the generated wallet to the local folder `/scratch/tls_wallet`:

   ```
   podman cp adb-free:/u01/app/oracle/wallets/tls_wallet /scratch/tls_wallet
   ```

   This copies the wallet to the folder: `/scratch/tls_wallet`.

2. Set the value of the `TNS_ADMIN` environment variable to the wallet directory.

   For example:

   ```
   export TNS_ADMIN=/scratch/tls_wallet
   ```

3. If you want to connect to a remote host where the Free Container Image is running, replace `localhost` in `$TNS_ADMIN/tnsnames.ora` with the remote host FQDN.

   For example:

   ```
   sed -i 's/localhost/example.com/g' $TNS_ADMIN/tnsnames.ora
   ```

4. Connect to the Autonomous Database instance.

   For example use `sqlplus` to connect to the **Transaction Processing** workload Autonomous Database instance:

   ```
   sqlplus admin/password@myatp_low_tls
   ```

   For example, use `sqlplus` to connect to the **Data Warehouse** workload Autonomous Database instance:

   ```
   sqlplus admin/password@myadw_low_tls
   ```

Connect to an Autonomous Database Container Image without a Wallet (TLS)

Describes the steps to connect to an Autonomous Database container image without a Wallet (TLS).

To connect without a wallet, you need to update your client's truststore with the self-signed certificate generated at container start.

Perform the following steps to setup a TLS connection and connect.

1. Copy `/u01/app/oracle/wallets/tls_wallet/adb_container.cert` from the container and update your system truststore.

   For example:

   ```
   podman cp adb-free:/u01/app/oracle/wallets/tls_wallet/adb_container.cert
   adb_container.cert
   sudo cp adb_container.cert /etc/pki/ca-trust/source/anchors
   sudo update-ca-trust
   ```

2. Connect to the Autonomous Database instance.

   For example, use `sqlplus` to connect to the Transaction Processing workload Autonomous Database instance:

   ```
   sqlplus admin/password@myatp_low
   ```

   For example, use `sqlplus` to connect to the Data Warehouse workload Autonomous Database instance:

   ```
   sqlplus admin/password@myadw_low
   ```

Migrate Data Between Autonomous Database Free Containers

When a new version of the Free Container Image is available, you can migrate data from a container to another container.

For example, use existing data that you created in a container by migrating that data to the latest version of Free Container Image when a new update is available.

1. Create a podman volume.

   For example:

   ```
   podman volume create adb_container_volume
   ```

2. Verify the volume mountpoint.

   The mountpoint is a podman managed directory location.

   ```
   podman inspect adb_container_volume

   [
        {
             "Name": "adb_container_volume",
             "Driver": "local",
             "Mountpoint": "/share/containers/storage/volumes/
   adb_container_volume/_data",
             "CreatedAt": "2023-09-11T21:23:34.305877073Z",
             "Labels": {},
             "Scope": "local",
             "Options": {},
             "MountCount": 0,
             "NeedsCopyUp": true,
             "NeedsChown": true
        }
   ]
   ```

3. Start the source container, mounting the volume to `/u01/data` in the container.

   For example:

   ```
   podman run -d \
   -p 1521:1522 \
   -p 1522:1522 \
   -p 8443:8443 \
   -p 27017:27017 \
   -e WORKLOAD_TYPE='ATP' \
   -e WALLET_PASSWORD=*** \
   -e ADMIN_PASSWORD=*** \
   --cap-add SYS_ADMIN \
   --device /dev/fuse \
   --name source_adb_container \
   --volume adb_container_volume:/u01/data \
   container-registry.oracle.com/database/adb-free:latest-23ai
   ```

> **Note:**
>
> Autonomous Database Serverless also supports 19c container images.
>
> - For 19c container image names, specify `latest` tag as your database version.
> - In the above example, to start the Autonomous Database Free Container Image 19c , use the following command in the image name tag:
>
>   `container-registry.oracle.com/database/adb-free:`**`latest`**
>
> Throughout the document, the image name tag refers to the `latest-23ai` version.

4. This example assumes you have previously created your data in schema, `app_user`.
5. Export the data from `app_user`'s schema to the container volume.
   a. Connect as `ADMIN` and create `ORA_EXP_DIR` directory pointing to `/u01/data`.

```
sqlplus admin/**************@myadw_high

SQL> exec DBMS_CLOUD_CONTAINER_ADMIN.create_export_directory('/u01/
data');


SQL> select directory_path from dba_directories where
directory_name='ORA_EXP_DIR';

DIRECTORY_PATH
--------------------------------------------------------------------------------
--------
/u01/data
```

   b. Run the export job in schema mode and for both `ADMIN` and `APP_USER` schemas.

```
SET scan off
SET serveroutput ON
SET escape off

DECLARE
    h1 NUMBER;
    s VARCHAR2(1000):=NULL;
    errorvarchar VARCHAR2(100):= 'ERROR';
    tryGetStatus NUMBER := 0;
    success_with_info EXCEPTION;
    PRAGMA EXCEPTION_INIT(success_with_info, -31627);
BEGIN
    h1 := dbms_datapump.OPEN(operation => 'EXPORT', job_mode =>
'SCHEMA', job_name => 'EXPORT_MY_ADW_4', version => 'COMPATIBLE');
    tryGetStatus := 1;
    dbms_datapump.add_file(handle => h1, filename =>
'EXPORT_MY_ADW.LOG', directory => 'ORA_EXP_DIR', filetype =>
DBMS_DATAPUMP.KU$_FILE_TYPE_LOG_FILE);
    dbms_datapump.metadata_filter(handle => h1, name => 'SCHEMA_EXPR',
```

```
VALUE => 'IN(''ADMIN'', ''APP_USER'')');
    dbms_datapump.add_file(handle => h1, filename => 'MY_ADW_%L.DMP',
directory => 'ORA_EXP_DIR', filesize => '500M',  filetype =>
DBMS_DATAPUMP.KU$_FILE_TYPE_DUMP_FILE);
    dbms_datapump.start_job(handle => h1, skip_current => 0, abort_step
=> 0);
    dbms_datapump.detach(handle => h1);
    errorvarchar := 'NO_ERROR';
EXCEPTION
    WHEN OTHERS THEN
    BEGIN
        IF ((errorvarchar = 'ERROR')AND(tryGetStatus=1)) THEN
            DBMS_DATAPUMP.DETACH(h1);
        END IF;
    EXCEPTION
    WHEN OTHERS THEN
        NULL;
    END;
    RAISE;
END;
/
```

6. Verify the export.

   List the files in the container `/u01/data`.

   ```
   podman exec -it source_adb_container bash
   cd /u01/data
   ```

   Verify the export log (`export log`), check for any errors and successful completion.

7. Stop and remove the source container.

   ```
   podman stop source_adb_container
   podman rm source_adb_container
   ```

   > **Note:**
   >
   > The `adb_container_volume` will live across container restarts and removals

8. Start a destination container mounting the same volume to `/u01/data` in the container.

   ```
   podman run -d \
   -p 1521:1522 \
   -p 1522:1522 \
   -p 8443:8443 \
   -p 27017:27017 \
   -e WORKLOAD_TYPE='ATP' \
   -e WALLET_PASSWORD=*** \
   -e ADMIN_PASSWORD=*** \
   --cap-add SYS_ADMIN \
   --device /dev/fuse \
   --name dest_adb_container \
   ```

```
--volume adb_container_volume:/u01/data \
container-registry.oracle.com/database/adb-free:latest-23ai
```

9. Import data in destination container.

   Connect as `ADMIN` and create `ORA_EXP_DIR` directory pointing to `/u01/data`.

```
SQL> exec DBMS_CLOUD_CONTAINER_ADMIN.create_export_directory('/u01/data');

PL/SQL procedure successfully completed.

SQL> select directory_path from dba_directories where
directory_name='ORA_EXP_DIR';

DIRECTORY_PATH
--------------------------------------------------------------------------------
-----
/u01/data
```

10. Run the import PL/SQL commands.

```
SET scan off
SET serveroutput ON
SET escape off

DECLARE
    h1 NUMBER;
    s VARCHAR2(1000):=NULL;
    errorvarchar VARCHAR2(100):= 'ERROR';
    tryGetStatus NUMBER := 0;
    success_with_info EXCEPTION;
    PRAGMA EXCEPTION_INIT(success_with_info, -31627);
BEGIN
    h1 := dbms_datapump.OPEN(operation => 'IMPORT', job_mode => 'SCHEMA',
job_name => 'IMPORT_MY_ADW_4', version => 'COMPATIBLE');
    tryGetStatus := 1;
    dbms_datapump.add_file(handle => h1, filename => 'IMPORT_MY_ADW.LOG',
directory => 'ORA_EXP_DIR', filetype =>
DBMS_DATAPUMP.KU$_FILE_TYPE_LOG_FILE);
    dbms_datapump.metadata_filter(handle => h1, name => 'SCHEMA_EXPR',
VALUE => 'IN(''ADMIN'', ''APP_USER'')');
    dbms_datapump.add_file(handle => h1, filename => 'MY_ADW_%L.DMP',
directory => 'ORA_EXP_DIR', filesize => '500M',  filetype =>
DBMS_DATAPUMP.KU$_FILE_TYPE_DUMP_FILE);
    dbms_datapump.start_job(handle => h1, skip_current => 0, abort_step =>
0);
    dbms_datapump.detach(handle => h1);
    errorvarchar := 'NO_ERROR';
EXCEPTION
    WHEN OTHERS THEN
    BEGIN
        IF ((errorvarchar = 'ERROR')AND(tryGetStatus=1)) THEN
            DBMS_DATAPUMP.DETACH(h1);
        END IF;
    EXCEPTION
    WHEN OTHERS THEN
```

```
            NULL;
        END;
        RAISE;
    END;
    /
```

11. Verify the import log.
12. Verify the import data.

4
Tutorials

- Load and Analyze Your Data with Autonomous Database
  Learn about Autonomous Database Serverless and learn how to create an Autonomous Database in just a few clicks. Then load data into your database, query it, and visualize it.
- Self-service Tools for Everyone Using Oracle Autonomous Database
  Introduces the suite of data tools built into Oracle Autonomous Database.
- Manage and Monitor Autonomous Database
  Connect using secure wallets and monitor your Autonomous Database instances.
- Integrate, Analyze and Act on All Data using Autonomous Database
- Load the Autonomous Database with Data Studio
- QuickStart Java applications with Autonomous Database
  Follow these easy steps to connect your Java applications to Autonomous Database using the Oracle JDBC driver and Universal Connection Pool.

Load and Analyze Your Data with Autonomous Database

Learn about Autonomous Database Serverless and learn how to create an Autonomous Database in just a few clicks. Then load data into your database, query it, and visualize it.

 Load and analyze data in Autonomous Database Workshop

- Provision a new Autonomous Database instance.
- Run Queries on the sample data sets
- Upload files to the Oracle Cloud Infrastructure Object Storage, create sample tables, load data into them from files on the OCI Object Storage
- Query files on the Oracle Cloud Infrastructure Object Storage directly without loading them to your database
- Visualize your data using Oracle Analytics Desktop

Self-service Tools for Everyone Using Oracle Autonomous Database

Introduces the suite of data tools built into Oracle Autonomous Database.

 Data Studio - Self-service tools for everyone using Oracle Autonomous Database

This is an overview Data Studio workshop to introduce you to different tools using a sample data analysis task. After this workshop you will have a working knowledge on how to use these tools.

The tools covered are:

- Catalog

- Data Load
- Data Transforms
- Data Analysis
- Data Insight

After completing this workshop, you can look at the individual in-depth workshops for each tool.

Manage and Monitor Autonomous Database

Connect using secure wallets and monitor your Autonomous Database instances.

Manage and Monitor Autonomous Database Workshop

- Get comfortable with Oracle's public cloud services
- Provision a new Autonomous Database instance Serverless
- Run sample queries against the sample data sets
- Load data from the object store
- Query external data from the object store
- Scale an Autonomous Database instance

Integrate, Analyze and Act on All Data using Autonomous Database

This workshop uses Oracle Cloud's modern data platform for analytics to enable MovieStream to better understand their customers, predict churn and provide targeted recommendations and promotions.

- MovieStream stores its data across Oracle Object Storage and Autonomous Database.
- Data is captured from various sources into a landing zone in object storage. This data is then processed, cleansed, transformed, and optimized, and stored in a gold zone on object storage.
- After the data is curated, it is loaded into an Autonomous Database where users can take advantage of a variety of Autonomous Database capabilities: advanced Oracle SQL to glean customer insights, Oracle Machine Learning to predict churn, and Oracle Spatial and Graph to provide targeted promotions.

Innovate with Autonomous Database's built-in Select AI, Graph, ML and Spatial analytics

Load the Autonomous Database with Data Studio

Data Studio is built into every Autonomous Database, and provides tools to load, transform and prepare data from a wide range of source systems.

This workshop provides practical examples of how to load data from several different systems, as well as how to link to, and feed data in from, data in cloud storage systems.

Load the Autonomous Database with Data Studio

QuickStart Java applications with Autonomous Database

Follow these easy steps to connect your Java applications to Autonomous Database using the Oracle JDBC driver and Universal Connection Pool.

QuickStart Java applications with Oracle Autonomous Database

5
Features

- Use Data Lake with Autonomous Database
  Learn the benefits of using Data Lake with Autonomous Database.
- The Data Studio Overview Page
  The Data Studio comprises of the Data Load, the Data Analysis, the Data Insights, Catalog, Data Marketplace and the Data Share tool.
- High Availability
  You can configure and manage Autonomous Data Guard and maintain continuous availability using disaster recovery.
- Security
  Autonomous Database stores all data with inbuilt security. Only authenticated users and applications can access the data when they connect to the database.
- Performance Monitor and Management
  Oracle Autonomous Database Serverless includes several features that automatically monitor, analyze and optimize the performance of your Autonomous Database.
- Using Oracle Graph with Autonomous Database
  Oracle Graph with Autonomous Database enables you to create graphs from data in your Autonomous Database. With graphs you can analyze your data based on connections and relationships between data entities.
- Use Oracle Machine Learning with Autonomous Database
- Creating Analytic Views
  You can create Analytic Views and view information about them. You can also edit and perform other actions on them.
- Use Select AI for Natural Language Interaction with your Database
- JSON Document Stores
  Autonomous Database has full support for data represented as JSON documents. In Autonomous Databases, JSON documents can coexist with relational data.
- Cache Messages with Singleton Pipes and Using Pipes for Persistent Messaging
- Full-Text Search Queries
  The PL/SQL package `DBMS_CLOUD` enables you to build a text index on the object store files, which allows you to search the text and use wildcards with your search. You can use `JSON_TEXTCONTAINS` to do full-text JSON search on files stored in Object Storage.
- Use Oracle Workspace Manager on Autonomous Database
  Workspace Manager provides an infrastructure that enables applications to create workspaces and group different versions of table row values in different workspaces.
- Use and Manage Elastic Pools on Autonomous Database
  Use an elastic pool to consolidate your Autonomous Database instances, in terms of their allocation of compute resources, and to provide up to 87% cost savings.
- Migrate to Autonomous Database
  Database migration is the process of migrating your database from a source database to a target database.

- Use Database In-Memory with Autonomous Database
  The Database In-Memory feature available on Autonomous Database can significantly improve performance for real-time analytics and mixed workloads.

Use Data Lake with Autonomous Database

Learn the benefits of using Data Lake with Autonomous Database.

- About Data Lake with Autonomous Database
  Oracle Autonomous Database is a versatile solution for accommodating any type of data and workload.
- What is a Data Lake?
  Data lakes are centralized repositories designed to store vast amounts of raw data in their native format until the data is needed for analysis.
- Key Data Lake Features of Autonomous Database
  Oracle Autonomous Database is designed to seamlessly support data lake workloads, eliminating the need for management or installation. It delivers robust capabilities to handle various data formats across different cloud environments, ensuring flexible and comprehensive data analysis.

About Data Lake with Autonomous Database

Oracle Autonomous Database is a versatile solution for accommodating any type of data and workload.

Autonomous Database provides cost-efficient storage, with a cost per TB comparable to object stores, while supporting diverse data types like JSON, Graph, and Vector. With Autonomous Database, businesses can consolidate their data onto a single platform. They can leverage converged capabilities such as Oracle Machine Learning (OML), Graph, Spatial, Vector, and Blockchain to manage their data comprehensively.

For organizations that already have existing data lakes on other platforms, Oracle Autonomous Database integrates seamlessly, allowing businesses to benefit from Autonomous Database's advanced features without disrupting their current setups.

To learn more, try the LiveLabs Build a Data Lake with Autonomous Data Warehouse.

What is a Data Lake?

Data lakes are centralized repositories designed to store vast amounts of raw data in their native format until the data is needed for analysis.

They are highly flexible and scalable, making them a powerful complement to traditional data warehouses by enabling organizations to store and process various types of data, including structured, semi-structured, and unstructured.

Key attributes of a Data Lake:

- **Open File and Table Formats**

  Data lakes store data in open file formats, such as CSV, Parquet, and table formats like Iceberg. This ensures interoperability and flexibility in data processing by allowing multiple engines to write and read these datasets.
- **Support for Multiple Data Processing Engines**

Data lakes are compatible with various data processing engines, such as Apache Spark, Presto, and Hive, enabling diverse analytical workloads.

- **Schema-on-Read**

  Data lakes often use a schema-on-read approach, meaning there is no need to define a schema upfront. This allows for rapid data ingestion, where data can be loaded without prior structuring, much like object stores that “capture data now, and ask questions later".

- **Support for Unstructured Data**

  Beyond structured data, data lakes can store unstructured data such as images (JPG), documents (PDF, Word), and other binary data, offering a comprehensive storage solution.

Key Data Lake Features of Autonomous Database

Oracle Autonomous Database is designed to seamlessly support data lake workloads, eliminating the need for management or installation. It delivers robust capabilities to handle various data formats across different cloud environments, ensuring flexible and comprehensive data analysis.

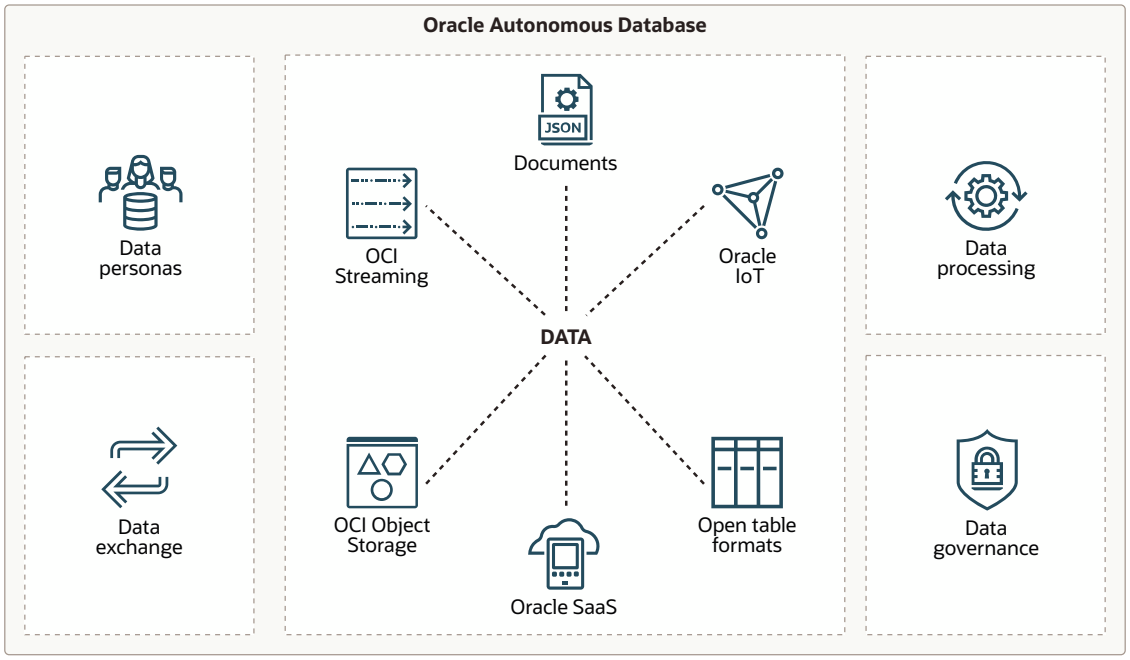
- Ready for Data Lake Workloads
  Oracle Autonomous Database is fully ready for data lake workloads out-of-the-box, requiring no additional components. This readiness extends to key data lake tasks such as data transformation, metadata management, and integration with popular data lake tools—all available from day one without extra setup.
- Multi-Cloud Support
  For organizations that already have existing data lakes on other platforms, Autonomous Database integrates seamlessly, allowing businesses to benefit from Autonomous Database advanced features without disrupting their current setups.
- End-To-End Data Format Support
  Oracle Autonomous Database is designed with the flexibility to handle a broad spectrum of data formats, making it a universal solution for diverse data sources and workloads.
- Enhanced Capabilities: Autonomous Database for Unstructured Data Management
  While Oracle Database is recognized for its powerful processing of structured and semi-structured data, Autonomous Database extends its capabilities to handle unstructured datasets as well.
- Flexible Metadata Management
  Oracle Autonomous Database provides users with various ways to define metadata for their datasets, making data management more adaptable and efficient.
- Federated Metadata Support
  Autonomous Database supports a federated metadata catalog, enabling users to unify metadata from different sources into a single view, providing a unified interface for metadata management.
- Collaboration
  After users finish their analysis, they often need to share their results with others. Oracle Autonomous Database makes sharing easy by offering several ways to collaborate, providing unique advantages over other databases, such as integrated security features, open protocols, and seamless cloud connectivity.
- Broad Compatibility with Oracle Database Tools
  The Autonomous Database environment is fully compatible with a wide array of Oracle database tools.

Ready for Data Lake Workloads

Oracle Autonomous Database is fully ready for data lake workloads out-of-the-box, requiring no additional components. This readiness extends to key data lake tasks such as data transformation, metadata management, and integration with popular data lake tools—all available from day one without extra setup.

This comprehensive readiness is what makes Autonomous Database stand out, offering an integrated, hassle-free experience that accelerates time-to-insight for data lake workloads. This means users can immediately start handling data lake tasks without any setup or configuration, making it a true plug-and-play solution for data lake environments. This built-in capability simplifies operations, reduces maintenance costs, and ensures higher reliability with fewer errors.

Autonomous Database provides a set of tools for all user types, from developers to business analysts, making the platform universal and accessible.



Developers can use tools such as the PL/SQL API for advanced operations, scripting, and automation, allowing seamless integration with existing tools and creating customized database solutions efficiently. See Autonomous Database Supplied Package Reference, for more information.

For business users, Data Studio can be used—a web-based interface for simplifying data interaction, exploration, and visualization. Data Studio enables non-technical users to derive insights, create reports, and collaborate effectively, reducing complexity and supporting informed decision-making. See The Data Studio Overview Page, for more information.

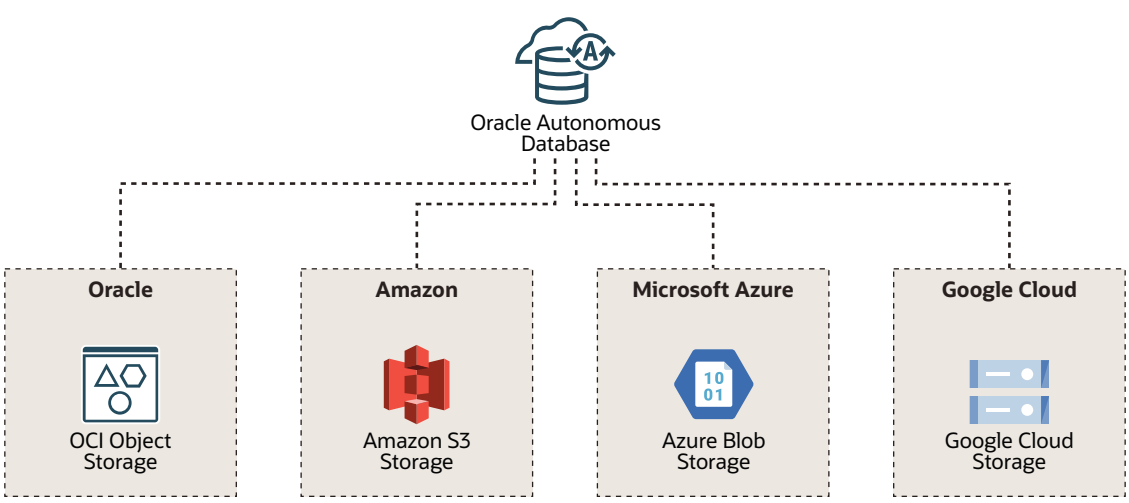
Multi-Cloud Support

For organizations that already have existing data lakes on other platforms, Autonomous Database integrates seamlessly, allowing businesses to benefit from Autonomous Database advanced features without disrupting their current setups.

Provide Autonomous Database access to your data lake by granting the necessary privileges and access for your data lake to be connected to Autonomous Database. Once you've provided the necessary credentials, Autonomous Database can seamlessly connect to data lakes across various cloud environments, including AWS, Azure, Google Cloud, and Oracle OCI object store.

This capability allows you to securely access and manage your data, leveraging the native security features of each cloud provider. With this multi-cloud support, you gain the flexibility to deploy and scale your data lake across different cloud platforms while maintaining a unified and secure environment.

Oracle Autonomous Database supports native security for other clouds, to learn more, see Use Amazon Resource Names (ARNs) to Access AWS Resources, Use Azure Service Principal to Access Azure Resources and Use Google Service Account to Access Google Cloud Platform Resources for your corresponding cloud platform.



End-To-End Data Format Support

Oracle Autonomous Database is designed with the flexibility to handle a broad spectrum of data formats, making it a universal solution for diverse data sources and workloads.

Whether your data resides in structured, semi-structured, or unstructured formats, Autonomous Database seamlessly supports them across various cloud environments. This allows businesses to ingest, store, and analyze data without worrying about format compatibility.

Autonomous Database provides native support for traditional formats like CSV and JSON, as well as advanced formats such as AVRO, Parquet, and ORC. See Query External Data with Autonomous Database, for more information. Autonomous Database supports the following file formats: CSV, JSON, XML, AVRO, ORC, Parquet, Delta Sharing, Iceberg, Word, PDF.

With the added support for the Iceberg Table format, Autonomous Database offers enhanced capabilities for large-scale data lake environments. Iceberg allows for optimized, high-performance querying, better version control, and easier data management, making it a good fit for large, evolving datasets. See Query Apache Iceberg Tables, for more information.

Enhanced Capabilities: Autonomous Database for Unstructured Data Management

While Oracle Database is recognized for its powerful processing of structured and semi-structured data, Autonomous Database extends its capabilities to handle unstructured datasets as well.

These capabilities include managing and analyzing a wide range of formats like JPG, PDF, Word documents, and more. With these advancements, Autonomous Database brings a comprehensive solution for businesses dealing with unstructured data sources.

- **AI-Driven Insights with Retrieval Augmented Generation (RAG)**: Autonomous Database integrates advanced AI models, enabling Vector Search for unstructured data. This allows for efficient retrieval of relevant information across massive datasets using AI, enhancing search accuracy and speed. See Select AI with Retrieval Augmented Generation (RAG), for more information.
- **Full-Text Indexing**: Autonomous Database supports the creation of full-text indexes on unstructured files, making it possible to perform advanced text searches on documents such as PDFs, Word files, and more. This capability greatly improves how unstructured content can be queried, indexed, and analyzed. See Use Full-Text Search on Files in Object Storage
- **Parse and Load Unstructured Data**: Autonomous Database enhanced parsing and data ingestion features allow users to load unstructured data seamlessly, automatically transforming it into a tabular format, ready to be load into database. See Perform Table Extraction from Image, for more information.
- **AI as a Source of Data (Prompt-to-Table)**: Leveraging AI, Autonomous Database enables prompt-to-table functionality, allowing users to generate data directly from AI models and load it into tables. This opens up possibilities for extracting valuable insights from AI-generated outputs and using them as a new source of structured data. See Loading Data from AI Source

These expanded capabilities position Autonomous Database as a powerful tool for handling the growing demands of unstructured data, while also tapping into AI-powered solutions, making it a versatile and future-proof platform for modern data challenges.

Flexible Metadata Management

Oracle Autonomous Database provides users with various ways to define metadata for their datasets, making data management more adaptable and efficient.

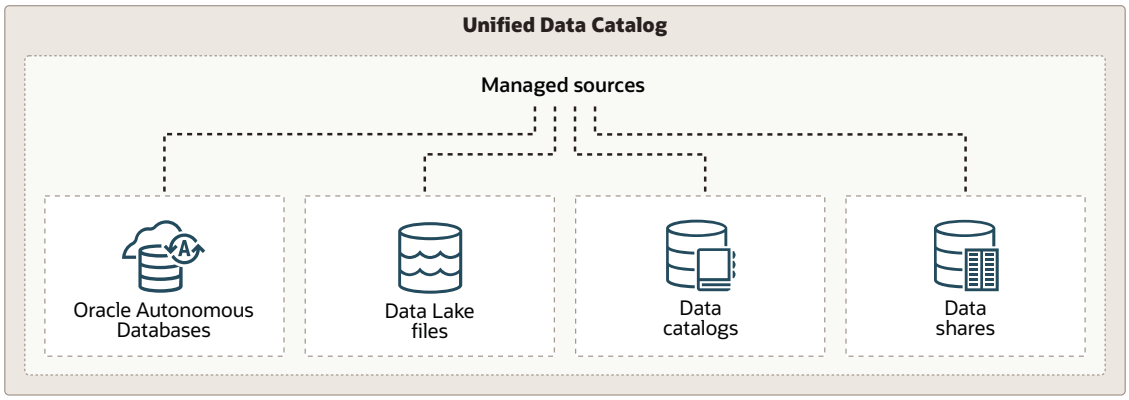
- **Catalog-Based Metadata Integration**
  Users can bring metadata from various catalogs into a centralized view, making it easier to control and maintain data consistency across the organization. Supported catalogs include:
  - **OCI Data Catalog**: A tool within Oracle Cloud Infrastructure (OCI) that helps users discover, organize, and manage data assets. It offers a clear view of all data assets, helping users maintain compliance, ensure data quality, and facilitate collaboration across teams. See Example: MovieStream Scenario , for more information.
  - **AWS Glue**: A managed ETL (extract, transform, load) service from Amazon Web Services that includes a data catalog for organizing and managing metadata. See Query External Data with AWS Glue Data Catalog, for more information.
- **Manual Metadata Definition**
  Users can also define metadata directly at the table level for datasets in object stores such as Oracle Cloud Infrastructure (OCI) Object Storage or Amazon S3. This allows for customized organization of data for individual files or groups of files, tailored to user

requirements. Autonomous Database can also automatically infer metadata, such as column names and data types, to save time and reduce errors. For example, when uploading a CSV file, the system can automatically detect headers as column names and assign appropriate data types like number or varchar2 based on the content. This helps users quickly prepare their data for analysis without manual intervention, reducing setup time and minimizing the chance of errors.

Federated Metadata Support

Autonomous Database supports a federated metadata catalog, enabling users to unify metadata from different sources into a single view, providing a unified interface for metadata management.

This approach simplifies metadata management across various environments by connecting data sources across multiple clouds and platforms. Whether using catalog-based metadata or defining it manually, all information is available in a unified catalog for easy browsing. For example, an organization can use this federated view to manage data assets from both AWS and Oracle Cloud, ensuring consistent governance and discoverability across platforms.



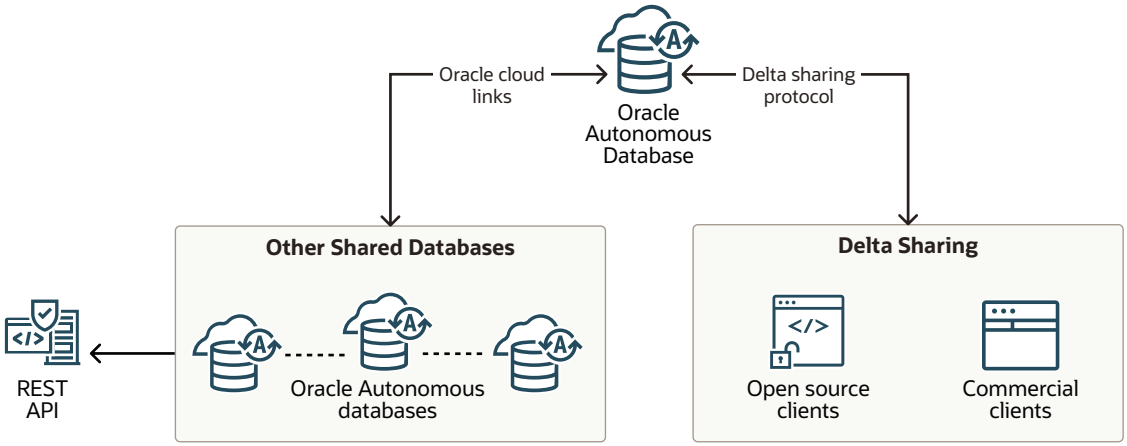
Collaboration

After users finish their analysis, they often need to share their results with others. Oracle Autonomous Database makes sharing easy by offering several ways to collaborate, providing unique advantages over other databases, such as integrated security features, open protocols, and seamless cloud connectivity.

These options are made to be flexible and secure, so they fit different collaboration needs:

- **Delta Sharing Protocol**: This lets you share data outside of Oracle using an open protocol called Delta Sharing. It supports secure data sharing with external partners, without needing complex integration, making it ideal for cross-cloud and cross-platform analytics. This way, data can be used smoothly in different analytics tools that are not part of Oracle. See Share Data Versions Using Object Storage, for more information.
- **Cloud Links**: You can share data between different Autonomous Database instances using secure cloud links. For example, Cloud Links are particularly effective for connecting different databases. This ensures consistent data availability and reduces latency for applications needing quick and reliable access to data across multiple databases, without the need to copy or duplicate. It keeps collaboration smooth for teams that are spread out and need to work together. See Share Live Data Using Direct Connection

- **Table Hyperlinks**: You can share data directly by creating special URLs that give access to the data without needing a separate login. Users can control the permissions and set expiration times for these URLs, ensuring secure and flexible sharing options. This feature is built specifically for REST clients. See Generate a Table Hyperlink for a Table or a View, for more information.



Broad Compatibility with Oracle Database Tools

The Autonomous Database environment is fully compatible with a wide array of Oracle database tools.

Any tool you already use to interact with Oracle databases—whether for data visualization, analytics, ETL, or administration—can also be leveraged seamlessly to analyze datasets within Autonomous Database . This compatibility ensures a frictionless experience, allowing users to integrate Autonomous Database into their existing workflows without needing to adopt new tools or processes, thereby maximizing efficiency and reducing the learning curve.

See The Data Studio Overview Page, for information on a few of the tools available to use with Oracle databases.

The Data Studio Overview Page

The Data Studio comprises of the Data Load, the Data Analysis, the Data Insights, Catalog, Data Marketplace and the Data Share tool.

The Data Studio Tools enables you to load data from cloud and other diverse sources, analyzes it and gain insights from it. You can share the result of the analysis with other users. It is a one-stop application of your analytics tool from multiple data sources. This tool makes sure that there is seamless transition between different applications. The multiple ways of navigation do not impact the progress of your work. For instance, if you are working on data analysis and decide you need some additional data, you can navigate to the Data Load page, bring in the new contents, and return to your analysis in progress. Select **Data Studio** tab from the Launchpad.

Data Studio Overview page

Click the **Data Studio Overview** menu from the **Data Studio** to view the Data Studio Overview page consists of four sections:

1. A navigation pane which consists of the following tools in the menu: Data Load, Data Analysis, Data Insights, the Catalog and the Data Share tools.
2. A widget with slides which displays information on the tools and their purpose. For example, the Data Load tool loads data from files on your computer or from cloud storage. Click **<** or **>** button to navigate back and forth the slides.
3. The right pane toolbar displays the help links to the Data Studio tool documentation.
4. A Recent Objects section which displays recently updated or created objects. Each entry has a context menu available at the end of the row. Click on the entries to view the details of the selected object.

Load your data from diverse sources such as files, databases and cloud storage, which are all consolidated in a single location. Some of the cloud data warehouses the Data Studio tool supports are Oracle Cloud Infrastructure (OCI), Amazon S3, Microsoft Azure Blob Storage and Google Cloud Storage.

Once your data is available, analyze it, generate insights, and create reports.

On the right of the Data Studio Overview page, refer to the Getting Started panel to know more about this tool.

> **Note:**
>
> You will lose your data once you click the Database Actions in the header or click the selector icon.

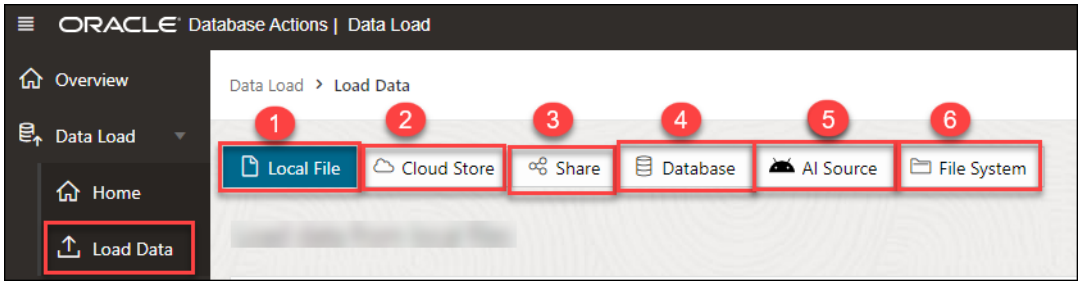
Let us run analysis with data from different sources and generate insights from it.

Load Data

Select the **Data Load** menu from the navigation pane in the Data Studio Overview page to load data from files on your local device, from remote databases, or from cloud storage buckets. Clicking Data Load opens Data Load Dashboard page.

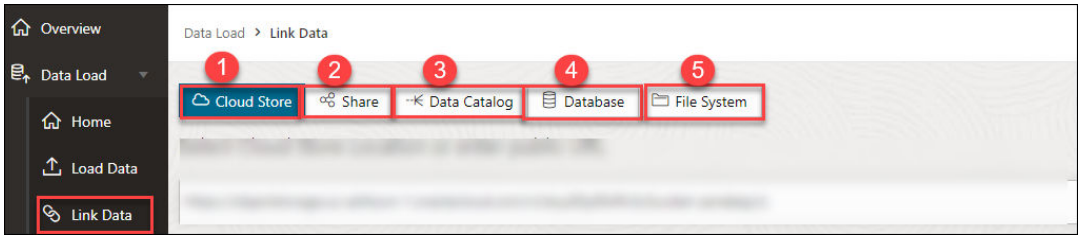
The Data Load menu has the following submenus:

- **Home**: This menu takes you to the home page of the Data Load tool in the Launchpad. You can perform any of the following activities from the submenus such as Load local file, Load cloud store, Load Database Tables, Link Cloud Store, Link Database Tables, Create Live feed and search for cloud location and then select the Home menu. This does not break the continuity of the action you perform. Refer to The Data Load Page for more details.
- **Load Data**:
  Click **Load Data** to view the following options to load data:



  - **Load Local File**: This option enables us to load data from files on your local device. Refer to the Loading Data From Local Files section for more details.
  - **Load Cloud Store**: This menu enables you to load data from a cloud store location to a table in your Autonomous Database. Refer to the Loading Data from Cloud Storage section for more details.
  - **Load Share**: This menu enables you to load data from Shares where you can select tables from a Share. See Loading Data from a Share.
  - **Load Database**: This menu enables you to load data from tables in another database into your Oracle Autonomous Database. Refer to the Loading Data from other Databases section for more details.
  - **Load AI Source**: This menu enables you to load data from an AI source. See Loading Data from AI.
  - **Load File System**: This menu enables you to load data from file system directories to your Autonomous Database.

- **Link Data**:
  Click **Link Data** to view the following options of linking data:



  - **Link Cloud Store**: This menu enables you to link to data in cloud storage buckets. Refer to the Linking to Objects in Cloud Storage section for more details.
  - **Link Share**: This menu enables you to link a share with your Autonomous Database. See Linking to Share.
  - **Linking Data Catalog** : This menu enables you to link Data Catalog to your Autonomous Database. See Linking Data Catalog.
  - **Link Database**: This menu enables you to link to data in tables in another database from your Oracle Autonomous Database. Refer to the Linking to Other Databases section for more details.
  - **Linking File System**: This menu enables you to link file system directories to your Autonomous Database. See Linking to File System.

- **Live Feed**: This menu enables us to create a live table feed on demand, on a schedule, or as the result of a notification. Refer to the Feeding Data section for more details.
- **Connections**: You can view all the existing connections from this option. This option enables you to edit or delete a cloud storage link and manage catalogs and shares. You can also create new cloud storage links. Refer to the Managing Connections chapter for more details.

After you load the data using the data load job, it becomes available in the Data Load page. You can later inspect, review and delete the table on clicking the **Actions** icon. You could also reload the cart from the Reload Cart button to reload cart using the current cart in case you want to make changes to the table. This data will be the source data of the Analytic View. Use breadcrumbs at the top of the page to navigate back to the home page.

Follow the steps described in The Data Load Page chapter in this book to load data from Oracle Cloud Infrastructure to your Autonomous Database.

Data Analysis

Analytic views provide metadata that highlight the parts of your data that are most important to your organization. By creating an AV, you can improve the value of automatic insights and the analysis of your data because the system has this metadata.

Select the **Data Analysis** menu from the navigation pane in the Data Studio Overview page to create Analytic Views and analyze the data from the Analytic View. The table you loaded using the Load Data feature will be used to create the Analytic View. You can select both hierarchies and measures from Analytic Views. Analytic Views utilize diagrams, symbols, and text, that represents how the data flows and the connection in between them. The Data Analysis tool helps you find the structural errors in the Analytic View you create. It spots anything from null values in a join column to duplicated fields. You can develop a visual representation of your analysis in the form of tables, Pivot tables and bar charts. You can also export the Analytic View to Tableau or Power BI for better visualization.

Create Analytic Views by referring to The Data Analysis Tool chapter in this book.

The **Related Insights** panel in the right works with the measures associated with the Analytic View to visualize data in the form of bar charts.

Table AI Assist

Select the **Table AI Assist** menu from the navigation pane in the Data Studio Overview page to effortlessly enhance your local data tables with the power of AI. You can use natural language prompts to augment and correct your information with this tool.

Refer to The Table AI Assist Tool chapter of this book for more information on the Table AI Assist page.

Automatic Insights

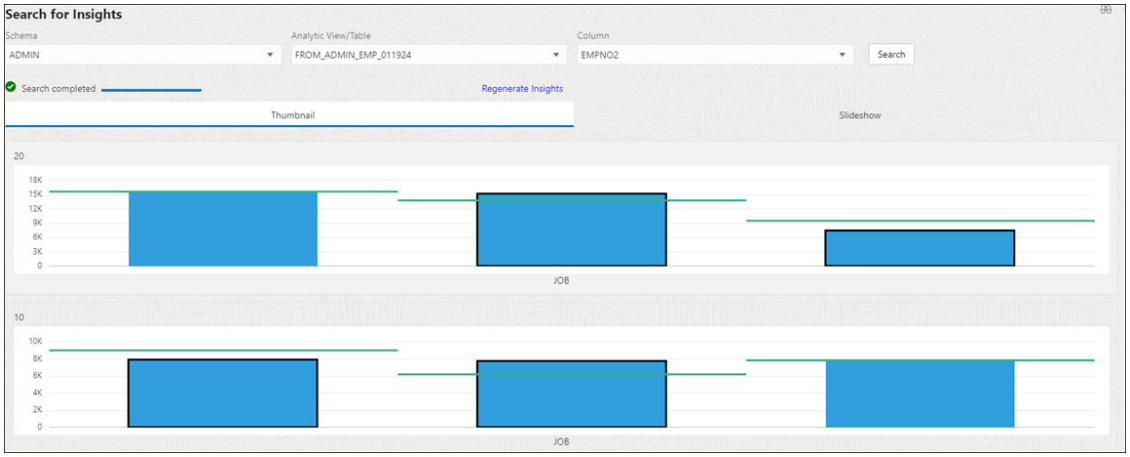
Select the **Automatic Insights menu** from the navigation pane in the Data Studio Overview page to generate insights automatically. The insights appear in the Data Insights dashboard as a series of bar charts.

By selecting the user, Analytic View, and the column name, you can view the data points of the actual values that deviate considerably from the forecast values. This type of predictive analytics will enable you to understand a pattern better and make better decisions. Then you can act on your data to optimize it or to improve its performance.

For example, consider that the fact table for the insights records values about different companies, and the measures of the fact table are Value (USD), Acquisition date and Acquisition year. Then you can generate insights for the Value measure. The dashboard will have a series of charts labeled Acquisition date and Acquisition year. Let’s say you want to view the value of your company **A**. Click the chart whose sale you want to display for a particular range of Acquisition year. Clicking a chart displays more details about it.

Generate automatic insights by referring to The Data Insights Page in this book.

Below is the chart that displays the values and insights.



According to the insights generated by the tool, you notice that your company is lagging from the expected value by a significant number. In order to lead you need to change a few factors that affect the sales of your organization.

The insights that Data Insights generates for the analytic view are more useful than those for a table because of the additional metadata that an analytic view provides.

Catalog

Select the **Catalog** menu from the navigation pane in the Data Studio Overview page to view information about the upstream dependencies of the entity, how the entity was created and how it is linked to other entities.

This page lists details about the database objects such as cloud storage links and tables you create from the Data Analysis tool and the Data load tool. It also displays lists of the database dictionary objects such as tables, columns, and Analytic Views.

Use the Catalog page to locate objects in the catalog and perform tasks specific to those objects.

Refer to The Catalog Tool chapter of this book for more information on the Catalog page.

Data Marketplace

Select the **Marketplace** menu from the navigation pane in the Data Studio Overview page to easily exchange metadata from data providers to data consumers, allowing easy data discovery and use.

Refer to The Data Marketplace Tool chapter of this book for more information on the Data Marketplace.

Data Share

Select the **Data Share** menu from the navigation pane in the Data Studio Overview page to provide data shares and consume data shares.

See The Data Share Tool

High Availability

You can configure and manage Autonomous Data Guard and maintain continuous availability using disaster recovery.

- Use Application Continuity on Autonomous Database
  Autonomous Database provides application continuity features for making connections to the database.
- Use Standby Databases with Autonomous Data Guard for Disaster Recovery
- Backup and Restore Autonomous Database Instances
  This section describes backup and recovery tasks on Autonomous Database.
- Use OraMTS Recovery Feature on Autonomous Database
  Use Oracle MTS (OraMTS) Recovery Service to resolve an in-doubt Microsoft Transaction Server transaction.
- Use Backup-Based Disaster Recovery
  Backup-Based Disaster Recovery provides a lower-cost disaster recovery option for Autonomous Database (the RTO is higher with this option, as compared to using Autonomous Data Guard).
- Track Table Changes with Flashback Time Travel
  Use Flashback Time Travel to view past states of database objects or to return database objects to a previous state without using point-in-time media recovery.
- Use OCI Full Stack Disaster Recovery with Autonomous Database
  You can enable the Oracle Cloud Infrastructure Full Stack Disaster Recovery on Autonomous Database.

Use Application Continuity on Autonomous Database

Autonomous Database provides application continuity features for making connections to the database.

- About Application Continuity on Autonomous Database
  Application Continuity masks outages from end users and applications by recovering the in-flight work for impacted database sessions following outages. Application Continuity performs this recovery beneath the application so that the outage appears to the application as a slightly delayed execution.
- Configure Application Continuity on Autonomous Database
  To configure Application Continuity you must enable application continuity for the database service your application uses and configure the failover type and the drain timeout. In addition, you must set several connection string parameters that enable high availability.
- Client Configuration for Continuous Availability on Autonomous Database
  You do not need to restart applications for planned maintenance activities when you enable Application Continuity and you follow the coding best practices.
- Notes for Application Continuity on Autonomous Database
  There are restrictions for Application Continuity on Autonomous Database, as follows:

About Application Continuity on Autonomous Database

Application Continuity masks outages from end users and applications by recovering the in-flight work for impacted database sessions following outages. Application Continuity performs this recovery beneath the application so that the outage appears to the application as a slightly delayed execution.

Your applications achieve *continuous availability* when planned maintenance, unplanned outages, and load rebalances of the database are hidden from the application. The combination of application coding best practices, application continuity configuration, and Autonomous Database ensures that your applications are continuously available.

The best approach for hiding planned maintenance activities from your applications is to transparently drain or failover applications. Oracle's connection pools and mid-tiers, including the WebLogic Server, Oracle Universal Connection Pool (UCP), OCI session pool and ODP.NET Unmanaged Provider are Fast Application Notification (FAN) aware and therefore are notified when maintenance is underway on Autonomous Database to allow graceful draining of work before maintenance. Application Continuity runs during planned maintenance to failover those sessions that do not drain in the predefined drain interval (5 minutes on Autonomous Database).

In order to hide unplanned outages resulting from a component or communication failure Oracle provides:

- **Notification.** FAN is the first step to hiding outages. FAN notifies clients and breaks them out of their current network wait when an outage occurs. This avoids stalling applications for long network waits. For Autonomous Database, FAN is handled at the driver and by the Autonomous Database cloud connection manger.

  FAN notification automatically triggers closing idle connections, opening new connections in the new service location, and allows a configurable time for active work to complete in the soon-to-be-shutdown service location. The major third-party JDBC mid-tiers, such as IBM WebSphere, allow for the same behavior when configured with UCP. For JDBC-based applications that cannot use UCP, Oracle provides solutions using Oracle Drivers and connection tests. On Autonomous Database FAN for planned maintenance is sent in-Band.

- **Recovery.** After the client is notified, failover handling with Transparent Application Continuity (TAC) or Application Continuity (AC) re-establishes a connection to the Autonomous Database and replays in-flight, uncommitted, work when possible. By replaying in-flight work, the application can usually continue executing without knowing that any failure happened.

You enable Application Continuity on Autonomous Database in one of two configurations, depending on the application:

- **Application Continuity (AC)**

  Application Continuity hides outages for thin Java-based applications, and Oracle Database Oracle Call Interface and ODP.NET based applications with support for open-source drivers, such as Node.js, and Python. Application Continuity rebuilds the session by recovering the session from a known point which includes session states and transactional states. Application Continuity rebuilds all in-flight work. The application continues as it was, seeing a slightly delayed execution time when a failover occurs.

- **Transparent Application Continuity (TAC)**

  Transparent Application Continuity (TAC) transparently tracks and records session and transactional state so the database session can be recovered following recoverable outages. This is done with no reliance on application knowledge or application code changes, allowing Transparent Application Continuity to be enabled for your applications. Application transparency and failover are achieved by consuming the state-tracking information that captures and categorizes the session state usage as the application issues user calls.

See Overview of Application Continuity for more information on Application Continuity.

> **Note:**
>
> By default Application Continuity is disabled on Autonomous Database.

Configure Application Continuity on Autonomous Database

To configure Application Continuity you must enable application continuity for the database service your application uses and configure the failover type and the drain timeout. In addition, you must set several connection string parameters that enable high availability.

- Configure Your Service to Enable Application Continuity
  Use `DBMS_APP_CONT_ADMIN` to enable Application Continuity or Transparent Application Continuity:
- Use Fast Application Notification (FAN)
  When connecting to Autonomous Database, the Oracle database auto-configures FAN. For application deployments with Autonomous Database, Fast Application Notification (FAN) events for unplanned outages are directed to the connection manager (CMAN) and no client application configuration steps required to use FAN.
- Configure Connection String for High Availability
  To maintain high availability, Oracle recommends that you set certain connection string parameters when you connect to Oracle Autonomous Database.
- Configure Driver Specific Client Options
  Depending on your client and your driver, you need to assure that the client is properly configured to use Application Continuity when you connect to Autonomous Database

Configure Your Service to Enable Application Continuity

Use `DBMS_APP_CONT_ADMIN` to enable Application Continuity or Transparent Application Continuity:

- **Application Continuity (AC)**: Set this failover option using the procedure `DBMS_APP_CONT_ADMIN.ENABLE_AC`. The `ENABLE_AC` procedure takes three parameters: `SERVICE NAME` is the service name to change, `FAILOVER_RESTORE`, set to `LEVEL1` to select Application Continuity(AC), and `REPLAY_INITIATION_TIMEOUT`, is the replay timeout that specifies how many seconds after a request is submitted to allow that request to replay.

  For example, as the ADMIN user, to enable Application Continuity for the `TPURGENT` service:

  ```
  execute DBMS_APP_CONT_ADMIN.ENABLE_AC(
          'databaseid_tpurgent.adb.oraclecloud.com', 'LEVEL1', 600);
  ```

- **Transparent Application Continuity (TAC)**: Set this failover option using the procedure `DBMS_APP_CONT_ADMIN.ENABLE_TAC`. The `ENABLE_TAC` procedure takes three parameters: `SERVICE NAME` is the service name to change, `FAILOVER_RESTORE`, set to `AUTO` to select Transparent Application Continuity (TAC), and `REPLAY_INITIATION_TIMEOUT` is the replay timeout that specifies how many seconds after a request is submitted to allow that request to replay.

  For example, as the ADMIN user, to enable Transparent Application Continuity for the `TP` service with the replay timeout set to 20 minutes:

  ```
  execute DBMS_APP_CONT_ADMIN.ENABLE_TAC(
          'databaseid_tp.adb.oraclecloud.com', 'AUTO', 1200);
  ```

- **Disabled**: Disable failover using the procedure `DBMS_APP_CONT_ADMIN.DISABLE_FAILOVER()`.

For example, as the ADMIN user, to disable failover for the `TP` service:

```
execute DBMS_APP_CONT_ADMIN.DISABLE_FAILOVER(
        'databaseid_tp.adb.oraclecloud.com');
```

Find the Service Name Parameter for Application Continuity

Depending on your workload type, use a command similar to the following to `SELECT` from `DBA_SERVICES` on your database and identify the service where you want to enable Application Continuity:

- **Data Warehouse**

```
SELECT name, failover_type FROM DBA_SERVICES;
NAME                                                     FAILOVER_TYPE
-------------------------------------------------------- -------------
nvt21_adb1_low.adb.oraclecloud.com
nvt21_adb1_high.adb.oraclecloud.com
nvt21_adb1_medium.adb.oraclecloud.com
```

- **Transaction Processing** or **JSON Database**

```
SELECT name, failover_type FROM DBA_SERVICES;
NAME                                                     FAILOVER_TYPE
-------------------------------------------------------- ----------------
nvt21_adb1_tp.adb.oraclecloud.com
nvt21_adb1_tpurgent.adb.oraclecloud.com
nvt21_adb1_low.adb.oraclecloud.com
nvt21_adb1_high.adb.oraclecloud.com
nvt21_adb1_medium.adb.oraclecloud.com
```

Notice the `FAILOVER_TYPE` for the `high` service has the no value and indicates that Application Continuity is disabled.

Verify Application Continuity is Enabled for a Service

Depending on your workload type, check the output of the query on `DBA_SERVICES` to verify that Application Continuity is enabled.

- **Data Warehouse**

```
SELECT name, failover_type FROM DBA_SERVICES;
NAME                                                     FAILOVER_TYPE
-------------------------------------------------------- -------------
nvt21_adb1_low.adb.oraclecloud.com
nvt21_adb1_high.adb.oraclecloud.com                      AUTO
nvt21_adb1_medium.adb.oraclecloud.com
```

- **Transaction Processing** or **JSON Database**

```
SELECT name, failover_type FROM DBA_SERVICES;
NAME                                                     FAILOVER_TYPE
-------------------------------------------------------- -------------
nvt21_adb1_tp.adb.oraclecloud.com
nvt21_adb1_tpurgent.adb.oraclecloud.com                  TRANSACTION
nvt21_adb1_low.adb.oraclecloud.com
```

```
nvt21_adb1_high.adb.oraclecloud.com                    AUTO
nvt21_adb1_medium.adb.oraclecloud.com
```

The `FAILOVER_TYPE` value for the `high` service is now `AUTO`, indicating that Transparent Application Continuity (TAC) is enabled and the `FAILOVER_TYPE` value for the `tpurgent` service is now `TRANSACTION`, indicating that Application Continuity (AC) is enabled.

Use Fast Application Notification (FAN)

When connecting to Autonomous Database, the Oracle database auto-configures FAN. For application deployments with Autonomous Database, Fast Application Notification (FAN) events for unplanned outages are directed to the connection manager (CMAN) and no client application configuration steps required to use FAN.

FAN is automatically handled for you by the client driver and by the Autonomous Database Connection Manager (CMAN):

- For planned maintenance events, FAN is sent in-band, directly to the drivers. This requires that applications use Oracle Pools or TAC for request boundaries, or use connection tests.
- The Oracle Database and Oracle client drivers drain on connection tests and at request boundaries.

See Client Configuration for Continuous Availability on Autonomous Database for more information.

Configure Connection String for High Availability

To maintain high availability, Oracle recommends that you set certain connection string parameters when you connect to Oracle Autonomous Database.

Set the `CONNECT_TIMEOUT`, `RETRY_DELAY`, `RETRY_COUNT`, and `TRANSPORT_CONNECT_TIMEOUT` parameters in the connection string when you connect to Oracle Autonomous Database. The connect strings embedded in the Oracle-supplied `tnsnames.ora` file are preconfigured with appropriate values for most applications. In some cases, depending on your applications needs, you may need to change the preconfigured values for a connection string.

Use this TNS for all Oracle clients version 12.2 or higher:

```
alias =
(DESCRIPTION =
(CONNECT_TIMEOUT= 90)(RETRY_COUNT=50)
(RETRY_DELAY=3)(TRANSPORT_CONNECT_TIMEOUT=3)
(ADDRESS_LIST =
  (LOAD_BALANCE=on)
  (ADDRESS = (PROTOCOL = TCP)(HOST=scan-host)(PORT=1521)))
(CONNECT_DATA=(SERVICE_NAME = service-name)))
```

Use the following for JDBC connections using Oracle driver version 12.1 or earlier:

```
alias =
(DESCRIPTION =
(CONNECT_TIMEOUT= 15)(RETRY_COUNT=50)
(RETRY_DELAY=3)
(ADDRESS_LIST =
   (LOAD_BALANCE=on)
   (ADDRESS = (PROTOCOL = TCP)(HOST=scan-host)(PORT=1521)))
 (CONNECT_DATA=(SERVICE_NAME = service-name)))
```

Notes for connection strings:

- For JDBC and ODP clients, the pool connection wait time should be configured to be longer than the `CONNECT_TIMEOUT` in the connect string.
- Do not use Easy Connect Naming on the client because such connections do not have high-availability capabilities.

See Download Client Credentials (Wallets) for information on the `tnsnames.ora` file.

Configure Driver Specific Client Options

Depending on your client and your driver, you need to assure that the client is properly configured to use Application Continuity when you connect to Autonomous Database

- Configure JDBC Thin Driver
  Shows details to use Application Continuity with Autonomous Database with a client using the JDBC Thin Driver.
- Configure Oracle Call Interface (OCI) Driver
  Shows details to use Application Continuity with Autonomous Database with a client using the Oracle Call Interface (OCI) driver.
- Configure ODP.NET Unmanaged Provider Driver
  Shows details to use Application Continuity with Autonomous Database with a client using the ODP.NET Unmanaged Provider Driver.

Configure JDBC Thin Driver

Shows details to use Application Continuity with Autonomous Database with a client using the JDBC Thin Driver.

If your application uses the JDBC Thin Driver, follow these recommended practices:

1. Use JDBC Statement Cache for Coverage and Performance.

   For best coverage and performance, use the JDBC driver statement cache in place of an application server statement cache. This allows the driver to know that statements are closed and memory is to be freed at the end of requests.

   To use the JDBC statement cache, use connection property `oracle.jdbc.implicitStatementCacheSize (OracleConnection.CONNECTION_PROPERTY_IMPLICIT_STATEMENT_CACHE_SIZE)`. The statement cache is per connection. The value for the cache size matches your number of `open_cursors`. For example:

   `oracle.jdbc.implicitStatementCacheSize=`*`nnn`* where *`nnn`* is typically between 10 and 100 and is equal to the number of open cursors your application maintains.

2. Tune the Garbage Collector.

   For many applications the default Garbage Collector tuning is sufficient. For applications that return and keep large amounts of data you can use higher values, such as 2GB or larger. For example:

   ```
   java -Xms3072m -Xmx3072m
   ```

   It is recommended to set the memory allocation for the initial Java heap size (ms) and maximum heap size (mx) to the same value. This prevents using system resources on growing and shrinking the memory heap.

3. When using the Universal Connection Pool (UCP), disable Fast Connection Failover. For example:

```
PoolDataSource.setFastConnectionFailoverEnabled(false)
```

Configure Oracle Call Interface (OCI) Driver

Shows details to use Application Continuity with Autonomous Database with a client using the Oracle Call Interface (OCI) driver.

If the client application uses the Oracle Call Interface (OCI) Driver, follow this recommended practice:

- Replace `OCIStmtPrepare` with `OCIStmtPrepare2`. `OCIStmtPrepare()` has been deprecated since 12.2. All applications should use `OCIStmtPrepare2()`. Transparent Application Continuity (TAC) and Application Continuity (AC) allow `OCIStmtPrepare` but do not replay this statement.

Do not configure ONS servers in `oraaccess.xml`:

```
<ons>
   <servers>
   <!—Do not enter any values -->
   </servers>
</ons>
```

Also, for Autonomous Database Serverless, do not configure the <fan> section:

```
<fan>
<!-- only possible values are "trace" or "error" -->
    <subscription_failure_action>
   </subscription_failure_action>
</fan>
```

Configure ODP.NET Unmanaged Provider Driver

Shows details to use Application Continuity with Autonomous Database with a client using the ODP.NET Unmanaged Provider Driver.

The ODP.NET Unmanaged Provider driver automatically uses Application Continuity when Application Continuity is enabled on the database service that your application uses to connect to Autonomous Database.

When connecting an ODP.NET application to your Autonomous Database Serverless, do not configure ONS servers in `oraaccess.xml`:

```
<ons>
   <servers>
   <!—Do not enter any values -->
   </servers>
</ons>
```

Client Configuration for Continuous Availability on Autonomous Database

You do not need to restart applications for planned maintenance activities when you enable Application Continuity and you follow the coding best practices.

- Connect Using Database Services with Application Continuity Enabled
- Use Recommended Practices That Support Draining
  On Autonomous Database there is never a need to restart application servers when planned maintenance follows best practice.
- Steps for Using Application Continuity
  Perform these steps to use Application Continuity:
- Developer Best Practices for Continuous Availability
  Follow these best practices to code for continuous availability on Autonomous Database.

Connect Using Database Services with Application Continuity Enabled

Oracle database services provide transparency for the underlying Autonomous Database infrastructure.

The high availability and application continuity operations are predicated on the use of Autonomous Database connection services. To obtain application continuity, use a database service when you connect to your database.

The names of the predefined database services on Autonomous Database are different, depending on your workload, as described in Database Service Names for Autonomous Database.

Use Recommended Practices That Support Draining

On Autonomous Database there is never a need to restart application servers when planned maintenance follows best practice.

For planned maintenance, the recommended approach is to provide time for current work to complete before maintenance is started. On Autonomous Database this happens automatically and work is drained before starting maintenance activities when you follow these guidelines:

- FAN with Oracle Connection Pools or Oracle Drivers
- Connection tests

Use draining in combination with your chosen failover solution for those requests that do not complete within the allocated time for draining. Your failover solution will try to recover sessions that did not drain in the allocated time.

Return Connections to the Connection Pool

The application should return the connection to the connection pool on each request. It is best practice that an application checks-out a connection only for the time that it needs it. Holding a connection instead of returning it to the pool does not perform. An application should therefore check-out a connection and then check-in that connection immediately the work is complete. The connections are then available for later use by other threads, or your thread when needed again. Returning connections to a connection pool is a general recommendation regardless of whether you use FAN to drain, or connection tests to drain.

Use an Oracle Connection Pool

Using a FAN-aware, Oracle connection pool is the recommended solution for hiding planned maintenance. As the maintenance progresses and completes, sessions are moved and rebalanced. There is no impact to users when your application uses an Oracle Pool with FAN and returns connections to the pool between requests. Supported Oracle Pools include UCP, WebLogic GridLink, Tuxedo, OCI Session Pool, and ODP.NET Managed and Unmanaged providers. No application changes whatsoever are needed to use FAN other than making sure that your connections are returned to pool between requests.

Use UCP with a Third-Party Connection Pool

If you are using a third party, Java-based application server, the most effective method to achieve draining and failover is to replace the pooled data source with UCP. This approach is supported by many application servers including Oracle WebLogic Server, IBM WebSphere, IBM Liberty, Apache Tomcat, Red Hat WildFly (JBoss), Spring, and Hibernate, and others.

Use Connection Tests

If you cannot use an Oracle Pool with FAN, then the Autonomous Database or provided client drivers will drain the session. When services are relocated or stopped during maintenance, or there is a switchover to a standby site using Autonomous Data Guard, the Oracle Database and Oracle client drivers look for safe places to release connections according to the following rules:

- Standard connection tests for connection validity at borrow or return from a connection pool
- Custom SQL tests for connection validity
- Request boundaries are in effect and the current request has ended
- Use Connection Tests with Autonomous Database
  You can add, delete, enable or disable connection tests for Autonomous Database.
- Use Connection Tests with Thin Java Driver
- Use Connection Tests with OCI Driver

Use Connection Tests with Autonomous Database

You can add, delete, enable or disable connection tests for Autonomous Database.

Use the view `DBA_CONNECTION_TESTS` to show the available connection tests.

For example:

```
SQL> EXECUTE
   dbms_app_cont_admin.add_sql_connection_test('SELECT COUNT(1) FROM DUAL');
SQL> EXECUTE
   dbms_app_cont_admin.enable_connection_test(dbms_app_cont_admin.sql_test,
                                              'SELECT COUNT(1) FROM DUAL');
SQL> SELECT * FROM DBA_CONNECTION_TESTS;
```

Configure the same connection test that is enabled in your database at your connection pool or application server. Also configure flushing and destroying the pool on connection test failure to at least two times the maximum pool size or `MAXINT`.

Use Connection Tests with Thin Java Driver

If you would like to use connection tests that are local to the driver and cannot use UCP's full FAN support:

- Enable `validate-on-borrow=true`
- Set the Java system properties:
  - `-Doracle.jdbc.fanEnabled=true`
  - `-Doracle.jdbc.defaultConnectionValidation=SOCKET`

And then use one of the following tests:

- `java.sql.Connection.isValid(int timeout)`
- `oracle.jdbc.OracleConnection.pingDatabase()`
- `oracle.jdbc.OracleConnection.pingDatabase(int timeout)`
- A `HINT` at the start of your test SQL:
  - `/*+ CLIENT_CONNECTION_VALIDATION */`

Use Connection Tests with OCI Driver

If you would like to use the OCI driver directly, use `OCI_ATTR_SERVER_STATUS`. This is the only method that is a code change. In your code, check the server handle when borrowing and returning connections to see if the session is disconnected. During maintenance, the value of `OCI_ATTR_SERVER_STATUS` is set to `OCI_SERVER_NOT_CONNECTED`. When using OCI session pool, this connection check is done for you.

The following code sample shows how to use `OCI_ATTR_SERVER_STATUS`:

```
ub4 serverStatus = 0OCIAttrGet((dvoid *)srvhp,
          OCI_HTYPE_SERVER,
        (dvoid *)&serverStatus, (ub4 *)0, OCI_ATTR_SERVER_STATUS,
      errhp);if (serverStatus ==
        OCI_SERVER_NORMAL)printf("Connection is
        up.\n");else if (serverStatus ==
        OCI_SERVER_NOT_CONNECTED) printf("Connection is down.\n");
```

Steps for Using Application Continuity

Perform these steps to use Application Continuity:

- As a prerequisite, enable and configure Application Continuity or Transparent Application Continuity (TAC) for your database service on Autonomous Database. See Configure Application Continuity on Autonomous Database for more information.
- Oracle strongly recommends that you use the latest client drivers. Oracle Database 19c client drivers and later provide full support for Application Continuity (AC) and for Transparent Application Continuity (TAC). Use one of the following supported clients drivers:
  - Oracle JDBC Replay Driver 19c or later. This is a JDBC driver feature provided with Oracle Database 19c for Application Continuity

- Oracle Universal Connection Pool (UCP) 19c or later with Oracle JDBC Replay Driver 19c or later
- Oracle Weblogic Server 12c with Active GridLink, or third-party JDBC application servers using UCP with Oracle JDBC Replay Driver 19c or later
- Java connection pools or standalone Java applications using Oracle JDBC Replay Driver 19c or later
- Oracle Call Interface Session Pool 19c or later.SQL*Plus 19c (19.8) or later
- ODP.NET pooled, Unmanaged Driver 19c or later ("Pooling=true" default in 12.2 and later)
- Oracle Call Interface based applications using 19c OCI driver or later

Return Connections to the Connection Pool

The application should return the connection to the Oracle connection pool on each request. Best practice for application usage is to check-out (borrow) connections for only the time that they are needed, and then check-in to the pool when complete for the current actions. This is important for best application performance at runtime, for rebalancing work at runtime and during maintenance and failover events. This practice is also important for draining.

When using an Oracle connection pool, such as Universal Connection Pool (UCP) or OCI Session Pool, or ODP.Net Unmanaged Provider or when using WebLogic Active GridLink, following this practice embeds request boundaries that Application Continuity uses to identify safe places to resume and end capture. This is required for Application Continuity and is recommended for Transparent Application Continuity.

Transparent Application Continuity, in addition, will discover request boundaries if a pool is not in use or when replay is disabled. The conditions for discovering a boundary are:

- No open transaction
- Cursors are returned to the statement cache or cancelled
- No un-restorable session state exists (refer to Clean Session State between Requests in this paper)

Enable Mutables Used in the Application

Mutable functions are functions that can return a new value each time they are executed. Support for keeping the original results of mutable functions is provided for `SYSDATE`, `SYSTIMESTAMP`, `SYS_GUID`, and sequence.`NEXTVAL`. If the original values are not kept and different values are returned to the application at replay, replay is rejected.

If you need mutables for PL/SQL, issue `GRANT KEEP` as required.

For example:

```
SQL> GRANT KEEP DATE TIME to adb_user;
SQL> GRANT KEEP SYSGUID to adb_user;
SQL> GRANT KEEP SEQUENCE mySequence to adb_user on mysequence.myobject;
```

Side Effects

When a database request includes an external call such as sending MAIL or transferring a file then this is termed a side effect.

Side effects are external actions, they do not roll back. When replay occurs, there is a choice as to whether side effects should be replayed. Many applications choose to repeat side effects

such as journal entries and sending mail as duplicate executions cause no problem. For Application Continuity side effects are replayed unless the request or user call is explicitly disabled for replay. Conversely, as Transparent Application Continuity is on by default, TAC does not replay side effects. The capture is disabled, and re-enables at the next implicit boundary created by TAC.

Developer Best Practices for Continuous Availability

Follow these best practices to code for continuous availability on Autonomous Database.

Return Connections to the Connection Pool

The most important developer practice is to return connections to the connection pool at the end of each request. This is important for best application performance at runtime, for draining work and for rebalancing work at runtime and during maintenance, and for handing failover events. Some applications have a false idea that holding onto connections improves performance. Holding a connection neither performs nor scales.

Clean Session State between Requests

It is best practice to clean session state between database requests.

When an application returns a connection to the connection pool, cursors in FETCH status, and session state set on that session remain in place unless an action is taken to clear them. If your application is setting state, it is best practice to return your cursors to the statement cache and to clear application related session state to prevent leakage to later re-uses of that database session. Cleaning your session state ensures that TAC can discover boundaries.

To automatically clean your state between requests with Oracle Database 23ai, set the service attribute `RESET_STATE=LEVEL1`. Doing this will avoid state leakage and fetching from cursors by later usage of the connection pool.

If you are using Oracle Database 19c, use `DBMS_SESSION.RESET_PACKAGE` to clear PL/SQL global variables, use `TRUNCATE` to clear temporary tables, `SYS_CONTEXT.CLEAR_CONTEXT` to clear context and cancel your cursors by returning them to the statement cache.

If your application is stateless, such as REST, APEX, Microservice, and most web applications, it is best practice to use `RESET_STATE`.

Do not embed COMMIT in PL/SQL and Avoid Commit on Success and Autocommit

It is recommended practice to use a top-level commit, (`OCOMMIT` or `COMMIT()` or `OCITransCommit`). If your application is using `COMMIT` embedded in PL/SQL or `AUTOCOMMIT` or `COMMIT ON SUCCESS`, it may not be possible to recover following an outage or timeout. PL/SQL is not reentrant. Once a commit in PL/SQL has executed, that PL/SQL block cannot be resubmitted. Applications either need to unpick the commit which is not sound as that data may have been read, or for batch use a checkpoint and restart technique. When using `AUTOCOMMIT` or `COMMIT ON SUCCESS`, the output is lost.

If your application is using a top-level commit, then there is full support for Transparent Application Continuity (TAC), Application Continuity (AC), and TAF Select Plus. If your application is using `COMMIT` embedded in PLSQL or `AUTOCOMMIT` or `COMMIT ON SUCCESS`, it may not be possible to replay for cases where that the call including the `COMMIT` did not run to completion.

Use ORDER BY or GROUP BY in Queries

Application Continuity ensures that the application sees the same data at replay. If the same data cannot be restored, Application Continuity will not accept the replay. When a `SELECT` uses `ORDER BY` or `GROUP BY` order is preserved. In Autonomous Database the query optimizer most often uses the same access path, which can help in the same ordering of the results. Application Continuity also uses an `AS OF` clause to return the same query results where `AS OF` is allowed.

Considerations for SQL*Plus

SQL*Plus is often our go to tool for trying things out. SQL*Plus of course does not reflect our actual application that will be used in production, so it is always better to use the real application test suite to test your failover plan and to measure your protection. SQL*Plus is not a pooled application so does not have explicit request boundaries. Some applications do use SQL*Plus for example for reports. To use SQL*Plus with failover check the following:

1. FAN is always enabled for SQL*Plus. Use the recommended connect string that auto-configures ONS end points for you.
2. When using SQL*plus the key is to minimize round trips to the database: https://blogs.oracle.com/opal/sqlplus-12201-adds-new-performance-features
3. SQL*Plus is supported for TAC starting with Oracle Database 19c. For best results set a large arraysize. For example (set arraysize 1000). Avoid enabling serveroutput as this creates unrestorable session state.

Notes for Application Continuity on Autonomous Database

There are restrictions for Application Continuity on Autonomous Database, as follows:

- By default Application Continuity is disabled on Autonomous Database. See Configure Application Continuity on Autonomous Database for more information.
- For JDBC clients using Oracle Database JDBC driver and jars, draining for Java based applications on Autonomous Database requires that you use Oracle JDBC driver and Companion Jars 19.12.0.0 or newer, or apply a patch. See Patch Request 31112088 for more information.

Use Standby Databases with Autonomous Data Guard for Disaster Recovery

When you add an Autonomous Data Guard standby database, Autonomous Database provides data protection and disaster recovery for your database.

Autonomous Database Disaster Recovery Options

- **Autonomous Data Guard**: When you add an Autonomous Data Guard standby database, the system creates a standby database that continuously gets updated with the changes from the primary database. You can use Autonomous Data Guard with a standby in the current region, a local standby, or with a standby in a different region, a cross-region standby. You can also use Autonomous Data Guard with both a local standby and a cross-region standby.

  You can also create an Autonomous Data Guard standby in a different tenancy.

  Autonomous Data Guard is available for the following workload types:

  - Data Warehouse
  - Transaction Processing
- **Backup-Based Disaster Recovery**: uses backups to instantiate a peer database at the time of switchover or failover. This enables you to have a lower cost and higher Recovery Time Objective (RTO) disaster recovery option for your Autonomous Database, as compared with Autonomous Data Guard. For local Backup-Based Disaster Recovery, existing local backups are utilized. There are no additional costs for a local Backup-Based Disaster Recovery. Cross-Region Backup-Based Disaster Recovery incurs an additional cost. See Use Backup-Based Disaster Recovery for details on Backup-Based Disaster Recovery.

  Backup-Based Disaster Recovery is available for all workload types.

Topics

- About Standby Databases
  Provides information about enabling and using Autonomous Data Guard for disaster recovery on Autonomous Database.
- Enable Autonomous Data Guard
- Update Standby to Use a Backup Copy Peer
- Autonomous Data Guard Switchover and Failover Operations
  Describes switchover and failover operations with an Autonomous Data Guard standby database.
- Autonomous Data Guard Cross-Region and Cross Tenancy Operations
  Describes cross-region and cross tenancy operations with Autonomous Data Guard.
- Events and Notifications for a Standby Database
  You can use Oracle Cloud Infrastructure Events to be notified of and to specify rules to automatically respond to Autonomous Data Guard operations.
- Autonomous Database Cross-Region Paired Regions
- Use the API
  Provides links for details on using API operations to manage Autonomous Data Guard.
- Autonomous Data Guard Notes
  Provides notes for using Autonomous Database with an Autonomous Data Guard standby database.

About Standby Databases

Provides information about enabling and using Autonomous Data Guard for disaster recovery on Autonomous Database.

When you use Autonomous Data Guard, the system creates a standby database that is continuously updated with the changes from the primary database. You can use Autonomous Data Guard with a standby in the current region, a local standby, or with one or more standby databases in different regions, cross-region standby databases, or you can add both a local standby and one or more a remote standby databases.

You can also create an Autonomous Data Guard standby, either local or remote in a different tenancy.

> **Note:**
>
> Autonomous Data Guard is available with the Data Warehouse and Transaction Processing workload types. Autonomous Data Guard is not available with the JSON and APEX workload types.

By selecting from the disaster recovery options that Autonomous Database provides, you can choose the features and options that meet your Recovery Time Objective (RTO) and Recovery Point Objective (RPO) requirements.

By default, each Autonomous Database instance provides a local Backup-Based Disaster Recovery peer database.

To add automatic failover and to lower your Recovery Time Objective (RTO), you can use a local Autonomous Data Guard standby database.

To use the most resilient disaster recovery option that Autonomous Database offers, you can add a local Autonomous Data Guard standby database and one or more cross-region Autonomous Data Guard standby databases.

In addition, other options using Backup-Based Disaster Recovery enable you to provide lower cost and higher Recovery Time Objective (RTO) disaster recovery options, as compared to Autonomous Data Guard. See Use Backup-Based Disaster Recovery for details on Backup-Based Disaster Recovery.

Topics

- Autonomous Data Guard with Local Standby
  When you use an Autonomous Data Guard standby database in the current region, Autonomous Database provisions a local standby database and monitors the primary database; if the primary database goes down, the standby instance automatically assumes the role of the primary instance.
- Autonomous Data Guard with Cross-Region Standby
  When you add a standby database in a different region, if the primary instance goes down, Autonomous Data Guard provides a standby database that is physically separated in a remote region. The standby database is available to assume the role of the unavailable primary instance.
- Autonomous Data Guard Recovery Time Objective (RTO) and Recovery Point Objective (RPO)
- Autonomous Data Guard Operations
  Autonomous Data Guard provides a set of operations to manage a standby database, including: to enable, switchover, disconnect, or terminate a standby database.
- Autonomous Database Disaster Recovery Status
  Autonomous Database provides information about disaster recovery status on the Autonomous Database Details page.
- Autonomous Data Guard Events
  You can use Oracle Cloud Infrastructure events to respond when Autonomous Database changes its state due to an Autonomous Data Guard related event such as a failover or switchover operation.

Autonomous Data Guard with Local Standby

When you use an Autonomous Data Guard standby database in the current region, Autonomous Database provisions a local standby database and monitors the primary database; if the primary database goes down, the standby instance automatically assumes the role of the primary instance.

Local Autonomous Data Guard peer databases incur the additional cost of the base CPUs and the storage of the Primary database, including any auto-scaled storage usage, billed on the Primary database itself. Auto-scaled CPUs of the Primary database are not billed for additionally on the local Autonomous Data Guard peer database. See Oracle Autonomous Database Serverless Features Billing for more information.

Adding a local standby database provides an identical standby database that allows the following, depending on the state of the primary database:

- If your primary database goes down, Autonomous Data Guard converts the standby database to the primary database with minimal interruption. After failover completes, Autonomous Data Guard creates a new standby database for you.
- You can perform a switchover operation, where the primary database becomes the standby database, and the standby database becomes the primary database.

Autonomous Database does not provide access to a standby database in the current region. You perform all operations, such as scaling up the **ECPU count** (**OCPU count** if your database uses OCPUs) and enabling compute auto scaling on the primary database, and Autonomous Data Guard then performs the same actions on the local standby database. Likewise, you only perform actions such as stopping or restarting the database on the primary database.

A local standby database is created in the same region as the primary database (current region). For better resilience, the standby database is provisioned as follows:

- In regions with more than one availability domain, a local standby database is provisioned automatically in a different availability domain than the primary database.
- In regions with a single availability domain, a local standby database is provisioned automatically in a different fault domain than the primary database (that is, on a different physical machine).

See View Network Information on the OCI Console and Regions and Availability Domains for more information on availability domains.

All Autonomous Database features from the primary database are available when the local standby instance becomes the primary, after the system fails over or after you perform a switchover operation, including the following:

- **Database Options:** The **ECPU count** (**OCPU count** if your database uses OCPUs), Storage, Display Name, Database Name, Auto Scaling, Tags, and BYOL licensing options have the same values after a failover to the standby database or after you perform a switchover.
- **OML Notebooks:** Notebooks and users created in the primary database are available in the standby.
- **APEX Data and Metadata:** APEX information created in the primary database is copied to the standby.
- **ACLs:** The Access Control List (ACL) of the primary database is duplicated for the standby.
- **Private Endpoint:** The private endpoint from the primary database applies to the standby.

Oracle recommends that for a databases on a Private Endpoint, when you create the subnet you use the regional subnet option for optimum availability and latency. See Creating a Subnet for more information.

- **APIs or Scripts:** Any APIs or scripts you use to manage the Autonomous Database continue to work without any changes after a failover operation or after you perform a switchover.
- **Client Application Connections:** Client applications do not need to change their connection strings to connect to the database after a failover to the standby database or after you perform a switchover.
- **Wallet Based Connections:** You can continue using your existing wallets to connect to the database after a failover to the standby database or after you perform a switchover.

Autonomous Data Guard with Cross-Region Standby

When you add a standby database in a different region, if the primary instance goes down, Autonomous Data Guard provides a standby database that is physically separated in a remote region. The standby database is available to assume the role of the unavailable primary instance.

A cross-region standby database is a replica of the primary database and may be used for recovery in case of failure or when the primary is not available. Enabling Autonomous Data Guard with a cross-region standby provides a low RTO solution for disaster recovery in the event an entire region is not available or when the primary database is down for any reason.

Autonomous Data Guard cross-region standby databases incur the additional cost of the base CPUs and twice the storage of the Primary database, including any auto-scaled storage usage, billed on the remote peer database. Auto-scaled CPUs of the Primary are not billed for additionally on the remote peer database. The number of base CPUs is specified by the number of ECPUs (OCPUs if your database uses OCPUs) as shown in the **ECPU count** (**OCPU count**) field on the Oracle Cloud Infrastructure Console.

Autonomous Database allows you to create one or more remote disaster recovery peer databases, depending on your compute model:

- OCPU Compute Model: You can add one remote standby database in a paired region. Paired regions are remote regions where you can create a cross-region peer.
- ECPU Compute Model: You can add multiple remote disaster recovery peers, with up to one peer in each remote paired region. For example if your primary database is in the IAD region, you can add a standby database in PHX and a standby database in SJC, but you cannot add two remote disaster recovery peers in PHX.

Paired regions are remote regions where you can create a cross-region standby database. See Autonomous Database Cross-Region Paired Regions for more information on paired regions.

You perform almost all operations, such as scaling up the **ECPU count** (**OCPU count** if your database uses OCPUs) and enabling compute auto scaling on the primary database. Autonomous Data Guard then performs the same actions on the cross-region standby database.

After you add a remote standby database, Autonomous Database provides access to the remote standby database from the Oracle Cloud Infrastructure Console. Autonomous Database provides access to the remote standby database so that you can perform some operations independently on the remote standby, such as configuring networks and VCNs for private endpoints and adding tagging to define keys and values that are not replicated between the primary database and the remote standby.

> **Note:**
>
> Autonomous Data Guard does not perform automatic failover for a cross-region standby. If the primary database is unavailable and a local standby is unavailable, you can perform a manual failover to make a cross-region standby database assume the primary role.

You cannot connect to a cross-region standby while it operates as a standby database, and it is not available for read-only operations. You can connect to the database in the following cases:

- When the database assumes the primary role after a failover or switchover operation. See Perform a Switchover and Perform a Manual Failover for more information.
- After you convert a standby database to a snapshot standby database. See Convert Cross-Region Disaster Recovery Peer to a Snapshot Standby for more information.

The following areas have differences for failover or switchover from the primary database to a remote standby database, when compared to failover or switchover to a local standby:

- **Display Name**: The display name has a "*_region*" extension. Where *region* is the region name, such as `IAD` or `BOM`.

  If you created the cross-region peer before the introduction of support for multiple cross-region peers, the cross-region peer's display name has a "`_Remote`" extension.
- **OML Notebooks:** After a cross-region switchover or failover, OML notebooks from primary that was switched over or failed over are not present in the primary database (the current primary database after the role change). New OML notebooks can be created.
- **Private Endpoint**: You can independently configure and update private endpoints on a standby database before failover or before you perform a switchover. This allows you to have a private endpoint that is configured differently, after failover or after you perform a switchover. Autonomous Database does not keep the networking configuration synchronized from the primary to a remote standby.

  VCN Peering and domain forwarding are required for wallets to work across regions, with Autonomous Databases with a private endpoint with an Autonomous Data Guard standby, where the primary and the standby database are in different VCNs. See Remote VCN Peering using an RPC and DNS in Your Virtual Cloud Network for information on VCN peering and domain forwarding.
- **Network Access Control List**: By default a disaster recovery primary database and remote peer databases use the same network Access Control Lists (ACLs). Optionally, you can independently edit network ACLs on a remote peer database. This allows you to use different ACLs on a remote peer database.

  See Manage Remote Peer Network ACLs for more information.
- **Tags**: Tags are handled independently on a disaster recovery primary and on remote peer database. This means:
  - When you add, remove, or update a tag on a remote peer, the change applies only on the remote peer database.
  - When you add, update, or remove a tag on the primary, the tag is not added, updated, or removed on remote peer databases.
- **APIs or Scripts:** Any APIs or scripts you use to manage the Autonomous Database need to be updated to call the APIs on the primary database, the current primary database, after a failover or after you perform a switchover.

For mTLS connections, you must download a wallet from the primary database, the current primary database, after a failover or after you perform a switchover. See Cross-Region Disaster Recovery Connection Strings and Wallets for more information.

- **Client Applications**: Client applications need to connect using the connection strings and wallet you download from the primary database, the current primary database, after a failover or after you perform a switchover. See Cross-Region Disaster Recovery Connection Strings and Wallets for more information.
- **Wallet Based Connections:** You must download a wallet and connect using the connection strings from the primary database, the current primary database, to connect to the database after a failover or after you perform a switchover. See Cross-Region Disaster Recovery Connection Strings and Wallets for more information.
- **Autonomous Database Tools**: The tools have different URLs in the primary database, the current primary database, after a failover or after you perform a switchover (the tools URLs do not change for a switchover or failover to a local standby):
  - Database Actions
  - Oracle APEX
  - Oracle REST Data Services (ORDS)
  - Graph Studio
  - Oracle Machine Learning Notebooks
  - Data Transforms
  - MongoDB API
- **Oracle Cloud Infrastructure Object Storage Usage**: After you failover or switchover from the primary database to a standby database, on the primary database (the current primary database) the credentials and the URLs that provide access to Object Storage continue to work as they did before the failover or the switchover, providing access to the following:
  - External tables
  - External partitioned tables
  - External hybrid partitioned tables

> **Note:**
>
> This applies when the Object Storage is available. For rare scenarios when the Object Storage is not available, Oracle recommends having Object Storage backups or replication to a different region. If the Object Storage is not available (that is the Object Storage resource you used with the Primary before a switchover or failover), you may update your user credentials and parameters that set URLs for Object Storage so that the parameters specify values to access an available region's Object Storage. See Using Replication for more information.

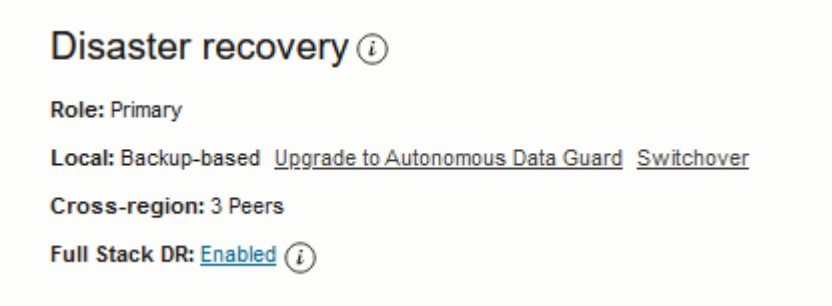
Cross Tenancy Autonomous Data Guard with Cross-Region Standby

You can enable cross tenancy Autonomous Data Guard with a cross-region standby. When you add a cross tenancy Autonomous Data Guard standby database in a different region, Autonomous Database provisions a cross-region standby database in the destination tenancy. With a cross tenancy Autonomous Data Guard standby, you can failover, switchover, or create a snapshot standby with a cross-region standby in a different tenancy. This feature allows you to use Autonomous Data Guard to migrate a database to a different tenancy.

See Use a Cross Tenancy Autonomous Data Guard Standby Database for more information.

OCI Full Stack Disaster Recovery with an Autonomous Data Guard Cross-Region Standby

When Full Stack Disaster Recovery is enabled, the Autonomous Database details page, under Disaster recovery, shows the **Full Stack DR** field as **Enabled**.



See Use OCI Full Stack Disaster Recovery with Autonomous Database for more information.

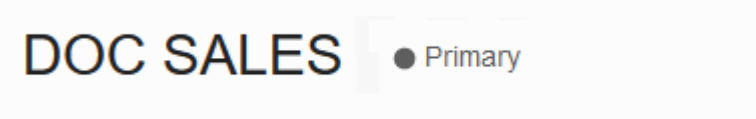
Topics

- Autonomous Data Guard Database Role
  After you add a cross-region standby database, each database has a designated role: primary, standby, or snapshot standby.
- Autonomous Data Guard Cross-Region Failover and Switchover
  You can have one local disaster recovery peer and optionally you can add one or more cross-region peers (multiple cross-region peers are allowed with the ECPU compute model). In both the local and cross-region cases, either peer can be a Backup-Based Disaster Recovery copy or an Autonomous Data Guard standby.
- Autonomous Data Guard Database Cross-Region Backup and Restore
  After you add an Autonomous Data Guard cross-region standby database, backup and restore from backup is handled as follows:
- Cross-Region Disaster Recovery Connection Strings and Wallets
  When you add an Autonomous Data Guard cross-region (remote) standby database, or when you use a cross-region Backup-Based Disaster Recovery peer, the wallet and connection string from the primary database contains only the hostname of the primary database.
- Autonomous Data Guard with Customer Managed Keys
  When you add an Autonomous Data Guard cross-region standby, there are special considerations when the primary database is using customer-managed keys, or if you want to switch to using customer-managed keys on the primary database.
- Replicating Backups to a Cross-Region Autonomous Data Guard Standby
  When you add a cross-region Autonomous Data Guard standby you can enable cross-region backup replication so that automatic backups from the primary are also available on a remote region.
- Autonomous Data Guard Cross-Region BYOL Licensing
  The **BYOL ECPU limit** you set on an Autonomous Data Guard Primary database does not apply to a cross-region or cross tenancy Autonomous Data Guard Standby database.

Autonomous Data Guard Database Role

After you add a cross-region standby database, each database has a designated role: primary, standby, or snapshot standby.

The role specifies the current state of a database, primary, standby, or snapshot standby, and this value changes after you perform a switchover or a failover or after you convert a standby database to a snapshot standby. You can view the Autonomous Database role in the icon that shows next to the display name on the Autonomous Database Information page. For example:





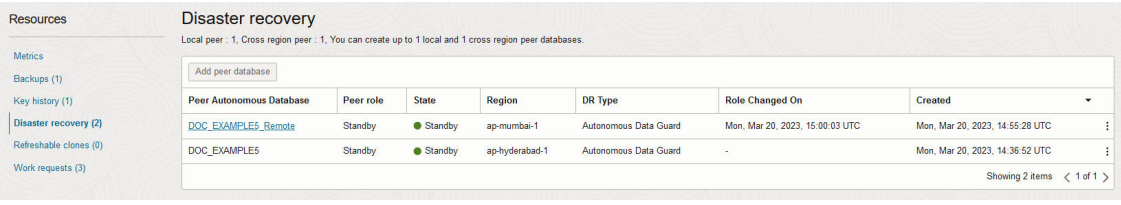
After you add a cross-region standby database, you can view the role in the **Disaster recovery** area on the details page. The role is one of:

- The **Role** shows **Primary** on the primary database.
- After a switchover or failover, on the same database the **Role** shows **Standby**.
- After you convert a cross-region peer to a snapshot standby, the **Role** shows as **Snapshot Standby**.

To view details for a peer, on the Autonomous Database Information page, under **Resources** select **Disaster recovery**:

- **Standby** (local): the **Peer role** column shows **Standby** and the database has the same display name in the **Peer Autonomous Database** column. The **Region** column shows the name of the current region.
- **Standby** (cross-region) the **Peer role** column shows **Standby** for a remote standby database and the database has the same name with an "*_region*" extension in the **Peer Autonomous Database** column. You can click the link to access the remote database. The **Region** column shows the name of the remote region.

  If you created the cross-region peer before the introduction of support for multiple cross-region peers, the cross-region peer's display name has a "`_Remote`" extension.
- **Snapshot standby**: the **Peer role** column shows **Snapshot standby**. The **Region** column shows the name of the remote region.



Autonomous Data Guard Cross-Region Failover and Switchover

You can have one local disaster recovery peer and optionally you can add one or more cross-region peers (multiple cross-region peers are allowed with the ECPU compute model). In both

the local and cross-region cases, either peer can be a Backup-Based Disaster Recovery copy or an Autonomous Data Guard standby.

With both a current region and one or more cross-region Autonomous Data Guard peer databases, depending on the state of the primary database, you have the following options:

- If your primary database goes down and the local standby database is available, Autonomous Data Guard automatically performs failover to convert the local standby database to the primary database, with minimal interruption. After failover completes, Autonomous Data Guard creates a new local standby database for you. If automatic failover is not possible, you have the option to perform a manual failover.

  Autonomous Data Guard continues to use the same cross-region peer databases.
- If your primary database goes down and the local standby database is not available, you can perform a manual failover to a cross-region peer database and the cross-region peer database you failover to becomes the primary database.

  In this case, after the failover completes, Autonomous Data Guard does not create a new local standby database (by default you have a backup copy peer).
- You can perform a switchover operation, where the primary database becomes the local standby database, and the local standby database becomes the primary database.

  Autonomous Data Guard continues to use the same cross-region peer databases.
- You can perform a switchover operation, where a cross-region peer database becomes the primary database (and the database that was the primary is recreated as a new standby database so that it becomes a peer database).

  A switchover changes the roles of the primary and a peer database. If you perform a switchover two times between the same two remote regions, the primary database returns to again be the primary database.

Autonomous Data Guard Database Cross-Region Backup and Restore

After you add an Autonomous Data Guard cross-region standby database, backup and restore from backup is handled as follows:

- If the primary database is restored from a backup, a new remote standby is created from the restored primary database.
- Automatic Backups are only taken on the primary database (the database showing **Role**: **Primary**). For example, after a switchover or failover, the database with the **Primary** role starts to perform automatic backups. A database with the **Standby** role no longer takes backups. If you switchover again, the database that becomes the **Primary** role database starts taking backups again.
- You cannot restore or clone from a backup when a peer database is in the **Standby** role. Backups are only taken on the database in the **Primary** role, and the restore operation is not available from the Oracle Cloud Infrastructure Console on a **Standby** database.

Cross-Region Disaster Recovery Connection Strings and Wallets

When you add an Autonomous Data Guard cross-region (remote) standby database, or when you use a cross-region Backup-Based Disaster Recovery peer, the wallet and connection string from the primary database contains only the hostname of the primary database.

In addition, the wallet and connection string from a remote peer database contains only the hostname of the remote database. This applies for both instance and regional wallets.

Oracle recommends that you configure your applications running on the Primary Role database to use the wallet or connection string downloaded from the Primary database. For

applications that run on a remote database, use the wallet or connection string downloaded from the remote database (where the remote database is the current primary database after a failover or after you perform a switchover). You can obtain these connection strings or the wallet by clicking **Database connection** on the Oracle Cloud Infrastructure Console.

For example, if your cross-region Autonomous Data Guard is setup with the primary in Ashburn (IAD) and a cross-region standby in Phoenix (PHX), Oracle recommends that your mid-tier applications running in IAD use the connection string or wallet from the primary database in IAD, and your corresponding applications that run in PHX after a failover or after you perform a switchover, use the connection string or wallet from the standby database in PHX. During regional failover or switchover, Oracle recommends failing over both your database and your mid-tier applications to the new Primary role database for optimum performance and to minimize any cross-regional latency.

See Download Client Credentials (Wallets) for more information.

If required by your application, you may manually construct connection strings containing both primary and remote database hostnames, to support connecting to either instance that is available and open for connections automatically, the primary or the remote database.

For details on the steps to manually create these connection strings see:

- Configure your mid-tier applications for cross-region disaster recovery
- Autonomous Data Guard Notes

Autonomous Data Guard with Customer Managed Keys

When you add an Autonomous Data Guard cross-region standby, there are special considerations when the primary database is using customer-managed keys, or if you want to switch to using customer-managed keys on the primary database.

> **Note:**
>
> Autonomous Database supports multiple customer-managed keys providers. Only Oracle Cloud Infrastructure Vault is supported for use with Autonomous Data Guard. Other vaults are not supported for customer-managed keys.

In order for a remote standby to be able to use the same master encryption key as the primary database, the master encryption key must be replicated to the remote region. Customer-Managed Encryption Keys are only supported with a single cross-region Autonomous Data Guard standby. Multiple cross-region standbys are not supported because Oracle Cloud Infrastructure Vault only supports replication to one remote region.

Consider the following cases:

- Adding an Autonomous Data Guard remote standby is allowed if the Autonomous Database is using customer-managed keys. When the database is using a customer-managed key and you add an Autonomous Data Guard cross-region standby, the **Region** list in the **Add peer database** dialog only shows the regions that contain the replicated vault and keys. If you don't see a remote region listed, you need to replicate your vault and keys to the region where you want your standby database (this must be a paired region).
- Switching to customer-managed keys is allowed on the primary when you have an Autonomous Data Guard cross-region standby. In the case when the database is using Oracle-managed keys and you switch to customer-managed keys on the primary, you only see the keys that are replicated in both the primary and the standby regions. The Manage Encryption Key **Vault** and **Master encryption key** lists only show vaults and keys that are

replicated across both the primary and the standby regions. If you don't see a key listed, replicate your vault and keys to a paired region.

See the following for more information:

- Manage Master Encryption Keys in OCI Vault
- Replicating Vaults and Keys
- Overview of Vault
- Autonomous Database Cross-Region Paired Regions

Replicating Backups to a Cross-Region Autonomous Data Guard Standby

When you add a cross-region Autonomous Data Guard standby you can enable cross-region backup replication so that automatic backups from the primary are also available on a remote region.

By default the backups taken on the primary are not replicated to a cross-region standby database. When you enable cross-region backup replication, up to 7 days of automatic backups for the primary are replicated to a cross-region standby database. When this feature is enabled, automatic backups are available in the remote region as follows:

- After a switchover or failover you can restore or clone to any timestamp in the past seven (7) days, or to any timestamp in the specified retention period when the retention period is set to less than seven days.
- All backups for the primary that are replicated to the remote region are deleted on the remote region peer after seven days, or after the retention period number of days when the retention period is set to less than seven days.
- You cannot modify the backup retention period for replicated backups, except if you modify the backup retention period on the primary to specify a value less than seven days. In this case, the retention period for replicated backups on the remote region matches the automatic backup retention period set on the primary.

Cross-region backup replication incurs an additional cost. See Oracle Autonomous Database Serverless Features Billing for more information.

See Add a Cross-Region Standby Database and Enable or Disable Backup Replication for an Existing Cross-Region Standby for more information.

Note the following for cross-region automatic backup replication:

- After a switchover or a failover, while the cross-region database is in the primary role, backups are taken on the current primary and are replicated to the current (remote) standby.
- In a remote region you can create a clone from a replicated backup while the database is in the **Standby** role.

Autonomous Data Guard Cross-Region BYOL Licensing

The **BYOL ECPU limit** you set on an Autonomous Data Guard Primary database does not apply to a cross-region or cross tenancy Autonomous Data Guard Standby database.

On a cross-region or cross-tenancy Standby you can independently set the **BYOL ECPU limit**, as required. Setting a value for **BYOL License limit** limits how many ECPUs will be covered by BYOL licenses.

For example, consider an 8 ECPU Autonomous Data Guard Primary database using BYOL licensing. When you add a cross-region or cross tenancy Standby, the Standby takes its licensing from the Primary (using BYOL licensing).

In this example, if you set the **BYOL License limit** to 4 (ECPUs) on the Primary, then 4 of the 8 ECPUs use BYOL licensing. However, the **BYOL License limit** you set on the Primary does not apply on a cross-region or cross tenancy Standby. The Standby uses Bring your own license (BYOL) licensing (without a **BYOL License limit**). If you separately set a **BYOL License limit** on the Standby, for example if you set the **BYOL License limit** value to 2 (ECPUs), 2 ECPUs on the Standby are billed using BYOL licensing and 6 ECPUs. Similarly, the **BYOL ECPU limit** you set on the Standby does not affect the Primary's **BYOL ECPU limit**.

See Choose Bring Your Own License Option While Provisioning or Cloning and Choose Bring Your Own License on Autonomous Database (ECPU Compute Model) for more information.

Autonomous Data Guard Recovery Time Objective (RTO) and Recovery Point Objective (RPO)

Autonomous Data Guard monitors the primary database and if the instance goes down, the local standby instance assumes the role of the primary instance, according to the Recovery Time Objective (RTO) and Recovery Point Objective (RPO).

If a local Autonomous Data Guard standby instance is not available and you have enabled cross-region disaster recovery, you can manually fail over to the cross-region standby.

If you do not add a cross-region Autonomous Data Guard standby, you have the option to add a cross-region Backup-Based Disaster Recovery peer. See Backup-Based Disaster Recovery Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for details on the RTO and RPO with Backup-Based Disaster Recovery.

The RTO is the maximum amount of time required to restore database connectivity to a standby database after a manual or automatic failover is initiated. The RPO is the maximum duration of potential data loss on the primary database.

Local Autonomous Data Guard Standby

When you add a local standby database Autonomous Data Guard provides these options for failover or switchover:

- **Automatic Failover or Switchover**:

  When you enable Autonomous Data Guard you can select a data loss limit. The default data loss limit for automatic failover is 0 (valid values are 0 to 3600 seconds). For example, a data loss limit of 0 means that Autonomous Data Guard only performs automatic failover when there is no data loss. This means if Autonomous Data Guard can verify that there is no data loss, it automatically fails over in case of a problem. When there is a problem and Autonomous Data Guard determines that the possible data loss is greater than the data loss limit, automatic failover does not happen and you have the option to perform a manual failover.

- **Manual Failover**: the RTO is two (2) minutes and the RPO 10 seconds

Cross-Region Autonomous Data Guard Standby

When you add a cross-region standby database, the RTO and RPO numbers for failover to the Autonomous Data Guard cross-region standby are as follows:

- **Switchover**: the RTO is less than ten (10) minutes and RPO is zero (0).
- **Automatic Failover**: Not available
- **Manual Failover**: the RTO is less than ten (10) minutes and RPO is up to one (1) minute.

See the following for more information:

- Automatic Failover with a Standby Database for details on automatic failover.
- Perform a Switchover
- Perform a Manual Failover
- Service Level Objectives (SLOs)

Autonomous Data Guard Operations

Autonomous Data Guard provides a set of operations to manage a standby database, including: to enable, switchover, disconnect, or terminate a standby database.

Operation	Description
Convert to Snapshot Standby	Converting a disaster recovery peer to a snapshot standby opens the database in read-write mode and the cross-region disaster recovery peer temporarily stops refreshing data from the source database. See Convert Cross-Region Peer to Snapshot Standby for more information.
Disable Autonomous Data Guard	If you have a local standby database or a cross-region standby database, you can change the disaster recovery type to Backup-Based Disaster Recovery for the local standby or you can terminate a cross-region standby. In either case, disabling Autonomous Data Guard terminates the standby database. See Update Standby to Use a Backup Copy Peer or Disable a Cross-Region Standby Database for details.
Disconnect Standby	When you disconnect a cross-region standby, the standby is disassociated from the Primary database. This converts the database from a peer database to a standalone database. Following the disconnect operation you are not allowed to reconnect to the Primary. See Disconnect a Peer Database and Disconnect a Snapshot Standby for more information.
Enable Autonomous Data Guard	If you are using Backup-Based Disaster Recovery, you can update your disaster recovery type to local (current region) Autonomous Data Guard, or you can add an Autonomous Data Guard cross-region standby. See Enable Autonomous Data Guard and Add a Cross-Region Standby Database for details.
Failover - Automatic	After you add a local Autonomous Data Guard standby database, the system monitors the primary instance and automatically fails over to a local standby database in certain scenarios. See Automatic Failover with a Standby Database for more information.
Failover - Manual	If the primary database is not available you can perform a manual failover to change roles to make a standby database the primary database: • If a local standby is available, you can manually failover to the local standby (you do not have the option to failover to a remote standby if a local standby is available). • If a local standby is not available, you have the option to manually failover to a remote standby. See Perform a Manual Failover for details.
Switchover	When Autonomous Data Guard is enabled, switchover changes the roles of the primary and the standby, the standby database becomes the primary, and the primary database becomes the standby. If you have both a local standby database (current region), and a cross-region standby database (remote), you can choose to switchover to either the local standby or the remote standby. See Perform a Switchover for details.

Operation	Description
Terminate	If you want to terminate the primary instance, select **More actions** and **Terminate**. Terminating the primary instance also terminates a local standby database. If you have both a local standby database (current region), and a cross-region standby database, you must terminate the cross-region standby database before you terminate the primary database. See Terminate a Cross-Region Standby Database for details.

Autonomous Database Disaster Recovery Status

Autonomous Database provides information about disaster recovery status on the Autonomous Database Details page.

In the **Disaster recovery** area:

The **Role** field shows the role of the current database, as follows:

- When you have either a local backup copy peer or a local Autonomous Data Guard standby, the Oracle Cloud Infrastructure Console shows the **Role** field value **Primary**. Autonomous Database does not provide access to a local standby database (or to a local backup copy peer).
- When using either a cross-region backup copy peer or a cross-region Autonomous Data Guard standby, the Oracle Cloud Infrastructure Console shows the **Role** field value **Primary** if you are viewing the primary database and shows **Standby** if you are viewing the details for a standby database.
- **Switchover**: Provides a link so that you can perform a switchover operation.
- **Failover**: When the primary database is not available and you have a local standby and automatic failover was not successful, the failover link allows you initiate a manual failover.

  When the primary database is not available and you have a cross-region standby and failover to a local standby is not possible, the failover link allows you initiate a manual failover to the remote standby database.

To view the peer Autonomous Database information, under **Resources** click **Disaster recovery**. This area lists the peer autonomous database information. The **State** column shows the state of a standby database, as follows:

- **Provisioning**
  - This state shows when you enable Autonomous Data Guard and indicates that a standby database is provisioning (until the standby database state changes to **Standby**).
  - This state shows after a failover to a local standby when a local standby database is being recreated.
  - This state shows if a restore from backup operation is performed on the primary database, the local standby is recreated and the **State** column shows Provisioning.
- **Standby**: Indicates that a standby is available and ready for either a switchover or a failover operation.

> **Note:**
>
> When a standby database is stopped, the standby state shows **Standby**. A standby database never shows the **Stopped** state.

- **Role Change in Progress**: Indicates a failover or switchover operation started.

Autonomous Data Guard Events

You can use Oracle Cloud Infrastructure events to respond when Autonomous Database changes its state due to an Autonomous Data Guard related event such as a failover or switchover operation.

Autonomous Database events include the following:

- Begin automatic failover
- End automatic failover
- Begin disable Autonomous Data Guard
- Begin enable Autonomous Data Guard
- Begin failover
- Begin switchover
- End disable Autonomous Data Guard
- End enable Autonomous Data Guard
- End failover with failover result of success or failure.
- End switchover with switchover result of success or failure.

Based on events you can perform actions or send notifications. See Events and Notifications for a Standby Database for more information on using events and producing notifications.

Enable Autonomous Data Guard

To enable Autonomous Data Guard you update the disaster recovery type to use a standby database.

By default and at no additional cost, Autonomous Database provides a local backup copy peer for each Autonomous Database instance. You enable Autonomous Data Guard by changing the disaster recovery type to use a standby database. Autonomous Data Guard provides a lower Recovery Time Objective (RTO), compared to using a backup copy peer, and provides for automatic failover to a local standby when the primary database is not available.

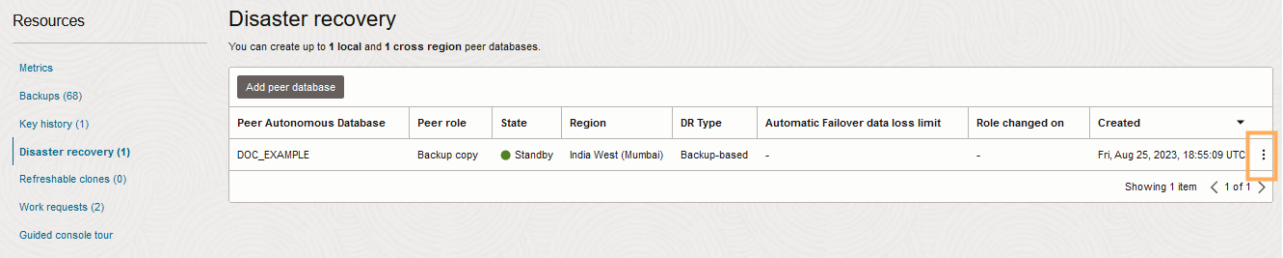
> **Note:**
>
> When you change your disaster recovery type to use a local standby database, you also have the option to add a second cross-region disaster recovery option, either a cross-region Autonomous Data Guard standby database or a cross-region backup copy peer.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload, click one of: Autonomous Data Warehouse or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To change your disaster recovery type to add a local Autonomous Data Guard standby database:

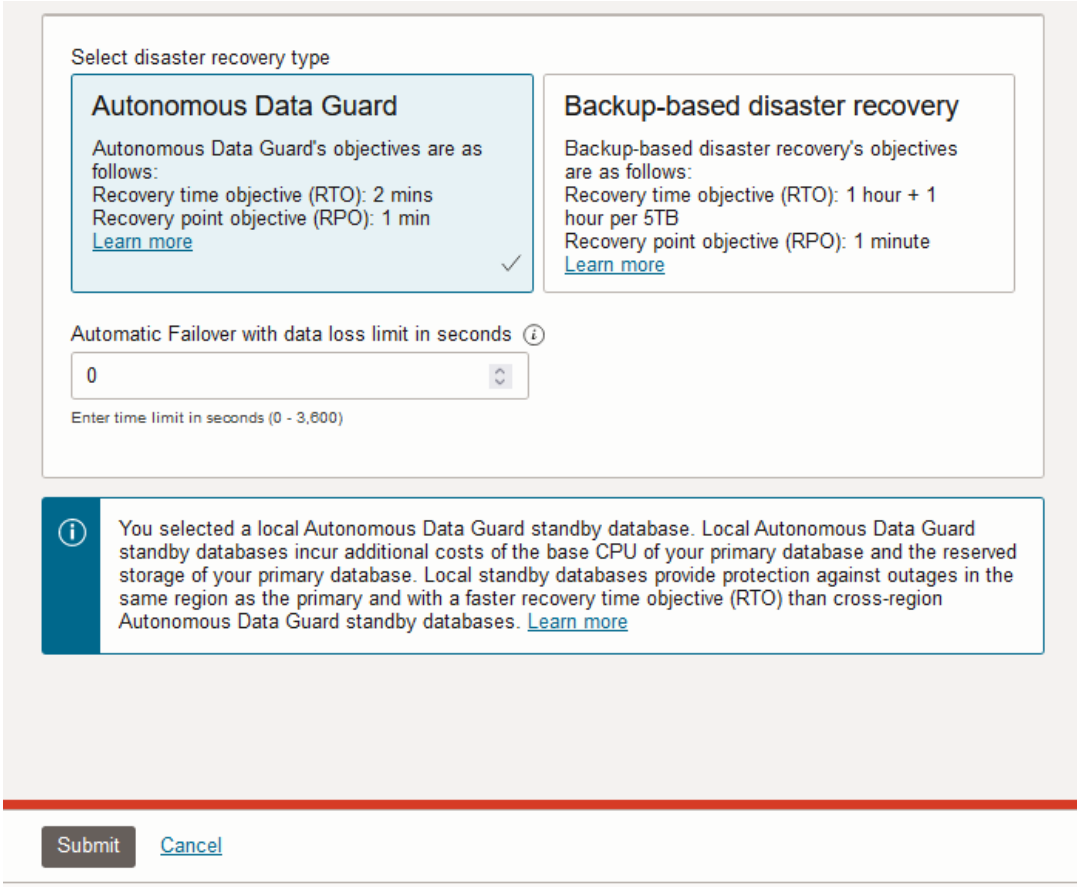
1. On the **Autonomous Database Details** page, in the **Resources** area select **Disaster recovery**.



2. In the row showing the database's disaster recovery details, click ⋮ at the end of the row and select **Update disaster recovery**.

   This shows the Update disaster recovery page.

3. Select Autonomous Data Guard.



4. In the **Automatic Failover with data loss limit in seconds** field, accept the default data loss limit of 0, or enter a custom value for the automatic failover data loss limit.

   See Automatic Failover with a Standby Database for more information.

5. Click **Submit**.

   The Autonomous Database Lifecycle state changes to **Updating**.

   The **State** column shows Provisioning while Autonomous Database provisions the standby database.

   After some time the **Lifecycle state** shows **Available** and the standby database provisioning continues.

   When provisioning completes the **DR type** column shows Autonomous Data Guard.

> **Note:**
>
> While you add a new standby database, the primary database is available for read/write operations. There is no downtime on the primary database.

Notes for changing the disaster recovery type to a local Autonomous Data Guard standby database:

- Autonomous Database generates an Enable Autonomous Data Guard work request. To view the request, under **Resources** click **Work Requests**.

The work request may complete to 100% before you see Autonomous Data Guard in the **DR type** column when you select **Disaster recovery** under Resources. The provisioning process takes several minutes.

- While you add a local standby database, when the **Lifecycle state** field shows **Updating**, the following actions are disabled for the primary database:
  - Move Resource. See Move an Autonomous Database to a Different Compartment for information on moving an instance.
  - Stop. See Stop Autonomous Database for information on stopping an instance.
  - Restart. See Restart Autonomous Database for information on restarting an instance.
  - Restore. See Restore and Recover your Autonomous Database for information on restoring.

Update Standby to Use a Backup Copy Peer

Describes the steps to change the disaster recovery type from Autonomous Data Guard standby database to Backup-Based Disaster Recovery.

These steps terminate the standby database. By default Autonomous Database provides a Backup-Based Disaster Recovery peer and these steps change the disaster recovery type from Autonomous Data Guard to backup copy.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload, click one of: Autonomous Data Warehouse or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To change the disaster recovery type:

1. On the Autonomous Database Details page, under **Resources**, select **Disaster recovery**.
2. In the row showing the disaster recovery details, click the Actions icon (three dots) and select **Update DR type**.

   This shows the Update disaster recovery type page.
3. Select **Backup-based disaster recovery**.
4. Click **Submit**.
5. In the Disable Autonomous Data Guard dialog, click **Disable Autonomous Data Guard**.

   While the standby database is terminating, the **Lifecycle state** changes to **Updating**.

   Autonomous Database generates the Disable Autonomous Data Guard work request. To view the request, under **Resources**, click **Work Requests**.

Autonomous Data Guard Switchover and Failover Operations

Describes switchover and failover operations with an Autonomous Data Guard standby database.

- Perform a Switchover
  When you perform a switchover, the primary database becomes the standby database and the standby database becomes the primary database, with no data loss.
- Automatic Failover with a Standby Database
- Perform a Manual Failover

Perform a Switchover

When you perform a switchover, the primary database becomes the standby database and the standby database becomes the primary database, with no data loss.

A switchover is typically done to test failover to the standby database for audit or certification reasons, or to test your application's failover procedures when you have added an Autonomous Data Guard standby database.

For switchover to a standby database, the Oracle Cloud Infrastructure Console on the primary database shows a **Switchover** link under **Disaster recovery** area when both the primary database and a standby database are available. You can perform a switchover when the primary database **Lifecycle state** shows **Available** or **Stopped**, and a standby database is available (the **State** field shows **Standby**).

To see either local or cross-region standby database state, under **Resources** click **Disaster recovery** and for the standby database listed in the **Peer Autonomous Database** column, check that the **State** shows **Standby**.

Using the Autonomous Database API, you can initiate the switchover operation at any time. See Use the API for more information.

> **Note:**
>
> When Full Stack Disaster Recovery is enabled, the Autonomous Database details page, under Disaster recovery, shows the **Full Stack DR** field as **Enabled**. See Use Full Stack Disaster Recovery to Perform Switchover or Failover Operations for details on performing a switchover or a failover when Full Stack Disaster Recovery is enabled.

- Perform a Switchover to a Local Standby
- Perform a Switchover to a Cross-Region Standby
- Notes for Performing a Switchover

Perform a Switchover to a Local Standby

When you perform a switchover, the primary database becomes the standby database and the standby database becomes the primary database, with no data loss.

To perform switchover to a local standby, you access the primary database from the Oracle Cloud Infrastructure Console.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.

- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload, click one of: Autonomous Data Warehouse or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To perform a switchover:

1. On the Autonomous Database Details page, under **Disaster recovery**, in the **Local** field, click **Switchover**.

   As an alternative, to initiate a switchover you can select **More actions** and **Switchover**.
2. In the **Confirm switchover to peer** panel, enter the peer database name to confirm that you want to switch over.
3. In the **Confirm switchover to peer** panel, click **Switchover**.

   When concurrent operations such as scaling are active, the confirmation also confirms either pausing or canceling the concurrent operation. See Concurrent Operations on Autonomous Database for more information.

   The database **Lifecycle state** changes to **Updating**. To see the state of the peer database, under **Resources** click **Disaster recovery**. The state **State** shows **Role change in progress**.

When the switchover completes, Autonomous Database reports the time of the last switchover when you hover over the ⓘ in the **Role** field.

See Notes for Performing a Switchover for more information.

Perform a Switchover to a Cross-Region Standby

When you perform a switchover, the primary database becomes the standby database and the standby database becomes the primary database, with no data loss.

> **Note:**
>
> For a cross-region switchover you must initiate the switchover from the Standby database.

You have several options to access the Standby database:

- Select the remote region in Oracle Cloud Infrastructure Console and then access the Standby database.
- Access the primary, and from the primary database you can access the standby from the Autonomous Database Details page by selecting **Disaster recovery** under **Resources** and clicking the link for the standby database in the **Peer Autonomous Database** column.

> **Note:**
>
> When Full Stack Disaster Recovery is enabled, the Autonomous Database details page, under Disaster recovery, shows the **Full Stack DR** field as **Enabled**. See Use Full Stack Disaster Recovery to Perform Switchover or Failover Operations for details on performing a switchover or a failover when Full Stack Disaster Recovery is enabled.

To perform a switchover:

1. On the cross-region standby database, on the Autonomous Database Details page, under **Disaster recovery**, in the **Role** field, click **Switchover**.

   As an alternative, to initiate a switchover you can select **More actions** and **Switchover**.
2. In the **Confirm switchover to peer** panel, enter the peer database name to confirm that you want to switch over.
3. In the **Confirm switchover to peer** panel, click **Switchover**.

   When concurrent operations such as scaling are active, the confirmation also confirms either pausing or canceling the concurrent operation. See Concurrent Operations on Autonomous Database for more information.

   The database **Lifecycle state** changes to **Updating**. To see the state of the peer database, under **Resources** click **Disaster recovery**. The state **State** shows **Role change in progress**.

When the switchover completes, Autonomous Data Guard does the following:

- The display name shows the standby indicator.
- The **Disaster recovery** resource information is updated to reflect the switchover. Under **Resources** select **Disaster recovery** to see the updated information.
- Autonomous Database reports the time of the last switchover when you hover over the in the **Role** field.

See Notes for Performing a Switchover for more information.

Notes for Performing a Switchover

Notes for performing a switchover on Autonomous Database.

- For a cross-region switchover, you must initiate the switchover from the Standby database.
- During the switchover most of the actions on the Oracle Cloud Infrastructure Console are not available and the Autonomous Database Information page shows the **Lifecycle state** with the value **Updating**.
- The switchover operation keeps the original state of the primary database. If the primary database is stopped when you perform a switchover, the primary database is stopped after the switchover.
- Autonomous Database generates the Switchover Autonomous Database work request. To view the request, under **Resources** click **Work Requests**.
- After a switchover or failover to the standby, the standby becomes the primary:
  - The Oracle Cloud Infrastructure Metrics display information about the primary database.

  - The graphs on the **Database Dashboard** card in Database Actions display information about the primary database.
  - The graphs and metrics do not contain information about the database that was the primary database, before the switchover or failover operation.
- You cannot cancel a cross-region switchover operation after the switchover begins and the **State** shows **Role change in progress**. Your options are:
  - Try or retry a switchover or a failover until the operation succeeds.
  - File a service request at Oracle Cloud Support or contact your support representative.

Automatic Failover with a Standby Database

After you add a local Autonomous Data Guard standby database, the system monitors the primary instance and automatically fails over to a local standby database in certain scenarios.

If automatic failover is not possible, you have the option to perform a manual failover. Automatic failover does not apply to a cross-region standby.

Autonomous Database automatically fails over to a local standby database as follows:

- When the primary database becomes unavailable and users are not able to connect, Autonomous Data Guard automatically fails over to a local Autonomous Data Guard standby database if a local standby database is available, based on the Recovery Time Objective (RTO) and when Autonomous Data Guard determines that the automatic failover data loss limit will not be exceeded.
- Autonomous Data Guard performs automatic failover to a local standby database when a local standby database is available and the system can guarantee either the default zero data loss or up to the data loss limit you specified a data loss limit when you enabled Autonomous Data Guard. You can specify an automatic failover data loss limit between 0 and 3600 seconds. If this target cannot be met then automatic failover does not occur and you have the option to perform a manual failover.

  See Enable Autonomous Data Guard for information on setting the automatic failover data loss limit.
- After automatic failover to a local standby database completes, Autonomous Database creates a new local standby database for you.

  Autonomous Data Guard for a particular standby, either local or remote, is not enabled while the system is provisioning the new standby database. After Autonomous Data Guard completes the provisioning step for the standby database and it becomes available, you then have a new standby database with Autonomous Data Guard enabled on it.
- After automatic failover completes, Autonomous Database reports the time of the last failover when you hover over the  in the **Role** field.
- After an automatic Autonomous Data Guard failover, if there was a regional failure, when the region comes back online the standby database is automatically reconnected or if required reprovisioned.

If the primary database has failed or is unreachable and the conditions for Autonomous Data Guard automatic failover have not been met, then Oracle Cloud Infrastructure console shows a banner indicating that automatic failover did not succeed, provides a reason, such as possible data loss above the default 0 RPO or above the limit you set for the automatic failover data loss limit, and provides a link for you to initiate manual failover. See Perform Manual Failover to a Local Standby Database for more information.

> **Note:**
>
> Autonomous Data Guard automatic failover is disabled when the **Lifecycle state** is either: **Restore in Progress** or **Upgrading**.

See Autonomous Data Guard Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for more information.

Perform a Manual Failover

When Autonomous Data Guard cannot automatically fail over to a local standby database, if a local standby database is available you can perform a manual failover to make the local standby database the primary database.

If a cross-region standby is available, you can perform a switchover to make the cross-region standby database the primary database. If the switchover fails, you can initiate a manual failover to the cross-region standby. It is possible for data loss to occur with a manual failover.

When you initiate a manual failover Autonomous Data Guard fails over to the standby database based on the Recovery Time Objective (RTO) and Recovery Point Objective (RPO) targets. See Autonomous Data Guard Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for more information.

- Perform Manual Failover to a Local Standby Database
- Perform Manual Failover to a Cross-Region Standby Database
  If a cross-region peer is available you can perform a switchover to make a cross-region peer the primary database. If the switchover fails, you can initiate a manual failover to the cross-region peer.
- Notes for Manual Failover with a Standby Database

Perform Manual Failover to a Local Standby Database

When Autonomous Data Guard cannot automatically fail over to a local standby database, if a local standby database is available you can perform a manual failover to make the local standby database the primary database.

When you add a local Autonomous Data Guard standby and automatic failover is not successful, Oracle Cloud Infrastructure console shows a banner with information about why the automatic failover was not successful. The Oracle Cloud Infrastructure console also shows a **failover** link in the **Role** field that you can click to initiate a manual failover to the local standby. The failover link only shows when the primary database is unavailable and a standby database is available. That is, the **Lifecycle state** field shows **Unavailable** and the local standby database is available.

Using the API, you can initiate manual failover at any time. See Use the API for information on using the API.

To see the standby database status, under **Resources** click **Disaster recovery** and for the standby database listed in the **Peer Autonomous Database** column, check that the **State** field shows **Available** or **Stopped**.

Perform the following prerequisite steps as necessary:

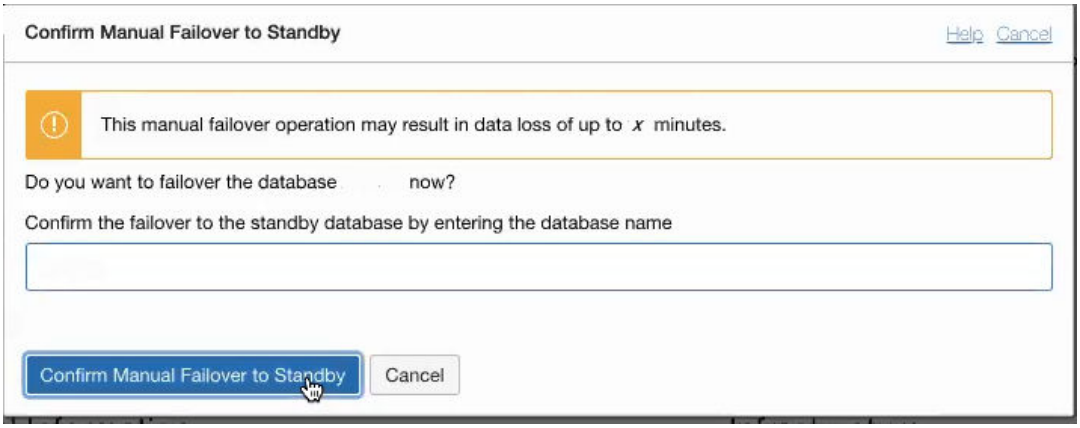
- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.

- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload, click one of: Autonomous Data Warehouse or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To initiate a manual failover when the primary database is unavailable and the local standby is available:

1. On the **Details** page, under **Disaster recovery**, in the **Local** field, click **Failover**.

   This shows the **Confirm Manual Failover to peer** dialog, along with information on possible data loss that may result if you perform the manual failover to standby.



2. In the **Confirm manual failover to peer** dialog, enter the Autonomous Database name to confirm that you want to failover.
3. In the **Confirm manual failover to peer** dialog, click **Confirm manual failover to peer**.

See Notes for Manual Failover with a Standby Database for information on steps that Autonomous Data Guard takes after failover completes.

Perform Manual Failover to a Cross-Region Standby Database

If a cross-region peer is available you can perform a switchover to make a cross-region peer the primary database. If the switchover fails, you can initiate a manual failover to the cross-region peer.

It is possible for data loss to occur with a manual failover. When you initiate a manual failover, Autonomous Database fails over to the peer database based on the Recovery Time Objective (RTO) and Recovery Point Objective (RPO) targets. See Autonomous Data Guard Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for more information.

With both a local Autonomous Data Guard standby and one or more cross-region Autonomous Data Guard standby databases, when automatic failover is not successful and the local standby database is available, Oracle recommends that you attempt a manual failover to the local standby first (not to a remote standby).

If a local standby is unavailable or a manual failover to the local standby fails, you can perform a manual switchover to a cross-region standby. If the switchover to the cross-region standby fails, on the standby database the Oracle Cloud Infrastructure Console shows a **Failover** link in the **Role** field that you can click to initiate a manual failover to the standby database.

Using the API, you can initiate manual failover at any time. See Use the API for information on using the API.

> **Note:**
>
> When Full Stack Disaster Recovery is enabled, the Autonomous Database details page, under Disaster recovery, shows the **Full Stack DR** field as **Enabled**. See Use Full Stack Disaster Recovery to Perform Switchover or Failover Operations for details on performing a switchover or a failover when Full Stack Disaster Recovery is enabled.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the  next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload, click one of: Autonomous Data Warehouse or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To initiate a manual failover to a cross-region standby:

1. On the standby database, perform a switchover. See Perform a Switchover to a Cross-Region Standby for details.
2. If the switchover attempt in Step 1 fails, on the standby database the **Role** field shows a **Failover** link. On the standby database, click the **Failover** link.

   This shows the **Confirm manual failover to standby** dialog, along with information on possible data loss that may result if you perform the manual failover to the standby database.
3. In the **Confirm manual failover to standby** dialog, enter the Autonomous Database name to confirm that you want to failover.
4. In the **Confirm manual failover to standby** dialog, click **Confirm manual failover to standby**.

   When concurrent operations such as scaling are active, the confirmation also confirms either pausing or canceling the concurrent operation. See Concurrent Operations on Autonomous Database for more information.

See Notes for Manual Failover with a Standby Database for information on steps that Autonomous Data Guard takes after failover completes.

Notes for Manual Failover with a Standby Database

When the failover completes, Autonomous Data Guard performs post failover steps.

- For a failover to the local standby, Autonomous Data Guard creates a new local Standby database after the failover. Autonomous Data Guard is not enabled when the system is provisioning the new Standby database and the **Role** field shows **Provisioning**. After Autonomous Data Guard completes the provisioning step, then you have a new local Standby database and Autonomous Data Guard is enabled.
- For a failover to a cross-region standby, if there was a local standby before the failover, Autonomous Data Guard makes the local standby available again for a switchover or failover on the current Primary.

- After a manual failover operation completes, you can see any data loss associated with the manual failover in the message on the Oracle Cloud Infrastructure console banner. The manual failover data loss is specified in minutes.

  This information also shows when you hover over the  in the **Role** field.
- After a manual Autonomous Data Guard failover, if there was a regional failure, when the region comes back online the standby database is automatically reconnected or if required reprovisioned.

Autonomous Data Guard Cross-Region and Cross Tenancy Operations

Describes cross-region and cross tenancy operations with Autonomous Data Guard.

- Add a Cross-Region Standby Database
  You can enable Autonomous Data Guard with a cross-region standby when Autonomous Database is available (Lifecycle state shows Available).
- Enable or Disable Backup Replication for an Existing Cross-Region Standby
- Convert Cross-Region Peer to Snapshot Standby
- Disconnect a Peer Database
  Shows you the steps to disconnect a peer database from the Primary database.
- Disable a Cross-Region Standby Database
  Disabling a cross-region Autonomous Data Guard standby database terminates the standby database. If you later add a cross-region Autonomous Data Guard standby, the system creates a new cross-region standby database.
- Use a Cross Tenancy Autonomous Data Guard Standby Database
  You can add a cross tenancy Autonomous Data Guard standby database. This creates an Autonomous Data Guard standby that resides in a different tenancy. The cross tenancy standby can be in the same region (local) or in a different region (cross-region).
- Manage Remote Peer Network ACLs
  You can independently modify network ACLs on a remote disaster recovery peer database.

Add a Cross-Region Standby Database

You can enable Autonomous Data Guard with a cross-region standby when Autonomous Database is available (Lifecycle state shows Available).

To add a standby database you must have adequate available resources. Adding an Autonomous Data Guard standby database will only be successful if adding the standby database does not cause you to exceed your Tenancy or compartment limits for CPU and Storage.

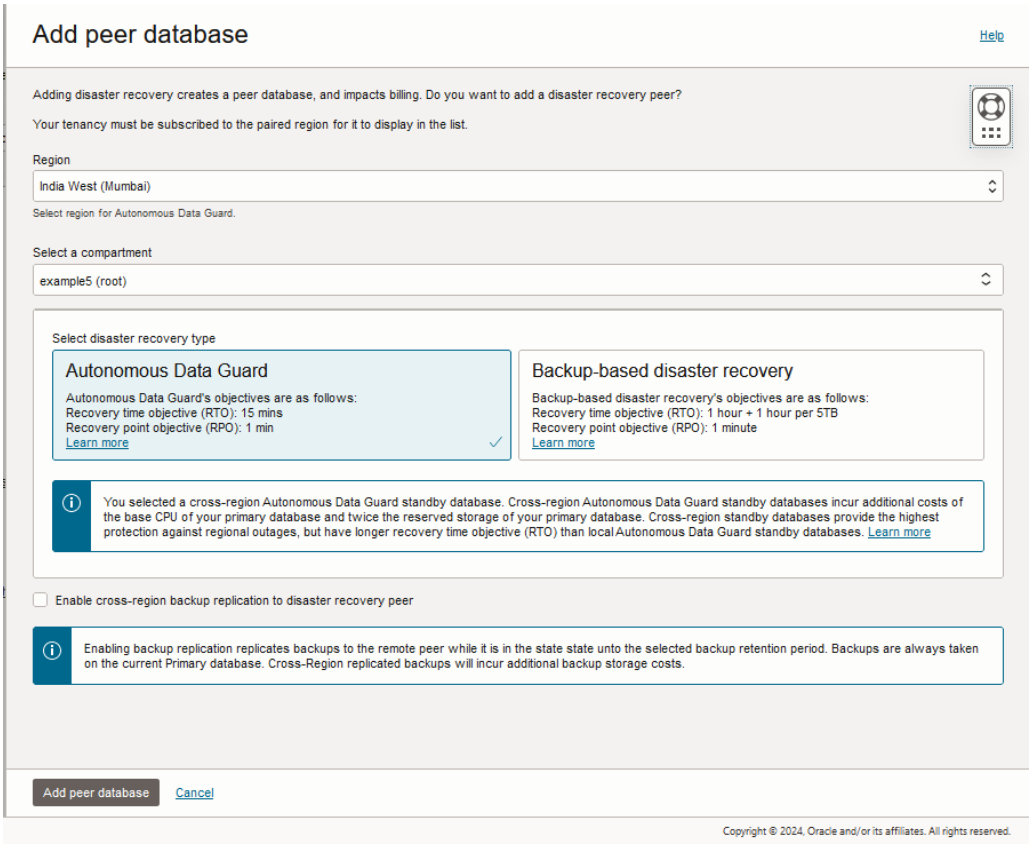
Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the  next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload, click one of: Autonomous Data Warehouse or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To add a cross-region Autonomous Data Guard standby:

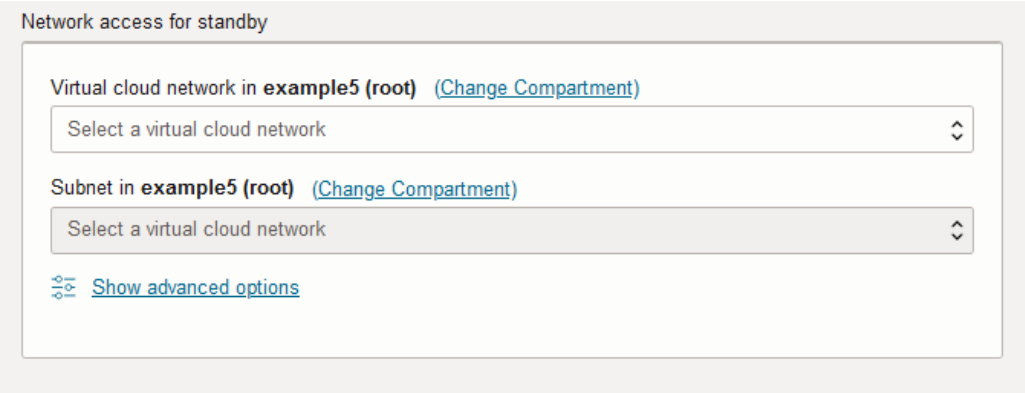
1. On the **Autonomous Database Details** page, under **Resources**, click **Disaster recovery**.
2. Under Disaster recovery, click **Add peer database**.
3. In the **Region** field, select a region.

   The region list shows the available remote regions where you can create a cross-region standby. When you add a standby database, the list of available regions only shows a remote region if your tenancy is subscribed to the remote region (you must be subscribed to the paired remote region). See Autonomous Database Cross-Region Paired Regions for more information.
4. When you select a region, Autonomous Database shows the **Select a compartment** list. From this list, choose a compartment for the cross-region standby.
5. Select the disaster recovery type. In addition, when the source database is configured with a private endpoint, enter the private endpoint information for the peer.
   a. Select the disaster recovery type: **Autonomous Data Guard**.



   b. If you want to enable cross-region backup replication, select **Enable cross-region backup replication to disaster recovery peer**.

      See Replicating Backups to a Cross-Region Autonomous Data Guard Standby for more information.
   c. When the source database is configured with a private endpoint, in the **Network access for standby** area enter the **Virtual cloud network** and the **Subnet**.



In these Network access for standby fields you specify the private endpoint's VCN and Subnet on the remote region where the standby is created.

> **Note:**
>
> If you change your network access on the source database to enable a private endpoint after the standby is created, you must manually access the standby to enable a private endpoint on the peer.

6. Click **Add peer database**.

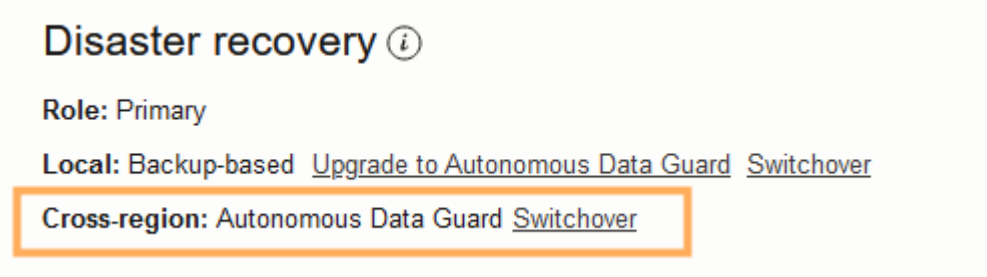
   The Autonomous Database **Lifecycle state** shows **Updating**. In the **Resources** area with **Disaster recovery** selected, the **State** field shows **Provisioning**.

   After some time the **Lifecycle state** shows **Available** and the standby database provisioning continues.

   > **Note:**
   >
   > While you add a standby database, the primary database is available for read/write operations. There is no downtime on the primary database.

When provisioning completes, on the **Autonomous Database Details** page under **Disaster recovery** the Oracle Cloud Infrastructure Console shows the following:

- **Role** shows **Primary**
- The **Local** field shows either **Backup-based** or **Autonomous Data Guard**. When the local field shows **Backup-based**, there is also a link, **Upgrade to Autonomous Data Guard**. Click this link to upgrade your local disaster recovery to Autonomous Data Guard.
- The **Cross-region** field shows:
  - **Autonomous Data Guard** when you have single cross-region standby database.
  - *x* peers: When you have multiple cross-region disaster recovery peers, where *x* is the number of cross-region peers.



When you add a cross-region standby, the standby database created in the remote region has the same display name as the primary database with an "_*region*" extension. Where *region* is the region name, such as `IAD` or `BOM`.

When you click **Disaster recovery** under **Resources**, the **Peer Autonomous Database** column shows the standby database name and a provides a link. Click the link to go to the Oracle Cloud Infrastructure Console for the remote standby database.

Notes for adding a cross-region standby database:

- Autonomous Database generates the Enable cross-region disaster recovery work request. To view the request, under **Resources** click **Work requests**.
- After you add a cross-region (remote) standby database, the wallet and connection string from the primary database will contain only the hostname of the primary database, and the wallet and connection string from the remote database will contain only the hostname of the remote database. This applies for both instance and regional wallets.

  See Cross-Region Disaster Recovery Connection Strings and Wallets for more information.
- If you select **Enable cross-region backup replication to disaster recovery peer**, it can take between several minutes and several hours to replicate the backups to the remote region, depending on the size of the backups. After backups are replicated, when you select **Backups** under **Resources** on the peer database's Oracle Cloud Infrastructure Console, you will see the list of replicated backups.
- While you add a standby database and the **Lifecycle state** shows **Updating**, the following actions are disabled for the primary database:
  - Move Resource. See Move an Autonomous Database to a Different Compartment for information on moving an instance.
  - Stop. See Stop Autonomous Database for information on stopping an instance.
  - Restart. See Restart Autonomous Database for information on restarting an instance.
  - Restore. See Restore and Recover your Autonomous Database for information on restoring.
  - Add cross-region disaster recovery peer. You need to wait until the update completes on the primary database before you can add an additional cross-region peer.
- See Cross-Region Autonomous Data Guard Notes and Notes for Customer-Managed Keys in OCI Vault with Autonomous Data Guard for information on using customer-managed keys and for additional notes for using Autonomous Data Guard with a cross-region standby.

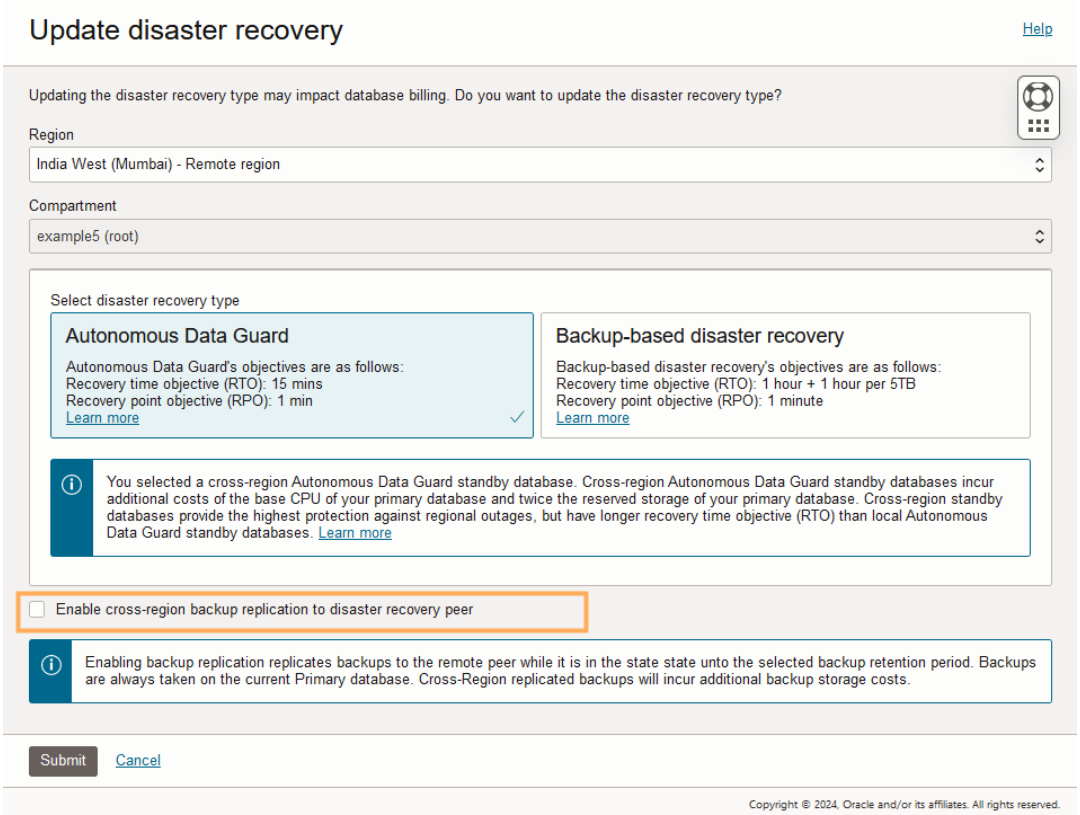
Enable or Disable Backup Replication for an Existing Cross-Region Standby

You can enable or disable backup replication on a Autonomous Data Guard cross-region standby.

To enable or disable backup replication for an existing cross-region Autonomous Data Guard standby:

1. On the **Autonomous Database Details** page, in the **Resources** area select **Disaster recovery**.
2. In a row that lists a cross-region standby, click  at the end of the row and select **Update disaster recovery**.

   This shows the Update disaster recovery page.



3. Enable or disable backup replication.
   a. If cross-region backup replication is disabled, select **Enable cross-region backup replication to disaster recovery peer** to enable the option.
   b. If cross-region backup replication is enabled, deselect **Enable cross-region backup replication to disaster recovery peer** to disable the option.
4. Click **Submit**.

   The Autonomous Database Lifecycle state changes to **Updating**.

If you select **Enable cross-region backup replication to disaster recovery peer**, it can take between several minutes and several hours to replicate the backups to the remote region, depending on the size of the backups. After backups are replicated, when you select **Backups** under **Resources** on the peer database's Oracle Cloud Infrastructure Console, you will see the list of replicated backups.

Convert Cross-Region Peer to Snapshot Standby

You can convert a cross-region peer database to a snapshot standby. This converts the peer to a read-write database for up to two days.

Snapshot standby CPU usage is billed based on the base CPU count and any additional CPU usage if compute auto scaling is enabled. The number of base CPUs is specified by the number of ECPUs (OCPUs if your database uses OCPUs) as shown in the **ECPU count** or **OCPU count** field on the Oracle Cloud Infrastructure Console.

Snapshot standby storage usage is billed based on the storage of the snapshot standby plus 1 x the storage of the source primary database.

You can create a snapshot standby for a cross-region peer. You cannot create a snapshot standby for a local disaster recovery peer.

- About Disaster Recovery Snapshot Standby Databases
  Converting a disaster recovery peer to a snapshot standby opens the database in read-write mode and the cross-region disaster recovery peer temporarily stops refreshing data from the source database.
- Convert Cross-Region Disaster Recovery Peer to a Snapshot Standby
  You can convert a cross-region disaster recovery peer to a snapshot standby.
- Convert Snapshot Standby Back to a Cross-Region Disaster Recovery Peer
  You can manually convert a snapshot standby back to a disaster recovery peer for the primary (source database). After the converting back, the snapshot standby returns to its role as a disaster recovery standby.
- Disconnect a Snapshot Standby
  You can disconnect a snapshot standby from the primary database.

About Disaster Recovery Snapshot Standby Databases

Converting a disaster recovery peer to a snapshot standby opens the database in read-write mode and the cross-region disaster recovery peer temporarily stops refreshing data from the source database.

While operating as a snapshot standby, updates from the source database are still sent to the snapshot standby, and you are protected if the source database region were to encounter a failure, however the updates are not applied to the snapshot standby until the database is converted back to a disaster recovery peer.

See Oracle Autonomous Database Serverless Features Billing for information on Snapshot standby billing.

- Snapshot Standby Features and Restrictions
  Provides information on snapshot standby features and restrictions.
- Snapshot Standby Operations
  After you create a snapshot standby you can perform almost all database operations on the snapshot standby. There are some operations that are not allowed on a snapshot standby.
- Snapshot Standby Reconnect Time
  A banner on the Oracle Cloud Infrastructure Console indicates the date and time when the snapshot standby automatically reconnects to the source database. At the time indicated on the banner, Autonomous Database converts a snapshot standby back to the **Standby** role.

Snapshot Standby Features and Restrictions

Provides information on snapshot standby features and restrictions.

While the database is in the **Snapshot standby** role, note the following:

- By converting a cross-region disaster recovery peer to a snapshot standby, you can use the snapshot standby for testing and querying data in the database. This allows you to test with no downtime on the primary (source) database, compared to testing by using switchover to the remote peer.
- You can use a snapshot standby to completely test your disaster recovery environment, including making any changes required to verify your standby environment, such as the mid-tier configuration. Using a snapshot standby you can make configuration changes or perform DML operations on the database as needed for complete testing and verification of the standby environment.

While the database is in the **Snapshot standby** role, note the following restrictions:

- The restore operation is not allowed on a snapshot standby database.
- No new backups are taken or displayed from the time a disaster recovery peer is converted to a snapshot standby. Existing backups that were available before the conversion to a snapshot standby are available. You can only use the backups that are available on a snapshot standby for a clone from backup operation.
- Cloning from a snapshot standby is only allowed to create a clone the same region as the snapshot standby. You cannot clone a snapshot standby across regions.

Notes for reconnecting a snapshot standby to the primary (source database):

- Reconnect to the primary, source database, when you are done with the tasks that require the snapshot standby to be open for read-write operations. If you do not manually reconnect within two days, the snapshot standby automatically reconnects to the primary.
- Oracle recommends you convert a snapshot standby back to a disaster recovery peer as soon as you are done with the operations that require the standby to be open for read-write operations. When you convert back to a to a disaster recovery peer, the accumulated changes from the source database are applied on the peer. If you keep the disaster recovery peer open as a snapshot standby for a longer period, assuming there are ongoing changes on the primary during this time, it takes longer to convert back to a disaster recovery peer.

When the snapshot standby reconnects to the primary database, Autonomous Database performs the following actions:

- The disaster recovery type you were using, and any associated billing returns to the type it was before you performed the convert disaster recovery peer to snapshot standby. That, the disaster recovery peer returns to the same type of disaster recovery peer, either Backup-Based Disaster Recovery (**Backup copy**) or **Autonomous Data Guard**, as shown in the DR Type column in the **Disaster Recovery** area.
- Any changes on the snapshot standby from the time it was converted to a snapshot standby, until the time it reconnects to the source are discarded. This means all changes, including metadata, that is inserted, updated, or deleted while the database operates as a snapshot standby are lost (discarded) when the snapshot standby reconnects to its source database.
- All changes that occurred on the primary are replicated to the remote region while the database operates as a snapshot standby, but the changes are not applied to the snapshot standby. The changes that occur on the primary during this period are applied to the snapshot standby when it is converted back to a disaster recovery peer.

- On the Oracle Cloud Infrastructure Console, the role updates from **Role: Snapshot standby** to **Role: Standby** or **Role: Backup copy**, depending on the disaster recovery type.

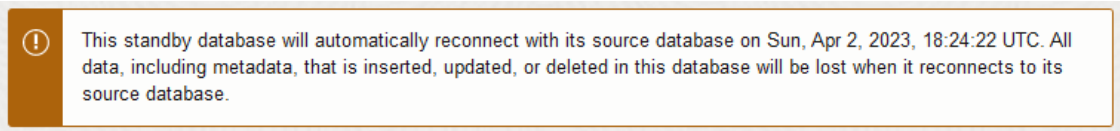
Snapshot Standby Operations

After you create a snapshot standby you can perform almost all database operations on the snapshot standby. There are some operations that are not allowed on a snapshot standby.

Operation	Description
Convert to Snapshot Standby	You can convert a cross-region peer to a snapshot standby. See Convert Cross-Region Disaster Recovery Peer to a Snapshot Standby for the steps to convert a peer database to snapshot standby.
Start or Restart	When a snapshot standby is stopped, as indicated by the Lifecycle state **Stopped**, you can start the database. When a snapshot standby is available, as indicated by the Lifecycle state **Available**, you can restart the database or stop the database.
Convert Snapshot Standby Back to Disaster Recovery Peer	When a snapshot standby is in the **Snapshot standby** role, the database operates as a read-write database. A snapshot standby has a two day (48 hour) limit up to which it can stay in the **Snapshot standby** role. If you do not manually convert the snapshot standby back within two days, the snapshot standby automatically converts back to a disaster recovery peer. See Convert Snapshot Standby Back to a Cross-Region Disaster Recovery Peer for more information.
Stop	When a snapshot standby is stopped, database operations are not available and charging for CPU usage on the snapshot standby stops.
Disconnect Peer	When you disconnect a snapshot standby, the snapshot standby is disassociated from the primary database. This converts the database from a snapshot database to a standalone database. Following the disconnect operation you are not allowed to reconnect to the Primary. See Disconnect a Snapshot Standby for more information.
Terminate	You are not allowed to terminate a snapshot standby. You can reconnect the snapshot standby to the primary. See Convert Snapshot Standby Back to a Cross-Region Disaster Recovery Peer for more information.
Create Clone	Cloning from a snapshot standby is only allowed to create a clone the same region as the snapshot standby. You cannot clone a snapshot standby across regions.
Create Refreshable Clone	You are not allowed to create a refreshable clone on a snapshot standby.
Disaster Recovery Peer	You are not allowed to add an Autonomous Data Guard standby database or a Backup-Based Disaster Recovery peer to a snapshot standby.

Snapshot Standby Reconnect Time

A banner on the Oracle Cloud Infrastructure Console indicates the date and time when the snapshot standby automatically reconnects to the source database. At the time indicated on the banner, Autonomous Database converts a snapshot standby back to the **Standby** role.



> **Note:**
>
> When a snapshot standby is not reconnected within 48 hours, the snapshot standby automatically reconnects to the source database.

Convert Cross-Region Disaster Recovery Peer to a Snapshot Standby

You can convert a cross-region disaster recovery peer to a snapshot standby.

> **Note:**
>
> All data, including metadata, that is inserted, updated, or deleted in the database during the disconnect period will be lost when a snapshot standby reconnects to its source database. All changes that occur on the primary during the disconnect period are applied to the standby when it reconnects to the source database.

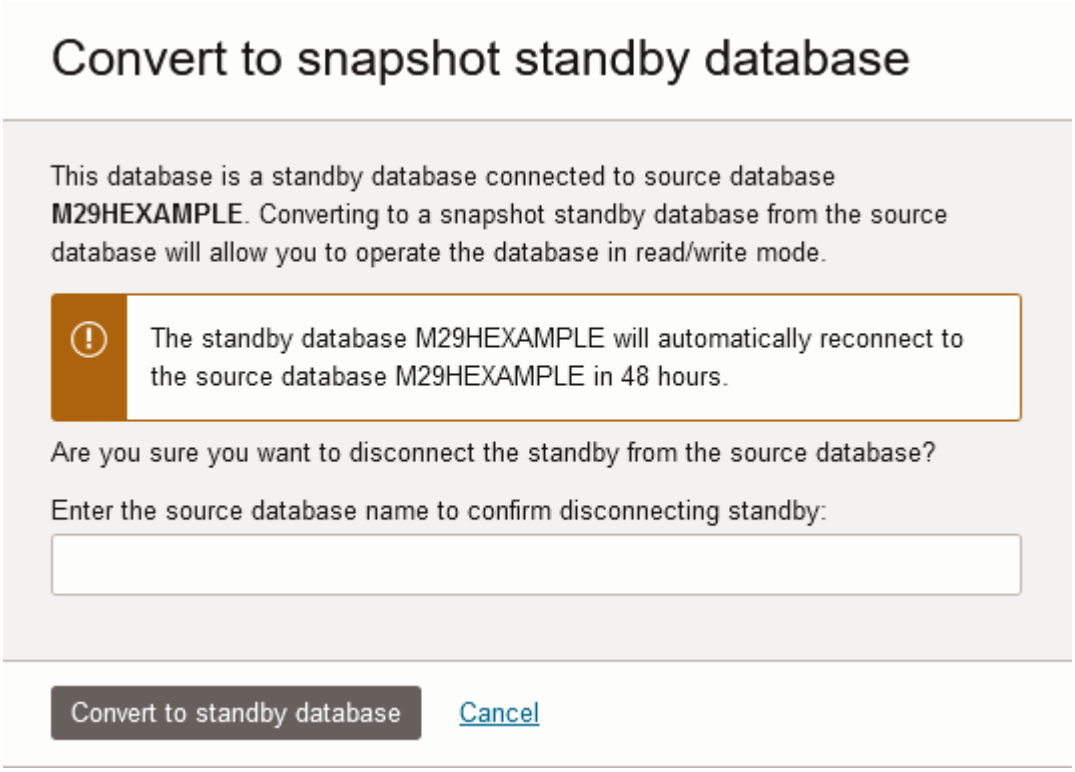
Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload, click one of: Autonomous Data Warehouse or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

1. On the Primary Region Autonomous Database instance, on the Autonomous Database Details page under **Resources**, select **Disaster recovery**.
2. Access the remote region peer.

   On the Primary Region Autonomous Database instance the **Disaster recovery** information area shows the **Peer Autonomous Database** column.

   Under **Peer Autonomous Database** column, click the link to access the cross-region peer.
3. On the cross-region peer, from the **More actions** drop-down list, select **Convert to snapshot standby database**.



4. On the **Convert to snapshot standby database** page, enter the source database name to confirm the disconnect.
5. Click **Convert to snapshot standby database**.

   The Autonomous Database Lifecycle state changes to **Updating**.

> **Note:**
>
> While you convert to a snapshot standby, the primary database is available for read/write operations. There is no downtime on the primary database.

When the operation completes, note the following:

- On the snapshot standby, there is a banner that indicates the date and time when the standby database will automatically reconnect with its source database
- On the snapshot standby, on the Autonomous Database details page, under **Disaster recovery**, the **Role** shows **Snapshot standby**.
- On the snapshot standby's source database, on the Autonomous Database details page, when you click **Disaster recovery** under **Resources**, the **Peer role** shows **Snapshot standby**.
- On the snapshot standby, you can scale CPU or storage, independent of the source database.
- When you scale CPU or storage on the primary, the changes to the primary do not affect the snapshot standby until it is converted back to a disaster recovery peer.
- Cloning is only allowed from a snapshot standby when the clone is in the same region. Cloning across regions is not allowed from a snapshot standby.

- The clone from backup operation is allowed on a snapshot standby. However, a snapshot standby does not backup any of the temporarily updated data in the read-write snapshot standby. When you clone from a backup on a snapshot standby, the operation creates a clone using a Primary side backup which is replicated to the snapshot standby. See Clone an Autonomous Database Instance for more information.
- The restore operation is not allowed on a snapshot standby.

Convert Snapshot Standby Back to a Cross-Region Disaster Recovery Peer

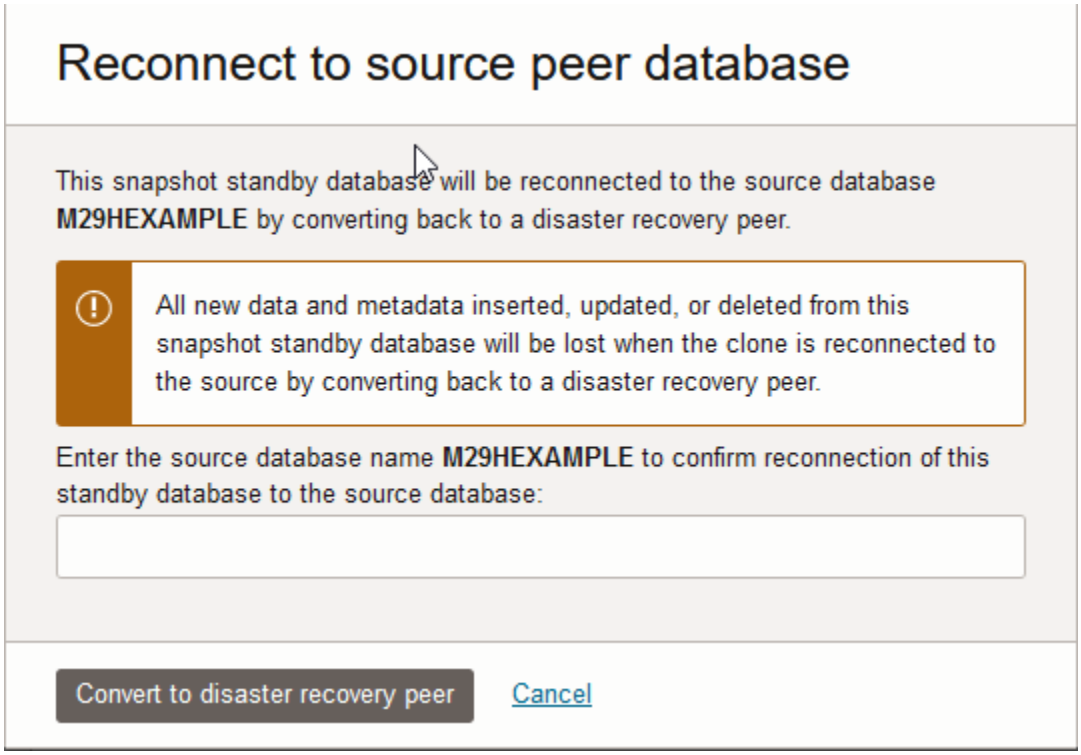
You can manually convert a snapshot standby back to a disaster recovery peer for the primary (source database). After the converting back, the snapshot standby returns to its role as a disaster recovery standby.

> **Note:**
>
> All data, including metadata, that is inserted, updated, or deleted in the snapshot standby database during the disconnect period will be lost when the snapshot standby reconnects to its source database.
> All changes on the primary that were sent to the snapshot standby but not applied during the disconnect period will be applied to the standby when it reconnects to the source database.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload, click one of: Autonomous Data Warehouse or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

1. On the snapshot standby, from the **More actions** drop-down list, select **Reconnect to source peer database**.



2. On the **Reconnect to source peer database** page, enter the source database name to confirm the reconnect.
3. Click **Convert to disaster recovery peer**.

   The Autonomous Database Lifecycle state changes to **Updating**.

> **Note:**
>
> While you reconnect the standby to the source database, the primary (source) database is available for read/write operations. There is no downtime on the primary database.

When the snapshot standby reconnects to the primary database, Autonomous Database does the following:

- The disaster recovery type you were using, and any associated billing returns to the type it was before you performed the convert disaster recovery peer to snapshot standby. That, the disaster recovery peer returns to the same type of disaster recovery peer, either Backup-Based Disaster Recovery (**Backup copy**) or **Autonomous Data Guard**, as shown in the DR Type column in the **Disaster Recovery** area.
- Any changes on the snapshot standby from the time it was converted to a snapshot standby, until the time it reconnects to the source are discarded. This means all changes, including metadata, that is inserted, updated, or deleted while the database operates as a snapshot standby are lost (discarded) when the snapshot standby reconnects to its source database.
- All changes that occurred on the primary are replicated to the remote region while the database operates as a snapshot standby, but the changes are not applied to the snapshot standby. The changes that occur on the primary during this period are applied to the snapshot standby when it is converted back to a disaster recovery peer.

- On the Oracle Cloud Infrastructure Console, the role updates from **Role: Snapshot standby** to **Role: Standby** or **Role: Backup copy**, depending on the disaster recovery type.

Disconnect a Snapshot Standby

You can disconnect a snapshot standby from the primary database.

When you disconnect a snapshot standby, the snapshot standby is disassociated from the primary database. This converts the database from a snapshot database to a standalone database. Following the disconnect operation you are not allowed to reconnect to the Primary.

The steps to disconnect a snapshot standby are the same as those to disconnect a standby database. See Disconnect a Peer Database for more information.

Notes for disconnecting a snapshot standby.

- There is no reconnect operation. After you disconnect a snapshot standby you are not allowed to reconnect to the Primary.
- The disconnect operation for a snapshot standby can only be performed on an Autonomous Database instance that uses the ECPU compute model.
- The disconnected database keeps any user inserted or updated data that was applied while the database was open in read/write mode as a Snapshot Standby.

  The disconnect operation does not apply any recent logs that were sent from the Primary.
- After the disconnect operation, the standalone database is no longer associated with the database that was the Primary database. To use the database as a standalone database, you must know the name of the database that was disconnected from the Primary database.
- After the disconnect operation the database starts taking new backups as a standalone database. Any backups that were associated with the standby database or with the primary database are not available on the standalone database.

Disconnect a Peer Database

Shows you the steps to disconnect a peer database from the Primary database.

When you disconnect a cross-region peer, the database is permanently disassociated from the Primary database. This converts the database from a peer database to a standalone database. After you disconnect a peer database you are not allowed to reconnect to the Primary.

Perform the following prerequisite steps as necessary:

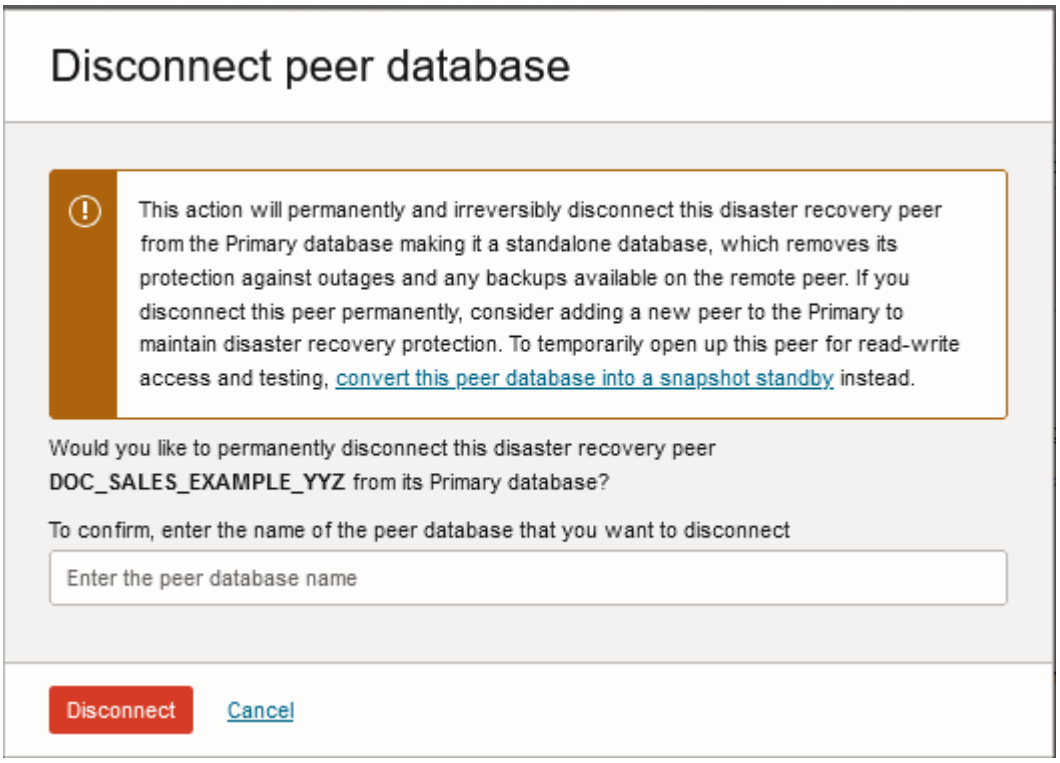
- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

You can perform this operation from the Primary database or from a cross-region peer.

To perform the disconnect operation on a cross-region standby (or on a cross-region backup copy):

1. On the cross-region peer, on the Autonomous Database details page, under **Disaster recovery** in the **Role** field click **Disconnect peer**.

This shows the disconnect peer dialog:



2. On the **Disconnect peer database** dialog, enter the peer database name to confirm the disconnect.
3. Click **Disconnect**.

   The Autonomous Database Lifecycle state changes to **Updating**.

To perform the disconnect operation on the Primary:

1. On the Primary database, on the Autonomous Database Details page, in the **Resources** area click **Disaster recovery**.
2. At the end of the row for the peer you want to disconnect, click  and select **Disconnect peer**.
3. On the **Disconnect peer database** dialog, enter the peer database name to confirm the disconnect.
4. Click **Disconnect**.
   The Autonomous Database Lifecycle state changes to **Updating**.

Notes for disconnecting a cross-region peer.

- The disconnect operation for a cross-region peer can only be performed on an Autonomous Database instance that uses the ECPU compute model.
- After the disconnect operation, the standalone database is no longer associated with the database that was the Primary database. To use the database as a standalone database, you must know the name of the database that was disconnected from the Primary database.

- There is no reconnect operation. After you disconnect a peer database you are not allowed to reconnect to the Primary.
- After the disconnect operation, the standalone database starts taking new backups as a standalone database.
- For a cross tenancy Autonomous Data Guard standby database, you can perform the disconnect operation on a local, same region, standby database, a cross-region standby database, or a snapshot standby database.

Disable a Cross-Region Standby Database

Disabling a cross-region Autonomous Data Guard standby database terminates the standby database. If you later add a cross-region Autonomous Data Guard standby, the system creates a new cross-region standby database.

You have two options to disable a cross-region Autonomous Data Guard standby:

You can update the cross-region disaster recovery type to use Backup-Based Disaster Recovery. This terminates the cross-region Autonomous Data Guard standby database and adds a cross-region backup copy peer.

You can terminate the cross-region Standby database.

- Update Cross-Region Standby to Use Backup-Based Disaster Recovery
  You can update the disaster recovery type from cross-region Autonomous Data Guard standby database to cross-region Backup-Based Disaster Recovery. This terminates the cross-region Autonomous Data Guard standby database.
- Terminate a Cross-Region Standby Database
  Describes the steps to terminate a cross-region standby database.
- Verify Cross-Region Autonomous Data Guard is Disabled
  Describes the steps to verify that an Autonomous Data Guard cross-region standby is disabled.

Update Cross-Region Standby to Use Backup-Based Disaster Recovery

You can update the disaster recovery type from cross-region Autonomous Data Guard standby database to cross-region Backup-Based Disaster Recovery. This terminates the cross-region Autonomous Data Guard standby database.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload, click one of: Autonomous Data Warehouse or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To change your disaster recovery type to use a cross-region backup-copy peer:

1. On the Primary Role Autonomous Database, on the Autonomous Database Details page under **Resources**, select **Disaster recovery**.
2. Access the remote standby.

   On the primary Autonomous Database instance the **Disaster recovery** information area shows the remote standby. The remote standby has the same display name as the primary

database, with a "*_region*" extension. Where *region* is the region name, such as `IAD` or `BOM`.

If you created the cross-region standby before the introduction of support for multiple cross-region peers, the remote standby has the same display name as the primary database, with an "`_Remote`" extension.

Under **Peer Autonomous Database**, click the link to access the cross-region standby.

3. On the remote standby database, click **Update disaster recovery** under **Disaster recovery** on the Oracle Cloud Infrastructure Console.
4. On the Update disaster recovery type page, select **Backup-based disaster recovery**.
5. Click **Submit**.

   The Autonomous Database Lifecycle state changes to **Updating**.

> **Note:**
>
> While you update the disaster recovery type, the primary database is available for read/write operations. There is no downtime on the primary database.

Autonomous Database generates a Change Disaster Recovery Configuration work request. To view the request, under **Resources** click **Work Requests**.

Terminate a Cross-Region Standby Database

Describes the steps to terminate a cross-region standby database.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload, click one of: Autonomous Data Warehouse or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

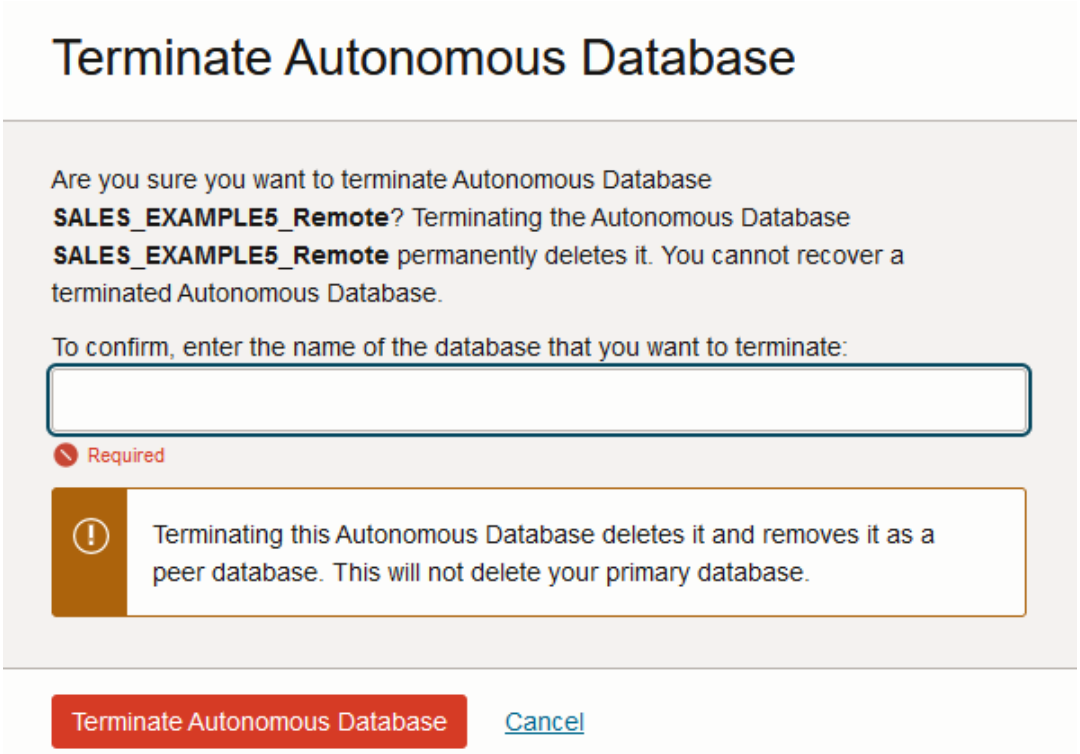
To terminate a cross-region standby database:

1. On the primary database, on the Autonomous Database Details page, under **Resources** select **Disaster recovery**.
2. Access the remote standby.

   The **Disaster recovery** information area shows the **Peer Autonomous Database**. The remote standby database has the same display name as the primary database, with an "*_region*" extension. Where *region* is the region name, such as `IAD` or `BOM`.

   If you created the cross-region peer before the introduction of support for multiple cross-region peers, the remote standby database has the same display name as the primary database, with an "`_Remote`" extension.

   Under **Peer Autonomous Database**, click the link to access the cross-region peer.
3. On the console for the remote standby database, from the **More actions** drop-down list, select **Terminate**.



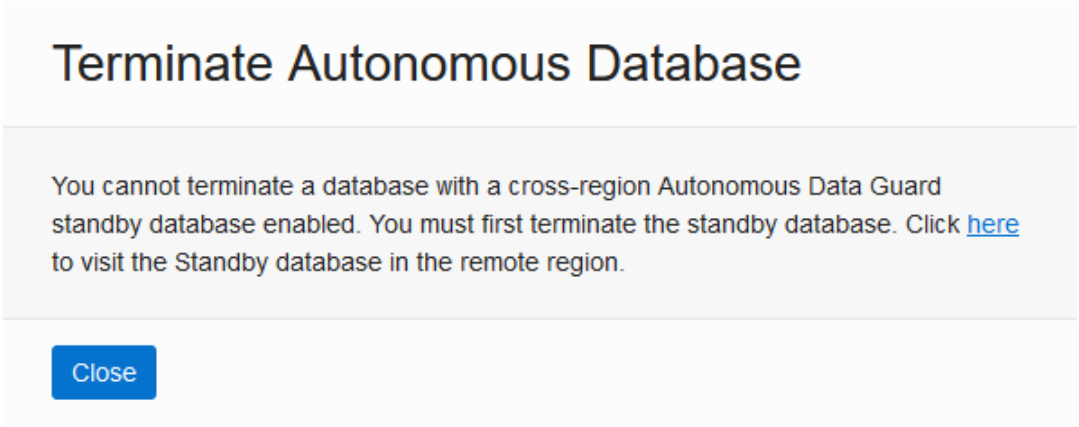
4. On the **Terminate Autonomous Database** page enter the database name to confirm that you want to terminate the cross-region standby database.
5. Click **Terminate Autonomous Database**.

   While the standby database is terminating, the **Lifecycle state** changes to **Terminating**.

> **Note:**
>
> If there is a local Autonomous Data Guard standby database, terminating the cross-region standby does not change the local disaster recovery option.

Note the following when your disaster recovery option includes a cross-region disaster recovery peer:

- A cross-region disaster recovery peer cannot be terminated from the primary database.
- You must terminate all cross-region disaster recovery peers before you terminate the primary database. If you attempt to terminate the primary when there is a cross-region disaster recovery peer, the system shows the following message:



In this case, first terminate the cross-region (remote) peers and then you can terminate the Primary role database.

After the cross-region disaster recovery peers are terminated, then terminate the Primary database. See Terminate an Autonomous Database Instance for more information.

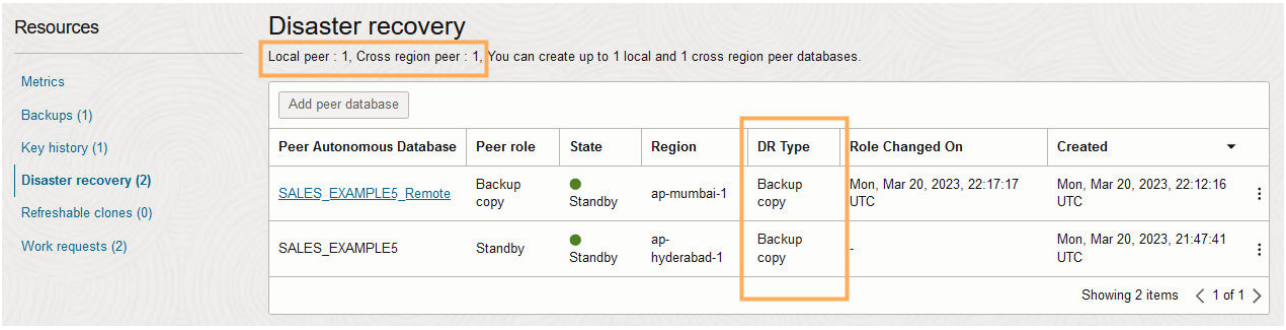
Verify Cross-Region Autonomous Data Guard is Disabled

Describes the steps to verify that an Autonomous Data Guard cross-region standby is disabled.

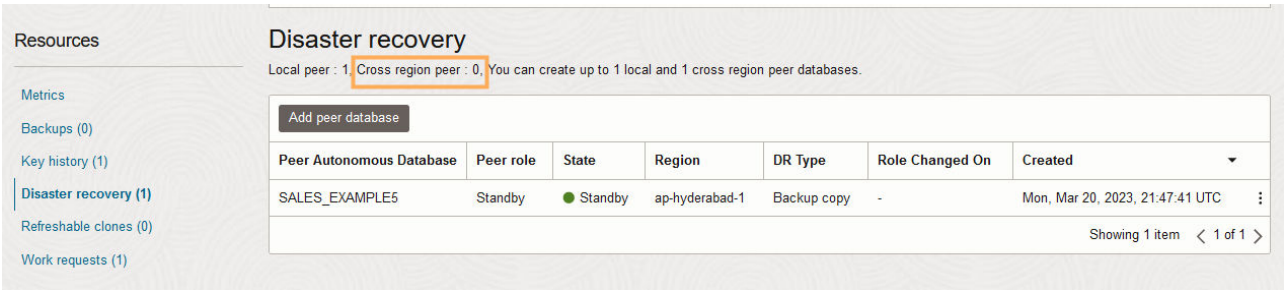
To verify that a cross-region Autonomous Data Guard standby database is disabled:

1. On the primary database, on the Autonomous Database Details page under **Resources**, select **Disaster recovery**.
2. Verify that either the disaster recovery type is changed to backup copy or the cross-region Autonomous Data Guard standby is terminated.

- If you updated the disaster recovery type to use cross-region backup based disaster recovery, verify that the remote peer autonomous database **DR Type** shows **backup copy**.



- If you manually terminated standby database, verify that the cross-region standby count is 0 (that is there is no remote peer).



Use a Cross Tenancy Autonomous Data Guard Standby Database

You can add a cross tenancy Autonomous Data Guard standby database. This creates an Autonomous Data Guard standby that resides in a different tenancy. The cross tenancy standby can be in the same region (local) or in a different region (cross-region).

- About Cross Tenancy Autonomous Data Guard
  When you add a cross tenancy Autonomous Data Guard standby, Autonomous Database provisions a standby database in a remote tenancy.
- Prerequisites to Create a Cross Tenancy Autonomous Data Guard Standby
  Describes prerequisites for creating a cross tenancy Autonomous Data Guard standby.
- Add a Cross Tenancy Autonomous Data Guard Standby
  Shows the steps to create a cross tenancy Autonomous Data Guard standby. This covers creating a cross tenancy standby when the source database and the standby database are in the same region (local) and when the source database and the standby database are in different regions (cross-region).
- Switchover, Failover, or Create a Snapshot Standby with a Cross Tenancy Autonomous Data Guard Standby
  Describes the procedure to switchover, failover, or create a snapshot standby with a cross tenancy standby.
- Terminate a Cross Tenancy Autonomous Data Guard Standby
  Shows the steps to terminate a cross tenancy standby.

About Cross Tenancy Autonomous Data Guard

When you add a cross tenancy Autonomous Data Guard standby, Autonomous Database provisions a standby database in a remote tenancy.

With a cross tenancy Autonomous Data Guard standby you can failover or switchover to a standby database in a different tenancy (or create a snapshot standby in a different tenancy). This feature also allows you to easily use an Autonomous Data Guard standby to migrate a database to a different tenancy.

Note the following for adding a cross tenancy Autonomous Data Guard standby:

- You can only create a cross tenancy Autonomous Data Guard standby using the CLI or the Autonomous Database REST APIs. This option is not available using the Oracle Cloud Infrastructure Console.
- Some of the features available when you enable an Autonomous Data Guard standby in the same tenancy, such as the automatic failover options, are not available with a cross tenancy Autonomous Data Guard standby.

  There are no automatic failover options for a cross tenancy Autonomous Data Guard standby and you can only use manual failover to failover to a cross tenancy standby.
- You can enable a cross tenancy Autonomous Database standby with either the ECPU or OCPU compute model. The standby database uses the same compute model as the primary database.
- A cross tenancy Autonomous Data Guard standby can be in the same region (local) or in a different region (cross-region).
- Only a single cross tenancy Autonomous Data Guard standby can be enabled, either local or cross-region. You cannot enable both a cross tenancy local and a cross tenancy cross-region Autonomous Data Guard standby.

  - In addition to a single cross tenancy Autonomous Data Guard standby, you can have one local or cross-region disaster recovery peer (either Autonomous Data Guard or Backup-Based Disaster Recovery).
- Customer-Managed Encryption Keys are not supported with a Cross Tenancy Autonomous Data Guard standby.

Prerequisites to Create a Cross Tenancy Autonomous Data Guard Standby

Describes prerequisites for creating a cross tenancy Autonomous Data Guard standby.

You must define OCI Identity and Access Management policies to authorize creating an Autonomous Data Guard standby in a different tenancy.

The OCI Identity and Access Management groups and policies you add support the following:

- A member of a group in the source tenancy allows a group in the destination tenancy to access (read) the source Autonomous Database instance on the source tenancy.

  You do not need to allow other actions on the source Autonomous Database instance (for example, start, stop terminate, or any write operations).
- A member of a group in the destination tenancy is allowed to read Autonomous Databases in the source tenancy.

  On the destination tenancy you also add a policy that allows a group to manage the Autonomous Databases in a compartment in the remote (destination) tenancy. For example, this policy allows the group to create the standby database.

1. Create a group on the remote tenancy that contains the user that will be allowed to add the Autonomous Data Guard standby.
   a. On the remote tenancy, in the Oracle Cloud Infrastructure Console click **Identity & Security**.
   b. Under **Identity** click **Domains** and select an identity domain (or create a new identity domain).
   c. Under **Identity domain**, click **Groups**.
   d. To add a group, click **Create group**.
   e. On the Create group page, enter a Name and a Description.

      For example, enter the Name: **CrossTenancyStandbyGroup**.
   f. On the Create group page, click **Create**.
   g. Click **Create** to save the group.
   h. On the **Group** page, click **Assign user to groups** and select the users you want to add to the group.
   i. Click **Add**.
   j. On the **Group** page, from the **Group information** tab copy the OCID for use in Step 2.
2. On the source tenancy, define OCI Identity and Access Management policies for the source Autonomous Database instance.
   a. On the source tenancy, in the Oracle Cloud Infrastructure Console click **Identity & Security**.
   b. Under **Identity**, click **Policies**.
   c. To write a policy, click **Create Policy**.
   d. On the Create Policy page enter a Name and a Description.

e. On the Create Policy page, select **Show manual editor**.

f. In the policy builder, add policies so that the group in the destination tenancy is allowed to create a standby using an Autonomous Database instance on the source tenancy as the source.

For example, define the following generic policies:

```
define tenancy RemoteTenancy as ocid1.tenancy.oc1..unique_id
define group GroupThatExistsInRemoteTenancy as
ocid1.group.region1..unique_id
admit group GroupThatExistsInRemoteTenancy of tenancy RemoteTenancy to
read autonomous-databases
  in compartment id ocid1.compartment.region..unique_id
```

This policy specifies the following:

- Line 1: the OCID is the OCID of the remote tenancy. This is the tenancy where you are going to create the standby.
- Line 2: the OCID is the OCID of the group to which the user who will create the standby belongs. This is the OCID for the group you created in Step 1.
- Line 3: The OCID is the OCID of the compartment where the source database resides.

Policies on Primary (Source) Tenancy To Add a Cross Tenancy Standby

```
define tenancy RemoteTenancy as ocid1.tenancy.oc1..aaa_example_rcyx2a
define group CrossTenancyStandbyGroup as
ocid1.group.oc1..aaa_example_6vctn6xsaq
admit group CrossTenancyStandbyGroup of tenancy RemoteTenancy to read
autonomous-databases
   in compartment id ocid1.compartment.oc1..aaaaaaaabexample6q
```

This policy specifies the `CrossTenancyStandbyGroup` group of the `RemoteTenancy` can read from Autonomous Databases in the specified compartment (on the source tenancy). To create a cross tenancy standby, the policy only needs to allow read on the source Autonomous Database instance.

g. Click **Create** to save the policy.

3. Define policies for the remote tenancy.

a. On the destination tenancy, in the Oracle Cloud Infrastructure Console click **Identity & Security**.

b. Under **Identity**, click **Policies**.

c. To write a policy, click **Create Policy**.

d. On the Create Policy page enter a Name and a Description.

e. On the Create Policy page, select **Show manual editor**.

f. In the policy builder, add policies so that a group is endorsed to manage Autonomous Databases on the source tenancy.

For example:

```
Define tenancy PrimaryTenancy as ocid1.tenancy.oc1..aaaaaaaabexampleyx2a
Endorse group CrossTenancyStandbyGroup to read autonomous-databases in
```

```
tenancy PrimaryTenancy
Allow group CrossTenancyStandbyGroup to manage autonomous-databases in
tenancy
```

This policy specifies the following:

- Line 1: The OCID is the primary tenancy OCID. This is the tenancy where the primary Autonomous Database instance resides.
- Line 2: Specifies that the CrossTenancyStandbyGroup group is endorsed in their own tenancy to read the source Autonomous Database in the primary tenancy.
- Line 3: Specifies that the CrossTenancyStandbyGroup group is allowed to create/manage an Autonomous Database in the specified Tenancy (the tenancy where the Standby is going to be created).

> **Note:**
>
> If these polices are revoked, creating a cross tenancy Autonomous Data Guard Standby is no longer allowed.

See Getting Started with Policies for more information.

Add a Cross Tenancy Autonomous Data Guard Standby

Shows the steps to create a cross tenancy Autonomous Data Guard standby. This covers creating a cross tenancy standby when the source database and the standby database are in the same region (local) and when the source database and the standby database are in different regions (cross-region).

You can only create a cross tenancy Autonomous Data Guard standby using the CLI or the Autonomous Database REST APIs. This option is not available using the Oracle Cloud Infrastructure Console.

First, perform the prerequisite steps to define the OCI Identity and Access Management policies to authorize cross tenancy Autonomous Data Guard.

See Prerequisites to Create a Cross Tenancy Autonomous Data Guard Standby for details.

To create a local (same region) cross tenancy Autonomous Data Guard standby:

The API call to create the standby is sent to the different tenancy in the local region.

1. Copy the OCID of the Primary database (on the primary's Tenancy).
2. Access the tenancy where you want to create the standby, the remote tenancy, and use the CLI or the REST API to create the standby database.

   For example, with the CLI:

   ```
   oci db autonomous-database create-cross-tenancy-disaster-recovery-details
   --compartment-id ocid1.tenancy.oc1..fcue4_example --source-id
   ocid1.autonomousdatabase.oc1.ap-mumbai-1.anrgexamples3q  --disaster-
   recovery-type ADG
   ```

   See update for more information.

   See UpdateAutonomousDatabase for more information.

To create a remote (cross-region) cross tenancy Autonomous Data Guard standby:

The API call to create the standby runs in the different tenancy in the source region.

1. Copy the OCID of the Primary database (on the primary's Tenancy).
2. Access the tenancy where you want to create the standby, the remote tenancy in the remote region, and use the CLI or the REST API to create the standby database.

   For example, with the CLI:

   ```
   oci db autonomous-database create-adb-cross-region-data-guard-details
    --compartment-id ocid1.tenancy.oc1..fcue4_example --source-id
   ocid1.autonomousdatabase.oc1.ap-mumbai-1.anrg6example2vws3q --disaster-
   recovery-type ADG
   ```

   See update for more information.

   See UpdateAutonomousDatabase for more information.

Notes for adding a cross tenancy standby database:

- Autonomous Database generates the Enable cross-region disaster recovery work request. To view the request, under **Resources** click **Work requests**.
- After you add a cross tenancy standby database either local or cross-region, the wallet and connection string from the primary database will contain only the hostname of the primary database, and the wallet and connection string from the remote tenancy database will contain only the hostname of the remote tenancy database. This applies for both instance and regional wallets.

  See Cross-Region Disaster Recovery Connection Strings and Wallets for more information.
- While you add a standby database and the **Lifecycle state** shows **Updating**, the following actions are disabled for the primary database:
  - Move Resource. See Move an Autonomous Database to a Different Compartment for information on moving an instance.
  - Stop. See Stop Autonomous Database for information on stopping an instance.
  - Restart. See Restart Autonomous Database for information on restarting an instance.
  - Restore. See Restore and Recover your Autonomous Database for information on restoring.
- See Cross-Region Autonomous Data Guard Notes and Notes for Customer-Managed Keys in OCI Vault with Autonomous Data Guard for information on using customer-managed keys and for additional notes for using Autonomous Data Guard with the standby in a different tenancy.

Switchover, Failover, or Create a Snapshot Standby with a Cross Tenancy Autonomous Data Guard Standby

Describes the procedure to switchover, failover, or create a snapshot standby with a cross tenancy standby.

- Switchover: The steps to perform a switchover to a cross tenancy standby are the same whether the standby is local (same region) or in a different region (cross-region). You follow the same steps as for a switchover to a same tenancy cross-region standby. See Perform a Switchover to a Cross-Region Standby for more information.

- Failover: The steps to perform a failover to a cross tenancy standby are the same whether the standby is local (same region) or in a different region (cross-region). You follow the same steps as for a manual failover to a same tenancy cross-region standby. See Perform Manual Failover to a Cross-Region Standby Database for more information.
- Create a Snapshot Standby: The steps to create a snapshot standby with a cross tenancy standby are the same whether the cross tenancy standby is local (same region) or in a different region (cross-region). You follow the same steps as with a same tenancy cross-region standby. See Convert Cross-Region Disaster Recovery Peer to a Snapshot Standby for more information.

Terminate a Cross Tenancy Autonomous Data Guard Standby

Shows the steps to terminate a cross tenancy standby.

The steps to terminate either a same region or a cross-region cross tenancy standby are the same as those to terminate a cross-region same-tenancy standby. See Terminate a Cross-Region Standby Database for more information.

Manage Remote Peer Network ACLs

You can independently modify network ACLs on a remote disaster recovery peer database.

By default the disaster recovery primary and remote peer databases use the same network Access Control Lists (ACLs). Optionally, you can configure ACLs independently on remote peer databases. This provides an option to use different ACLs on remote peer databases.

If you modify the ACLs on a remote peer, Autonomous Database no longer keeps the ACL configuration synchronized between the primary and the remote peer. After you modify the ACLs on a remote peer, Autonomous Database manages the ACLs on the remote peer independently.

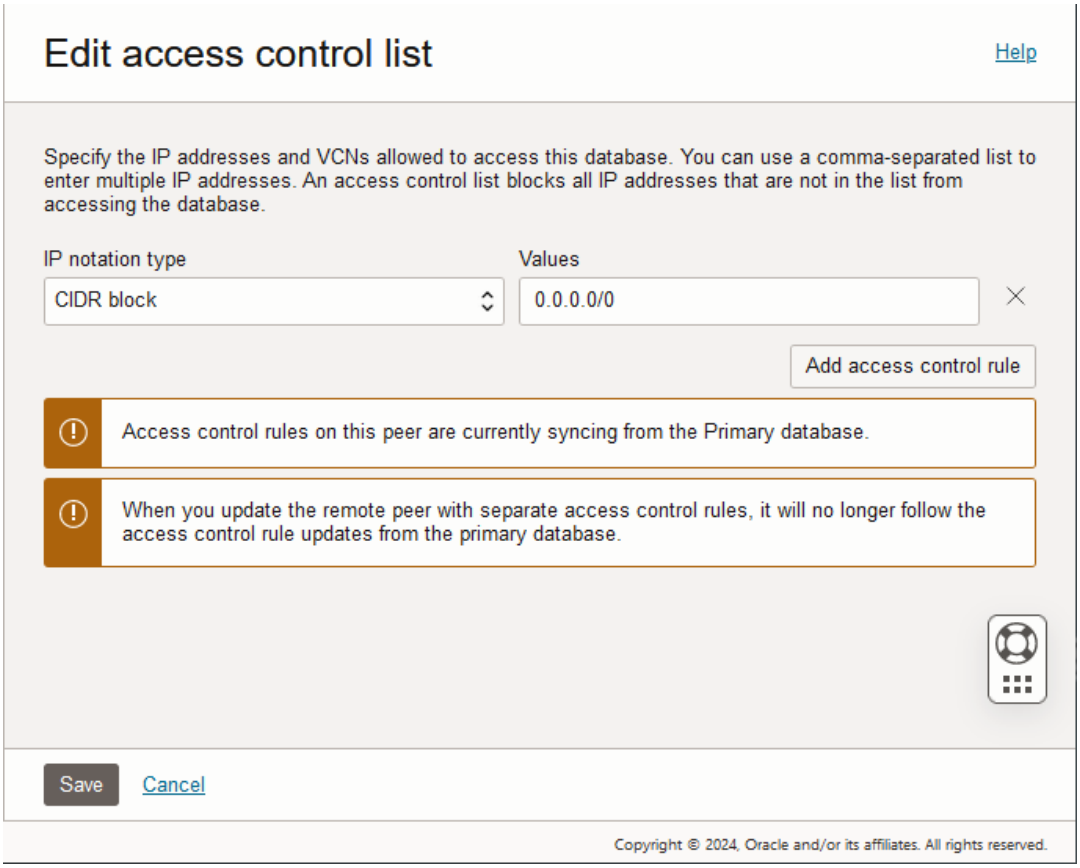
To use different network ACLs on a remote Autonomous Database peer:

1. On the primary database, on the Autonomous Database Details page, under **Resources** select **Disaster recovery**.
2. Access the remote peer.

   The **Disaster recovery** information area shows the **Peer Autonomous Database**. The remote peer database by default has the same display name as the primary database, with an "_*region*" extension. Where *region* is the region name, such as `IAD` or `BOM`.

   Under **Peer Autonomous Database**, click the link to access a cross-region peer.
3. On the remote peer database, edit the access control list.

   Before you change ACL values the dialog shows a message indicating that ACLs on the peer database are syncing from the Primary database. For example:



See Configure Access Control Lists for an Existing Autonomous Database Instance for more information.

4. Add, remove, or modify one or more ACLs.
5. Click **Save**.

After you modify ACLs, the ACLs on the primary and on the remote peer are managed separately.

If you want to restart the synchronization of ACLs between the primary and the remote peer, you have two options:

- Terminate the peer Autonomous Database and create a new cross-region disaster recovery peer database.

  See Disable a Cross-Region Standby Database for details on terminating a remote standby database.

  See Disable a Cross-Region (Remote) Peer for details on terminating a remote peer.
- Contact Oracle Cloud Support and file a service request or contact your support representative.

Events and Notifications for a Standby Database

You can use Oracle Cloud Infrastructure Events to be notified of and to specify rules to automatically respond to Autonomous Data Guard operations.

> **Note:**
>
> When Autonomous Data Guard is enabled with a remote standby, events are only sent to the instance with the **Primary** role.

Event Description	Type
Automatic Failover Begin	`com.oraclecloud.databaseservice.automaticfailover.begin`
Automatic Failover End	`com.oraclecloud.databaseservice.automaticfailover.end`
Disable Data Guard Begin	`com.oraclecloud.databaseservice.disableautonomousdataguard.begin`
Disable Data Guard End	`com.oraclecloud.databaseservice.disableautonomousdataguard.end`
Enable Data Guard Begin	`com.oraclecloud.databaseservice.enableautonomousdataguard.begin`
Enable Data Guard End	`com.oraclecloud.databaseservice.enableautonomousdataguard.end`
Manual Failover Begin	`com.oraclecloud.databaseservice.failoverautonomousdatabase.begin`
Manual Failover End	`com.oraclecloud.databaseservice.failoverautonomousdatabase.end`
Switchover Begin	`com.oraclecloud.databaseservice.switchoverautonomousdatabase.begin`
Switchover End	`com.oraclecloud.databaseservice.switchoverautonomousdatabase.end`

See the following Oracle Cloud Infrastructure topics for information on events and notifications:

- Overview of Events
- Autonomous Database Event Types
- Notifications Overview

Autonomous Database Cross-Region Paired Regions

Autonomous Database provides paired regions for creating cross-region databases, including cross-region clones and cross-region disaster recovery peers.
The **Region** list shows paired regions when you add a remote disaster recovery peer or when you create a clone in a different region. The **Region** list only shows paired regions. That is, if your tenancy is not subscribed to a region that is paired to the source region, the region is not displayed. To subscribe to a region, see Managing Regions.

If you do not see a region that you are subscribed to in the **Region** list, open a support ticket to request that the region be added as a paired region for your source region. See Get Help, Search Forums, and Contact Support for more information.

Use the API

Provides links for details on using API operations to manage Autonomous Data Guard.

For information about using the API and signing requests, see REST APIs and Security Credentials.

For information about SDKs, see Software Development Kits and Command Line Interface.

Use these API operations to manage Autonomous Data Guard:

- To enable or disable Autonomous Data Guard, use UpdateAutonomousDatabase.
- To initiate a manual failover operation, use FailOverAutonomousDatabase.
- To initiate a switchover operation, use SwitchOverAutonomousDatabase.

Use these Terraform APIs to manage Autonomous Database resources:

For information about Terraform, see Terraform Provider and for information about Terraform APIs, see Data Source: oci_database_autonomous_database.

Autonomous Data Guard Notes

Provides notes for using Autonomous Database with an Autonomous Data Guard standby database.

- You cannot connect to a Standby database until it is made the Primary by a failover or a switchover. Thus, a Standby database cannot be opened for read-only access and cannot be used to offload queries from a Primary database.
- Autonomous Data Guard is available with the Data Warehouse and Transaction Processing workload types. Autonomous Data Guard is not available with the JSON and APEX workload types.
- Autonomous Data Guard is not available with Always Free Autonomous Databases.
- Autonomous Database does not provide access to a local Standby database:
  - You perform all operations, such as scaling up the **ECPU count** (**OCPU count** if your database uses OCPUs) enabling compute auto scaling on the Primary database and Autonomous Database performs the same actions on the local Standby database. Likewise, you only perform actions such as stopping or restarting the database on the Primary database.
  - A local Standby database is not available for use as a read-only database.

- The Number of ECPUs (OCPUs if your database uses OCPUs) allocated graph and the CPU utilization graph on the **Database Dashboard** card in Database Actions displays the ECPUs (OCPUs if your database uses OCPUs) allocated and the CPUs utilization for the Primary database. These graphs do not include information about a local Standby database or about a remote Standby database.

  The CPU Utilization metrics on the Oracle Cloud Infrastructure Console metrics page display the CPU utilization for the Primary database. Other metrics on this page also apply to the Primary database. These metrics do not include information about the local Standby database or remote Standby database.
- After a switchover or failover to a peer database, the peer database becomes the Primary and the graphs on the **Database Dashboard** in Database Actions and the Oracle Cloud Infrastructure Console metrics page display information about the Primary database. The graphs and metric do not contain information about the database that was the Primary before the switchover or failover operation.
- Automatic Failover to a local Standby is disabled during a Restore in Progress operation.
- Automatic Failover to a local Standby disabled when Upgrading a Database.
- When the **Lifecycle state** field for the Primary database shows **Stopped**, standby databases are also stopped. You may still perform a switchover when the Primary database is **Stopped**.
- Cross-Region Autonomous Data Guard Notes

Cross-Region Autonomous Data Guard Notes

The following are restrictions and limitations when you enable Autonomous Data Guard with a cross-region (remote) standby database:

- To disable Autonomous Data Guard with a cross-region standby database, you terminate the remote Standby database. See Terminate a Cross-Region Standby Database for more information.
- When a private endpoint is enabled or disabled on the Primary database, any previously configured Access Control List (ACL) on the Standby is enabled and the values are cleared. You must reset and verify the ACL on the Standby database after you disable a private endpoint on the Primary.
- You perform most operations, such as scaling up the **ECPU count** (**OCPU count** if your database uses OCPUs) and enabling compute auto scaling on the Primary database and Autonomous Database performs the same actions on the remote Standby database. Likewise, you only perform actions such as stopping or restarting the database on the Primary database.
- You can perform certain operations, such as configuring private endpoints on a remote Standby database.
- You can modify network configuration for ACLs on a remote Standby database. See Manage Remote Peer Network ACLs for more information.
- A remote Standby database is not available for use as a read-only database.
- Oracle Data Safe can be enabled on a database that has a cross-region Standby database enabled, but it only monitors the database within its region, and cannot monitor the standby in the event of a switchover or a failover.
- When you allow TLS authentication for the Primary database, Autonomous Data Guard enables TLS authentication in the cross- region Standby. Thus, when Autonomous Data Guard is enabled with a remote Standby, you can only allow TLS connections on the

Primary if both the Primary and the remote Standby are configured to support TLS connections. That is, the Primary and the remote Standby must either be configured with ACLs or with a private endpoint. See Network Configuration Prerequisites to Allow TLS Authentication for more information.

- See for the following information on using customer-managed keys with Autonomous Data Guard
  - Autonomous Data Guard with Customer Managed Keys
  - Notes for Customer-Managed Keys in OCI Vault with Autonomous Data Guard
- The **BYOL ECPU limit** you set on an Autonomous Data Guard Primary database does not apply to a cross-region or cross-tenancy Autonomous Data Guard Standby database. On a cross-region or cross-tenancy Standby you can independently set the **BYOL ECPU limit**, as required. Setting a value for **BYOL License limit** limits how many ECPUs will be covered by BYOL licenses.

  See Autonomous Data Guard Cross-Region BYOL Licensing for more information.
- When you enable Autonomous Data Guard with a cross-region standby database, the wallets for the primary and the standby specify different database hostnames and use different connection strings. Oracle recommends that applications use the connection string or wallet downloaded from the same region as the primary database.

  If you need to use a single connection string or wallet containing both the primary and the standby database hostnames, you may construct this manually.

  To manually construct a wallet that contains both the primary and the remote database connections strings:

  1. From the primary database's Oracle Cloud Infrastructure Console, click **Database connection** to download the primary's `wallet.zip`.
  2. From the remote standby database's Oracle Cloud Infrastructure Console, click **Database connection** to download the standby's `wallet.zip`.
  3. Unzip both wallet files and open the two `tnsnames.ora` files.
  4. Copy the remote database's connect descriptor into the primary database's connection string in the primary's `tnsnames.ora` file using your preferred retry delays.
  5. Zip the updated primary database wallet folder.

  With this updated `tnsnames.ora`, your primary database connection strings in the updated `wallet.zip` will contain both the primary and the standby hostnames, to support failover. An application using the updated wallet attempts to connect and retries connecting to the first listed database hostname, and if that connection fails due to the database being Unavailable, the application then automatically attempts to connect to the second database hostname.

  For example, if your Autonomous Data Guard is setup with the primary in Ashburn (IAD) and a cross-region standby in Phoenix (PHX), Oracle recommends that your mid-tier applications running in IAD use the connection string or wallet from that of the primary database in IAD, and your corresponding applications running in PHX use the connection string or wallet from that of the standby database in PHX. For a regional failover or switchover, Oracle recommends failing over both your database and your application or mid-tier, for optimum performance and to minimize any cross-regional latency.

  For example:

```
a6gxf2example9ep_high = (description_list=
     (failover=on) (load_balance=off)
     (description= (retry_count=15)(retry_delay=3)(address=(protocol=tcps)
```

```
(port=1522)(host=adb.us-ashburn-1.oraclecloud.com))
(connect_data=(service_name=mqssyowmexample_a6gxf2example9ep_high.adb.oracl
ecloud.com))(security=(ssl_server_dn_match=yes)))
     (description= (retry_count=15)(retry_delay=3)(address=(protocol=tcps)
(port=1522)(host=adb.us-phoenix-1.oraclecloud.com))
(connect_data=(service_name=mqssyowmexample_a6gxf2example9ep_high.adb.oracl
ecloud.com))(security=(ssl_server_dn_match=yes))))

a6gxf2example9ep_low =
     (description_list= (failover=on) (load_balance=off)
     (description= (retry_count=15)(retry_delay=3)(address=(protocol=tcps)
(port=1522)(host=adb.us-ashburn-1.oraclecloud.com))
(connect_data=(service_name=mqssyowmexample_a6gxf2example9ep_low.adb.oracle
cloud.com))(security=(ssl_server_dn_match=yes)))
     (description= (retry_count=15)(retry_delay=3)(address=(protocol=tcps)
(port=1522)(host=adb.us-phoenix-1.oraclecloud.com))
(connect_data=(service_name=mqssyowmexample_a6gxf2example9ep_low.adb.oracle
cloud.com))(security=(ssl_server_dn_match=yes))))

a6gxf2example9ep_medium =
     (description_list= (failover=on) (load_balance=off)
     (description= (retry_count=15)(retry_delay=3)(address=(protocol=tcps)
(port=1522)(host=adb.us-ashburn-1.oraclecloud.com))
(connect_data=(service_name=mqssyowmexample_a6gxf2example9ep_medium.adb.ora
clecloud.com))(security=(ssl_server_dn_match=yes)))
     (description= (retry_count=15)(retry_delay=3)(address=(protocol=tcps)
(port=1522)(host=adb.us-phoenix-1.oraclecloud.com))
(connect_data=(service_name=mqssyowmexample_a6gxf2example9ep_medium.adb.ora
clecloud.com))(security=(ssl_server_dn_match=yes))))
```

Backup and Restore Autonomous Database Instances

This section describes backup and recovery tasks on Autonomous Database.

- About Backup and Recovery on Autonomous Database
- Edit Automatic Backup Retention Period on Autonomous Database
  You can change the backup retention period for automatic backups on Autonomous Database with a retention period between 1 day and up to 60 days.
- Create Long-Term Backups on Autonomous Database
- View Backup Information and Backups on Autonomous Database
  You can view backup information and view the available backups on Autonomous Database.
- Restore and Recover your Autonomous Database
  From the Oracle Cloud Infrastructure Console you can restore the database using the **Restore** operation, where you initiate recovery for your database.
- Backup and Restore Notes
  Provides notes for automatic backups, long-term backups, and restoring your database.

About Backup and Recovery on Autonomous Database

Describes Autonomous Database backup options.

All backups taken and managed by Oracle in Autonomous Database are immutable. Immutable backups cannot be deleted or modified.

Backups are designed to be highly durable, as follows:

- In regions with multiple availability domains, backups are automatically replicated across availability domains.
- In regions with a single availability domain, backups are automatically replicated across fault domains.

Automatic Backups

Autonomous Database performs periodic backups including full, cumulative, and incremental backups, to ensure data reliability and recoverability. Each database instance is continuously backed up to achieve the Recovery Point Objective (RPO) of less than 10 seconds. See Recovery Time Objective and Recovery Point Objective for more information.

The choices for automatic backups are different depending on your compute model:

- **OCPU compute model**
  Autonomous Database automatically backs up your database for you. The retention period for backups is 60 days. You can restore and recover your database to any point-in-time in this retention period.
- **ECPU compute model**
  By default Autonomous Database automatically backs up your database for you and the retention period for backups is 60 days. Optionally, you can choose the backup retention period in days, from 1 to 60. You can restore and recover your database to any point-in-time in the specified retention period.

Billing for automatic backups depends on the compute model that you select when you provision or clone an Autonomous Database instance:

- **OCPU compute model**: the storage for automatic backups is included in the cost of database storage.
- **ECPU compute model**: the storage for backups is billed separately and in addition to database storage.

See How Is Autonomous Database Billed? for more information.

Immutable Backup Retention Option

To protect your data, you can enable the **Immutable backup retention** option to lock the automatic backup retention period. If locked, Autonomous Database does not allow changes to the retention period for automatic backups. This prevents any unauthorized removal of older backups by modification of the backup retention period.

> **Note:**
>
> After setting **Immutable backup retention**, you cannot disable this option or change the retention period. To disable **Immutable backup retention** or to change the backup retention period, file a Service Request at Oracle Cloud Support.

If you enable cross-region disaster recovery or backup replication and you set the **Immutable backup retention** option, the current primary's backup retention period remains locked before or after a switchover or failover.

- While provisioning or cloning an Autonomous Database Instance, you can select the immutable backup retention option.

  See Provision an Autonomous Database Instance and Clone an Autonomous Database Instance for more information.

- For an existing instance, you can update the settings of automatic backup retention and set the immutable backup retention option.

  See Edit Automatic Backup Retention Period on Autonomous Database for more information.

The immutable backup retention option is applicable for automatic backups only and does not affect long-term backups.

Long-Term Backups

Using the long-term backup feature you can create a long-term backup as a one-time backup or as scheduled long-term backup. You select the retention period for long-term backups in the range of a minimum of 3 months up to a maximum of 10 years.

See About Long-Term Backups on Autonomous Database for more information.

Recovery

You can initiate recovery for your Autonomous Database using the Oracle Cloud Infrastructure Console. Autonomous Database automatically restores and recovers your database to the point-in-time you specify or using the backup you select from a list of backups.

See Restore and Recover your Autonomous Database for more information.

Clone from a Backup

You can use either an automatic backup or a long-term backup to create a clone of your database.

See Clone an Autonomous Database from a Backup for more information.

Listing Backups

The list of backups available for recovery is shown on the Autonomous Database details page under **Backups**. Click **Backups** under **Resources** to show the backups.

The Oracle Cloud Infrastructure Console Autonomous Database details page, in the Backup area shows the **Total backup storage** field. This fields shows the total storage being billed, including for automatic backups and if there are long-term backups, this also includes the long term backup storage.

See View Backup Information and Backups on Autonomous Database for more information.

- About Long-Term Backups on Autonomous Database
  Using the long-term backup feature you can create a long-term backup as a one-time backup or as scheduled long-term backup. You select the retention period for long-term backups in the range of a minimum of 3 months up to a maximum of 10 years.

About Long-Term Backups on Autonomous Database

Using the long-term backup feature you can create a long-term backup as a one-time backup or as scheduled long-term backup. You select the retention period for long-term backups in the range of a minimum of 3 months up to a maximum of 10 years.

Autonomous Database takes scheduled long-term backups automatically according to the schedule you create: one-time, weekly, monthly, or yearly.

You can create a new database by cloning from a long-term backup. You cannot use a long-term backup to recover or restore the same database where the long-term backup was created.

Long-term backup billing:

- **ECPU compute model:** Long-term backup storage is billed in addition to your database storage. The total of your automatic backups and long-term backups will be billed additionally as backup storage.

  See ECPU Compute Model Billing Information for more information about how backup storage is billed.
- **OCPU compute model**: Using the OCPU compute model, long-term backups are billed in addition to your database storage and long term backup storage is billed at the same rate as database storage. Long-term backup storage is billed based on the storage usage, rounded to the nearest Terabyte. If your long-term backup storage usage is 0.5 TB, you are billed for 1 TB (at the same rate as database storage).

  For example, if you provision your Autonomous Database with 3 TB of storage and you create several long-term backups amounting to 1.2 TB, you are billed for a total of 5 TB of storage, as follows:

  – 3 TB (the provisioned storage)

  plus

  – 2 TB (1.2TB of long-term backups rounded up)

  See OCPU Compute Model Billing Information for more information about how backup storage is billed.

Each long term backup is a standalone backup that allows you to create a point-in-time clone, based on the long-term backup timestamp.

See Create Long-Term Backups on Autonomous Database for more information.

Edit Automatic Backup Retention Period on Autonomous Database

You can change the backup retention period for automatic backups on Autonomous Database with a retention period between 1 day and up to 60 days.

Editing the backup retention period is only available for instances using the ECPU compute model. See Compute Models in Autonomous Database for more information.

> **Note:**
>
> The storage for backups incurs additional costs. See ECPU Compute Model Billing Information for more information.

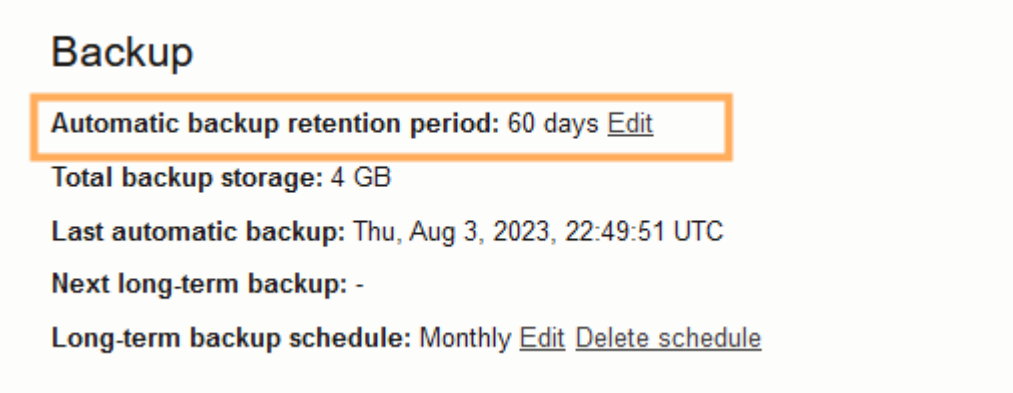
Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.

- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To change the automatic backup retention period:

1. On the Autonomous Database Details page, under **Backup**, in the **Automatic backup retention period** field, click **Edit**.



   This displays the Edit backup retention panel.

2. To edit the automatic backup retention period, enter the automatic backup retention period (in days).

   Enter a value in the text field or use the arrows to increase or decrease the backup retention period.

3. Select **Immutable backup retention** to lock the backup retention period.

   After setting the immutable backup retention option, you cannot disable this option or change the retention period. To disable immutable backup retention or to change the backup retention period, file a Service Request at Oracle Cloud Support.

   See About Backup and Recovery on Autonomous Database for more information.

4. Click **Save**.

Notes for specifying the automatic backup retention period:

- The backup retention period specifies how many days automatic backups are retained and the period for which you are billed for the automatic backups that are retained. If you change the backup retention period to a shorter period and automatic backups exist that are older, you are no longer billed for the automatic backups with timestamps beyond the specified retention period (and Autonomous Database may delete the older automatic backups).

  If you change the backup retention period to a longer period, for example if you change from 7 to 60 days, Autonomous Database starts retaining automatic backups up to 60 days, and billing for the retained backups. If older automatic backups already exist and were not deleted when you changed to a shorter backup retention period, the backups that have not been deleted show up and you are again billed for the automatic backups.

- The backup retention period for a remote peer database follows that of the source Primary database. Automatic backups are only taken on a peer database when it becomes the Primary database after a switchover or after a failover.

  See Use Standby Databases with Autonomous Data Guard for Disaster Recovery and Use Backup-Based Disaster Recovery for more information.

Create Long-Term Backups on Autonomous Database

You can create long-term backups on Autonomous Database with a retention period between three (3) months and up to ten (10) years.

When you create a long-term backup you select to create a one-time backup or set a schedule to automatically create backups that are taken weekly, monthly, or annually (yearly).

You can use a long-term backup to clone a new database. You may clone a new database from any timestamp between a long-term backup timestamp and up to 24 hours post (after) its timestamp. See Clone an Autonomous Database from a Backup for more information.

> **Note:**
>
> Long-term backups incur additional costs. See About Long-Term Backups on Autonomous Database for more information.

The Autonomous Database instance must have at least one automatic backup before you can create a long term backup. After you provision or clone an Autonomous Database instance, you may need to wait up to 4 hours before an automatic backup is available. See View Backup Information and Backups on Autonomous Database to find out if a backup exists.

Perform the following prerequisite steps as necessary:

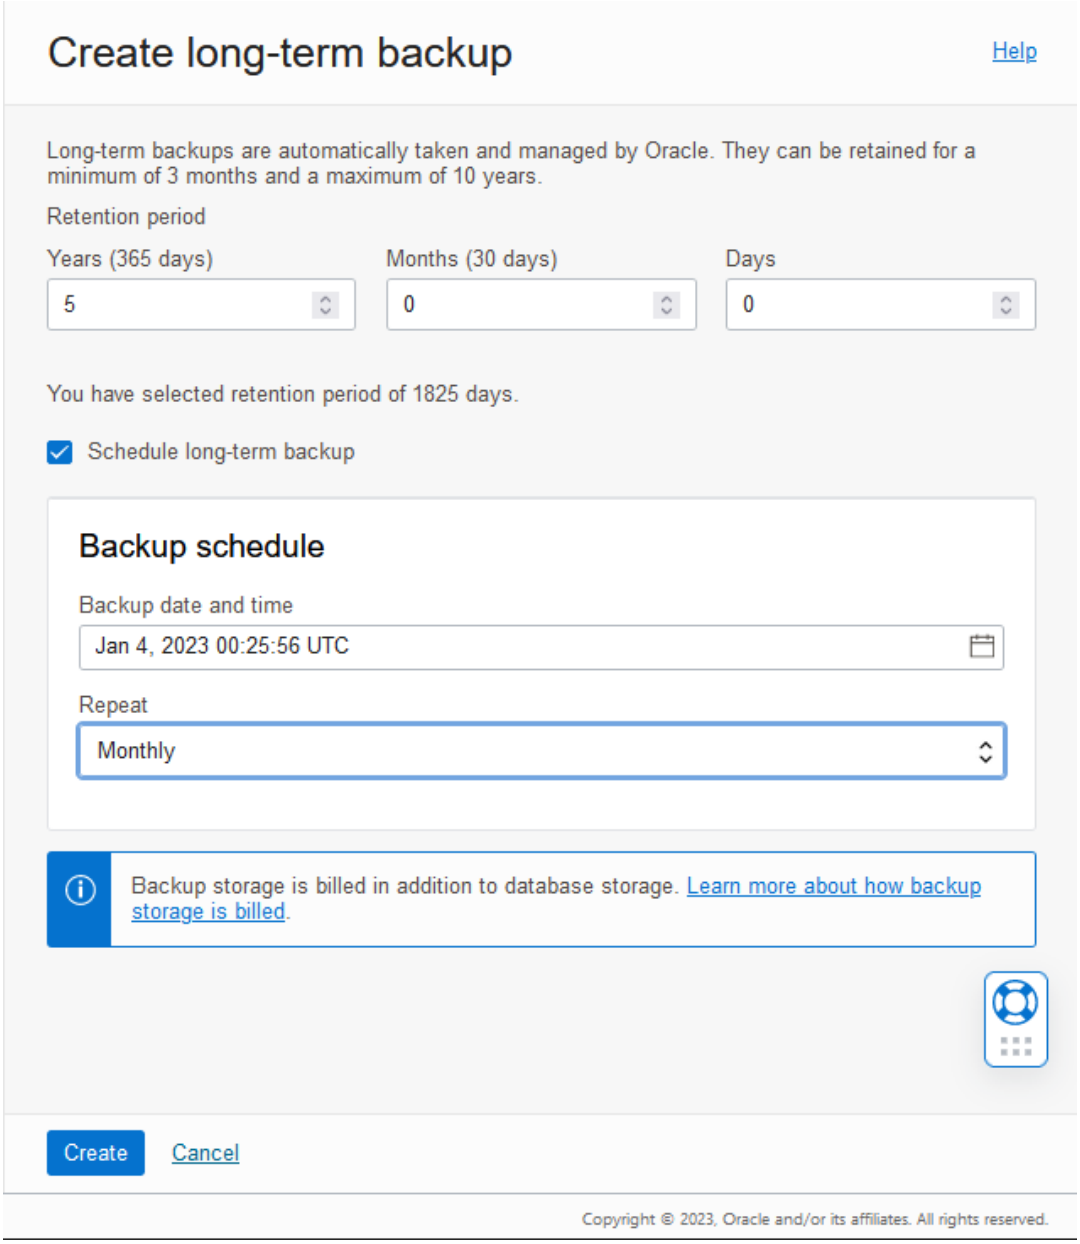
- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To create a long-term backup:

1. On the Autonomous Database Details page, under **Resources**, click **Backups**.
2. On the Autonomous Database Details page, under **Backups**, click **Create long-term backup**.
3. Set the long-term backup retention period.

   In the **Years**, **Months**, and **Days** fields, accept the default value, enter a value, or increment the values to increase or decrease the retention period for each of these choices.
4. If you want to begin the backup immediately, do not select **Schedule long-term backup** and skip to the next step.

   If you want to create a backup schedule, select **Schedule long-term backup**.



a. Enter a **Backup date and time**.

b. Select a repeat interval from the **Repeat** list.

When you select **One-time (does not repeat)**, this schedules a single long-term backup to be created on the specified date and time.

Selecting either **Weekly** specifies that Autonomous Database repeatedly creates long-term backups at the specified time at a weekly, every 7 days, interval. That is, backups are created and repeat at the same time weekly. For example, if the Backup date and Time is Jan 24, 2023 00:09:28 UTC and this is a Tuesday, and **Weekly** is selected, the long-term backup will happen every Tuesday.

Selecting **Monthly** specifies that Autonomous Database repeatedly creates long-term backups at the specified time each month. When you select **Monthly**, the backup will occur every month at the date you select. If you select a day of the month >= 29, Autonomous Database takes the backup on the last the last day of the month instead (to account for months with fewer days and leap years).

Selecting **Yearly** specifies that Autonomous Database repeatedly creates long-term backups at the specified time each year. When you select **Yearly**, the backup will occur every year at the date you select. If you select a day of the month >= 29, Autonomous Database takes the backup on the last the last day of the month instead (to account for months with fewer days and leap years).

5. Click **Create**.

   If you did not select **Schedule long-term backup**, Autonomous Database begins creating a long-term backup immediately.

- Edit Long-Term Backup Schedule
  You can edit the long-term backup schedule for an Autonomous Database instance.
- Delete Long-Term Backup Schedule
  You can delete the long-term backup schedule for an Autonomous Database instance.
- Delete a Long-Term Backup
  You can delete a long-term backup.

Edit Long-Term Backup Schedule

You can edit the long-term backup schedule for an Autonomous Database instance.

1. On the Autonomous Database Details page, in the **Long-term backup schedule** field under Backup, click **Edit**.

   This shows the Configure long-term backup schedule dialog.
2. Make changes to the long-term backup schedule.
3. Click **Save**.

Delete Long-Term Backup Schedule

You can delete the long-term backup schedule for an Autonomous Database instance.

1. On the Autonomous Database Details page, in the **Long-term backup schedule** field under Backup, click **Delete Schedule**.
2. In the Delete Schedule confirmation dialog, click **Delete**.

Delete a Long-Term Backup

You can delete a long-term backup.

To delete a long-term backup,

1. On the Autonomous Databases page select your Autonomous Database from the links under the **Display Name** column.
2. On the Autonomous Database Details page, under **Resources**, click **Backups**.
3. In the table, from the list of backups select the long-term backup you want to delete and click  at the end of a row.
4. In the list, select **Delete**.

View Backup Information and Backups on Autonomous Database

You can view backup information and view the available backups on Autonomous Database.

The Oracle Cloud Infrastructure Console shows information about backups. To view general backup information, view the Backup area on the Autonomous Database details page. Depending on your compute model, the Backup area includes the following information:

- **Automatic backup retention period**:
- **Total backup storage**: Shows the total storage being billed, including for automatic backups and if there are long-term backups, this also includes the long term backup storage.
- **Last automatic backup**: Shows the timestamp of the last automatic backup.
- **Next long-term backup**: When you configure scheduled long-term backups, this shows the timestamp for the next long-term backup.
- **Long-term backup schedule**: Shows the long-term backup schedule, with links to edit or delete a long-term backup schedule.

To list completed backups:

1. On the Autonomous Databases page select your Autonomous Database from the links under the **Display Name** column.
2. On the Autonomous Database Details page, under **Resources**, click **Backups**.

This shows the list of backups for the Autonomous Database instance.

If you click  at the end of a row for a backup in the Backups list, this shows the actions you can perform.

- **Actions for Automatic Backups**:
  - Restore: See Restore and Recover your Autonomous Database for more information.
  - Create clone: See Clone an Autonomous Database from a Backup for more information.
- **Actions for long-term backups**:
  - Create clone: Clone an Autonomous Database from a Backup for more information.
  - Edit retention period: This allows you to change the retention period for a long-term backup.
  - Delete: This allows you to remove a long-term backup.

Restore and Recover your Autonomous Database

From the Oracle Cloud Infrastructure Console you can restore the database using the **Restore** operation, where you initiate recovery for your database.

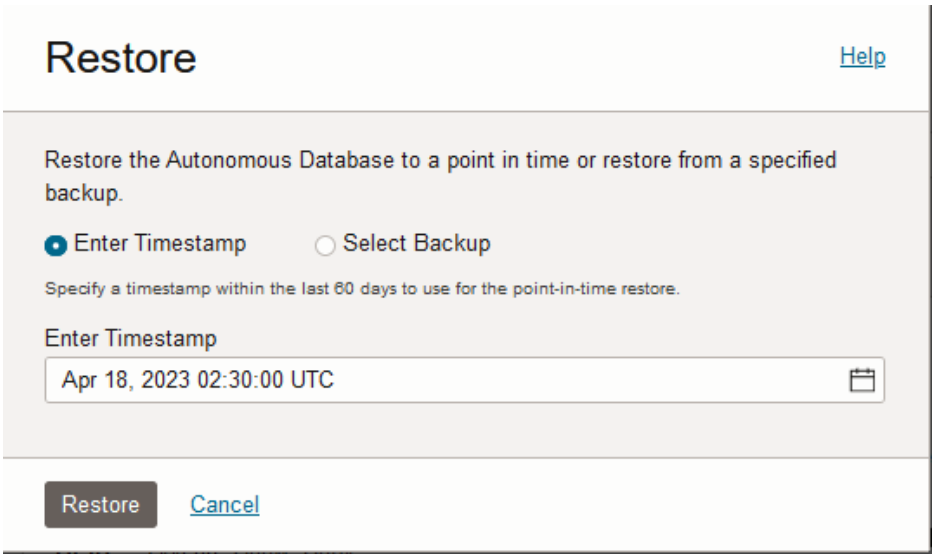
Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the  next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.

- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To restore and recover your database, do the following:

1. On the Autonomous Database Details page, from the **More actions** drop-down list, select **Restore** to display the **Restore** prompt.
2. In the Restore prompt, select **Enter Timestamp** or **Select Backup** to restore to a point in time or to restore from a specified backup.
   - **Enter Timestamp**: Enter a timestamp to restore to in the **Enter Timestamp** calendar field.



     Click the Calendar icon to show the date and timestamp calendar selector. If automatic backups are enabled on your instance, you can select a timestamp to restore and recover your database to any point-in-time in the retention period.

     Manually edit the values in the **Enter Timestamp** field if you prefer to use a more granular timestamp.
   - **Select Backup**: Select a backup from the list of backups. Limit the number of backups you see by specifying a period using the **From** and **To** calendar fields.
3. Click **Restore**.

> **Note:**
>
> Restoring Autonomous Database puts the database in the unavailable state during the restore operation. You cannot connect to a database in that state. The only lifecycle management operation supported in unavailable state is **Terminate**.

When concurrent operations such as scaling are active, the confirmation also confirms either pausing or canceling the concurrent operation. See Concurrent Operations on Autonomous Database for more information.

The details page shows Lifecycle state: **Restore In Progress**

4. When the restore operation finishes your Autonomous Database instance opens and the Lifecycle state shows **Available**.

At this point you can connect to your Autonomous Database instance and check your data to validate that the restore point you specified was correct. After checking your data if you find that the restore date you specified was not the one you needed, you can initiate another restore operation to another point in time.

> **Notes:**
>
> - The restore operation also restores the `DATA_PUMP_DIR` directory and user defined directories to the timestamp you specified for the restore; files that were created after that timestamp would be lost.
> - When you restore, the Oracle Machine Learning workspaces, projects, and notebooks are not restored.
> - For external tables, partitioned external tables, and the external partitions of hybrid partitioned tables a backup does not include the external files that reside on your Object Store. Thus, for operations where you use a backup to restore your database, such as **Restore** or **Clone from a backup** it is your responsibility to backup, and restore if necessary, the external files associated with external tables, external partitioned tables, or the external files for a hybrid partitioned table.
>
> See Clone an Autonomous Database from a Backup for information on using **Clone from a backup**.

Backup and Restore Notes

Provides notes for automatic backups, long-term backups, and restoring your database.

- **Files on Object Store**: For external tables, partitioned external tables, and the external partitions of hybrid partitioned tables, backups do not include the external files that reside on Object Store. Thus, for operations where you use a backup to restore your database, such as **Restore** or **Clone from a backup**, it is your responsibility to backup and restore if necessary, the external files associated with external tables, external partitioned tables, or the external files for a hybrid partitioned table.

  See Restore and Recover your Autonomous Database for information on **Restore**.

  See Clone an Autonomous Database from a Backup for information on **Clone from a backup**.
- **Stopped Database**: Automatic backups occur when a database is stopped. If you stop your database, you do not miss a regularly scheduled automatic backup.
- **Manual Backups**: The Oracle Cloud Infrastructure Console does not provide the option to create a manual backup. Oracle recommends that you do not create or use manual backups. If required, manual backups are available using the API. See CreateAutonomousDatabaseBackup for more information.
- **Sample Data Sets SH and SSB**: The sample data sets are not included in backups and do not impact the size or speed of either backup or restore operations.

Long-Term Backup Notes

- Long-term backups are retained and managed by Oracle for the specified retention period you select when you create or schedule long-term backups.
- Each long-term backup creates a full backup and the backup can only be used from the Autonomous Database instance to clone a new database.
- After the long-term retention period is over a backup will be deleted and is no longer retained.
- The encryption key option that is selected when a long-term backup is created applies when you create a clone from the long-term backup. If a customer-managed master encryption key is in use at the time a long-term backup is created, then the same customer-managed master encryption key must be available to access a database that you clone from the long-term backup.

  See Manage Encryption Keys on Autonomous Database for more information.
- While backing up a database, the database is fully functional. However, during a long-term backup certain lifecycle management operations, such as stopping the database may affect the backup. See Long-Term Backup with Concurrent Operations for more information.
- Cloning from a long-term backup creates a new database with the most current available database version, which is possibly not the original database version in use at the time when the long-term backup was created. A database cloned from a long-term backup uses latest available database version.

  For example, if a long-term backup is taken on a database with version 19c, and then 4 years later the long-term backup is used to clone a new database, the new database may only offer the latest database version (for example Oracle Database 23ai, if version 19c is not available).
- Long-term backups are only accessible while a database exists. If you terminate an Autonomous Database instance, you no longer have access to long-term backups. If you want to retain long-term backups and minimize costs:

  1. Drop all tables and scale down the database to the minimum storage.
  2. Stop the database.

  This preserves the long-term backups that were created for the Autonomous Database instance.

Use OraMTS Recovery Feature on Autonomous Database

Use Oracle MTS (OraMTS) Recovery Service to resolve an in-doubt Microsoft Transaction Server transaction.

- About OraMTS Recovery Service
  The Oracle MTS (OraMTS) Recovery Service resolves an in-doubt Microsoft Transaction Server transaction.
- Prerequisites to Enable OraMTS Recovery Service on Autonomous Database
  Lists the prerequisites to enable OraMTS Recovery Service for an Autonomous Database.
- Enable OraMTS Recovery Service on an Autonomous Database
  Describes steps to enable OraMTS Recovery Service on an Autonomous Database.
- Disable OraMTS Recovery Service on an Autonomous Database
  Describes steps to disable OraMTS Recovery Service for an Autonomous Database.

About OraMTS Recovery Service

The Oracle MTS (OraMTS) Recovery Service resolves an in-doubt Microsoft Transaction Server transaction.

An Microsoft Transaction Server (MTS) is a COM-based transaction processing system that runs on an internet or network server.

An Oracle MTS Recovery Service is automatically installed with Oracle Services For Microsoft Transaction Server. The Oracle MTS (OraMTS) Recovery Service resolves in-doubt transactions on the computer that started the failed transaction. A scheduled recovery job for each MTS-enabled database lets the OraMTS Recovery Service resolve in-doubt transactions.

To use MTS with an Oracle database, distributed transaction capabilities are necessary.

When any of these components fails, Oracle transactions connected to Microsoft Transaction Server become in-doubt transactions:

- Microsoft Transaction Server application
- Network
- Microsoft Distributed Transaction Coordinator (MS DTC)

See Using Microsoft Transaction Server with Oracle Database for more information.

The Oracle MTS (OraMTS) Recovery Service resolves an in-doubt Microsoft Transaction Server transaction in this order:

1. The DBMS recovery job detects an in-doubt MTS-related transaction.
2. The DBMS recovery job extracts the recovery service's endpoint address from the `XID` of the in-doubt transaction and requests the recovery service for the outcome of the MTS/MS DTC transaction.
3. The recovery service requests its MS DTC for transaction outcome.
4. The recovery service reports transaction outcome to the DBMS job process.
5. The DBMS recovery job commits or terminates the in-doubt transaction based on the outcome reported by MS DTC.

> **Note:**
>
> Each computer can only have one instance of Oracle MTS (OraMTS) Recovery Service installed.

Prerequisites to Enable OraMTS Recovery Service on Autonomous Database

Lists the prerequisites to enable OraMTS Recovery Service for an Autonomous Database.

To enable Oracle MTS Recovery Service on an Autonomous Database:

- You must configure your database on a private endpoint.
- For your OraMTS Recovery Service, you must deploy the VM in the same private network as the database.

- You must configure an OCI Private Load Balancer (LBaaS) and the Load Balancer (LBaaS) must be able to access the VM on port 2030. See Load Balancer Management for more information.
- Your database must be able to communicate with the Load Balancer (LBaaS) on port 443. To enable this you need an egress rule for port 443 in VCN's security list or in the network security group.
- Your Load Balancer (LBaaS) must also be able to receive the communication from the database. To enable this you need an ingress rule for your Load Balancer (LBaaS) for port 443.
- Reserve a domain name with a domain provider.
- Generate an SSL certificate for the domain.
- You must configure a secure HTTPS endpoint using OCI Load Balancer to ensure that the communication between the Autonomous Database and the MTS server uses HTTPS protocol with SSL encryption. See Configure Network Access with Private Endpoints and Submit an HTTP Request to a Private Host for more information.

Enable OraMTS Recovery Service on an Autonomous Database

Describes steps to enable OraMTS Recovery Service on an Autonomous Database.

To enable OraMTS Recovery Service on your Autonomous Database, you must be logged in as the `ADMIN` user or have the `EXECUTE` object privilege on `DBMS_CLOUD_ADMIN`.

Run `DBMS_CLOUD_ADMIN.ENABLE_FEATURE` to enable OraMTS Recovery Service on your Autonomous Database.

Example to Enable and Verify the OraMTS Recovery Service:

```
BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_FEATURE(
        feature_name => 'ORAMTS',
        params       => JSON_OBJECT('location_uri' VALUE 'https://
mymtsserver.mycorp.com')
   );
END;
/


SELECT property_value FROM database_properties WHERE property_name =
'ORAMTS_SERVER';
```

The first example enables the OraMTS Recovery Service on your Autonomous Database.

The `feature_name` parameter specifies the name of the feature to enable. The `ORAMTS` value indicates that you are enabling the OraMTS recovery service feature for your database.

The `location_uri` parameter specifies the HTTPS URL for the OraMTS server in a customer network.

The second example is a SQL statement that you can run to verify that the OraMTS Recovery Service is enabled for your Autonomous Database.

See ENABLE_FEATURE Procedure for more information.

Disable OraMTS Recovery Service on an Autonomous Database

Describes steps to disable OraMTS Recovery Service for an Autonomous Database.

To disable OraMTS Recovery Service on your Autonomous Database, you must be logged in as the `ADMIN` user or have the `EXECUTE` object privilege on `DBMS_CLOUD_ADMIN`.

Run `DBMS_CLOUD_ADMIN.DISABLE_FEATURE` to disable OraMTS Recovery Service on your Autonomous Database.

Example to disable the OraMTS Recovery Service:

```
BEGIN
DBMS_CLOUD_ADMIN.DISABLE_FEATURE(
    feature_name => 'ORAMTS');
END;
/
```

This disables the OraMTS Recovery feature on your Autonomous Database.

The `feature_name` parameter specifies the name of the feature to disable. The `ORAMTS` value indicates that you are disabling the OraMTS recovery service feature for your database.

See DISABLE_FEATURE Procedure for more information.

Use Backup-Based Disaster Recovery

Backup-Based Disaster Recovery provides a lower-cost disaster recovery option for Autonomous Database (the RTO is higher with this option, as compared to using Autonomous Data Guard).

Autonomous Database Disaster Recovery Options

- **Backup-Based Disaster Recovery**: uses backups to instantiate a peer database at the time of switchover or failover. This enables you to have a lower cost and higher Recovery Time Objective (RTO) disaster recovery option for your Autonomous Database, as compared with Autonomous Data Guard. For local Backup-Based Disaster Recovery, existing local backups are utilized. There are no additional costs for a local Backup-Based Disaster Recovery. Cross-Region Backup-Based Disaster Recovery incurs an additional cost.

  Backup-Based Disaster Recovery is available for all workload types.

- **Autonomous Data Guard**: When you add an Autonomous Data Guard standby database, the system creates a standby database that continuously gets updated with the changes from the primary database. You can use Autonomous Data Guard with a standby in the current region, a local standby, or with a standby in a different region, a cross-region standby. You can also use Autonomous Data Guard with both a local standby and a cross-region standby.

  You can also create an Autonomous Data Guard standby in a different tenancy.

  Autonomous Data Guard is available for the following workload types:

  - Data Warehouse
  - Transaction Processing

See Use Standby Databases with Autonomous Data Guard for Disaster Recovery for details on Autonomous Data Guard.

Topics

- About Backup-Based Disaster Recovery
  Backup-Based Disaster Recovery uses backups to instantiate a peer database at the time of switchover or failover. This enables you to have a lower cost and higher Recovery Time Objective (RTO) disaster recovery option for your Autonomous Database, as compared with Autonomous Data Guard.
- View Backup-Based Disaster Recovery Peer
  Backup-Based Disaster Recovery with a local peer is enabled by default for a newly created Autonomous Database instance, and for existing databases. A local Backup-Based Disaster Recovery peer does not incur any additional cost.
- Update Disaster Recovery Type
  Describes the steps to change to an alternative Disaster Recovery option.
- Backup-Based Disaster Recovery Switchover and Failover Operations
  Describes switchover and failover operations with an Backup-Based Disaster Recovery peer database.
- Backup-Based Disaster Recovery Cross-Region Operations
  Describes cross-region operations with Backup-Based Disaster Recovery.
- Backup-Based Disaster Recovery Events
  You can use Oracle Cloud Infrastructure events to respond when Autonomous Database changes its state due to a Backup-Based Disaster Recovery related event.
- Use the API
  Provides links for details on using API operations to manage Backup-Based Disaster Recovery.

About Backup-Based Disaster Recovery

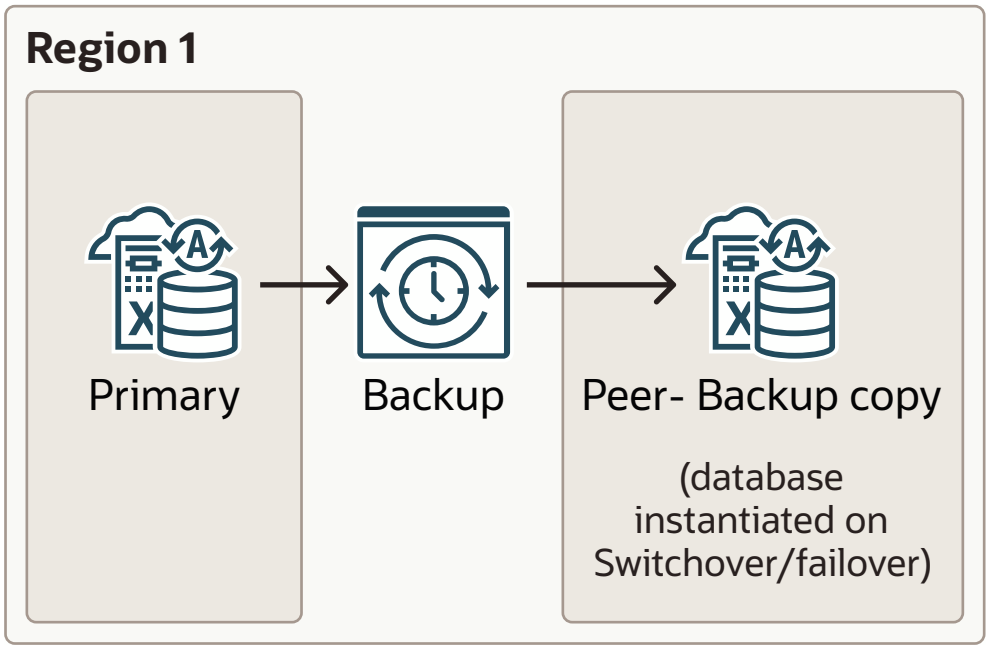
Backup-Based Disaster Recovery uses backups to instantiate a peer database at the time of switchover or failover. This enables you to have a lower cost and higher Recovery Time Objective (RTO) disaster recovery option for your Autonomous Database, as compared with Autonomous Data Guard.

You can use Backup-Based Disaster Recovery with a peer in the current region, a local peer, or with one or more disaster recovery peers in different regions, or you can add both a local disaster recovery peer and one or more a remote disaster recovery peers. You can also create a Backup-Based Disaster Recovery peer, either local or remote in a different tenancy.

> **Note:**
>
> Backup-Based Disaster Recovery (backup copy) is available in all Autonomous Database workload types. Backup-Based Disaster Recovery is not available with Always Free Autonomous Database.

Backup-Based Disaster Recovery with Local Peer

For local Backup-Based Disaster Recovery, existing local backups are utilized. There are no additional costs for local Backup-Based Disaster Recovery.



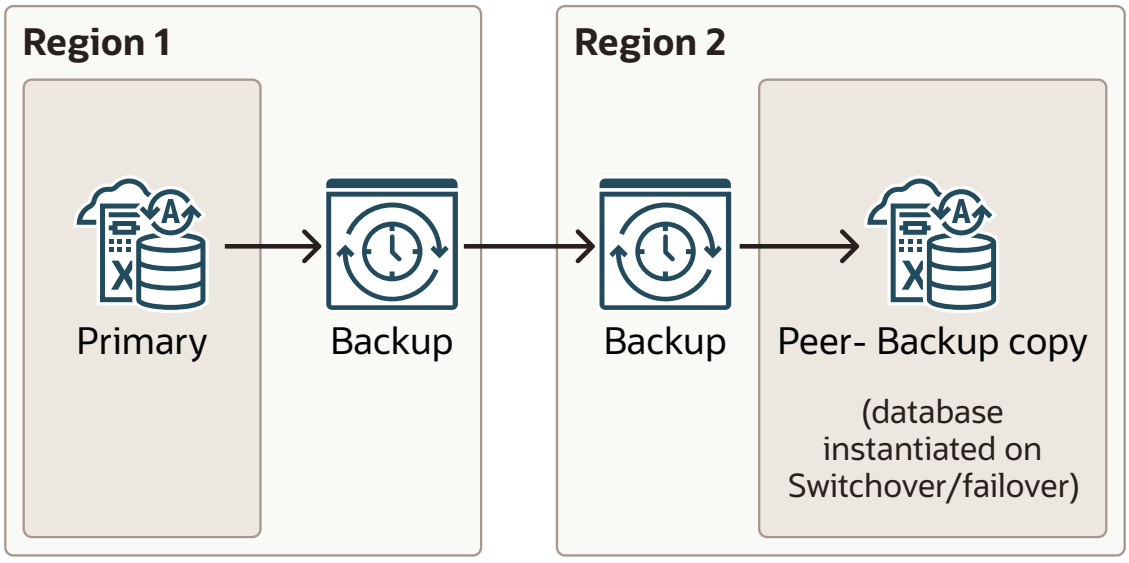
For better resilience for Backup-Based Disaster Recovery, a peer is instantiated as follows:

- In regions with more than one availability domain, a local peer is instantiated in a different availability domain than the primary database.
- In regions with a single availability domain, a local peer is instantiated in a different fault domain than the primary database (that is, on a different physical machine).

All Autonomous Database features from the primary database are available when a peer is instantiated and becomes the primary, after the system fails over or after you perform a switchover operation. See Autonomous Data Guard with Local Standby for more information.

Backup-Based Disaster Recovery with Cross-Region Peer

For backup-based disaster recovery with a cross-region peer, backups are copied to the remote region. Cross-region Backup-Based Disaster Recovery incurs additional costs.



Autonomous Database allows you to create one or more remote disaster recovery peer databases, depending on your compute model:

- OCPU Compute Model: You can add one remote disaster recovery peer in a paired region. Paired regions are remote regions where you can create a cross-region peer.
- ECPU Compute Model: You can add multiple remote disaster recovery peers, with up to one peer in each remote paired region. For example if your primary database is in the IAD region, you can add a remote peer in PHX and in SJC, but you cannot add two remote peers in PHX.

Paired regions are remote regions where you can create a cross-region disaster recovery peer. See Autonomous Database Cross-Region Paired Regions for more information on paired regions.

Topics

- Backup-Based Disaster Recovery Recovery Time Objective (RTO) and Recovery Point Objective (RPO)
  When you perform a failover using with Backup-Based Disaster Recovery enabled, the peer instance assumes the role of the primary instance, according to the Recovery Time Objective (RTO) and Recovery Point Objective (RPO).
- Replicating Backups to a Cross-Region Backup Based Disaster Recovery Peer
  When you add a cross-region Backup-Based Disaster Recovery peer you can enable cross-region backup replication for automatic backups.

Backup-Based Disaster Recovery Recovery Time Objective (RTO) and Recovery Point Objective (RPO)

When you perform a failover using with Backup-Based Disaster Recovery enabled, the peer instance assumes the role of the primary instance, according to the Recovery Time Objective (RTO) and Recovery Point Objective (RPO).

The RTO is the maximum amount of time required to restore database connectivity to a backup copy database after a failover is initiated. The RPO is the maximum duration of potential data loss, in minutes, on the primary database.

Backup-Based Disaster Recovery RTO and RPO numbers are:

Backup-Based Disaster Recovery Configuration	RTO	RPO
Local backup copy	one (1) hour + 1 hour per 5 TB	10 seconds
Cross-region (remote) backup copy	one (1) hour + 1 hour per 5 TB	1 min

Replicating Backups to a Cross-Region Backup Based Disaster Recovery Peer

When you add a cross-region Backup-Based Disaster Recovery peer you can enable cross-region backup replication for automatic backups.

By default, automatic backups are created and maintained at the current Primary database and are not replicated to a cross-region peer. Optionally, you can enable replication of the automatic backups to a cross-region peer.

When you enable cross-region backup replication, up to 7 days of automatic backups for the primary are replicated to a cross-region peer. When this feature is enabled automatic backups are available in the remote region as follows:

- After a switchover or failover you can restore or clone to any timestamp in the past seven (7) days, or to any timestamp in the specified retention period when the retention period is set to less than seven days.
- All backups for the primary that are replicated to the remote region are deleted on the remote region peer after seven days, or after the retention period number of days when the retention period is set to less than seven days.
- You cannot modify the backup retention period for replicated backups, except if you modify the backup retention period on the primary to specify a value less than seven days. In this case, the retention period for replicated backups on the remote region matches the automatic backup retention period set on the primary.

The cross-region backup replication incurs an additional cost. See Oracle Autonomous Database Serverless Features Billing for more information.

See the following for more information:

- Add a Cross-Region Disaster Recovery Peer
- Enable or Disable Backup Replication for an Existing Cross-Region Peer

Note the following for cross-region backup replication:

- After a switchover or a failover, while the cross-region database is in the primary role, backups are taken on the current primary and are replicated to the current (remote) peer.
- When using Backup-Based Disaster Recovery with a cross-region peer, this feature is supported for all workload types.

View Backup-Based Disaster Recovery Peer

Backup-Based Disaster Recovery with a local peer is enabled by default for a newly created Autonomous Database instance, and for existing databases. A local Backup-Based Disaster Recovery peer does not incur any additional cost.

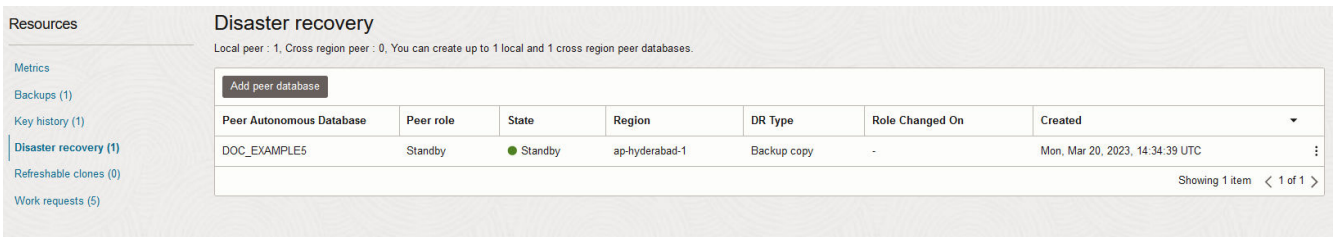
Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To view the disaster recovery information for your Autonomous Database, do the following:

On the Autonomous Database Details page, in the **Resources** area click **Disaster recovery**. The **DR type** field indicates the type of disaster recovery, either Backup-Based Disaster Recovery (Backup copy) or Autonomous Data Guard.

For example:



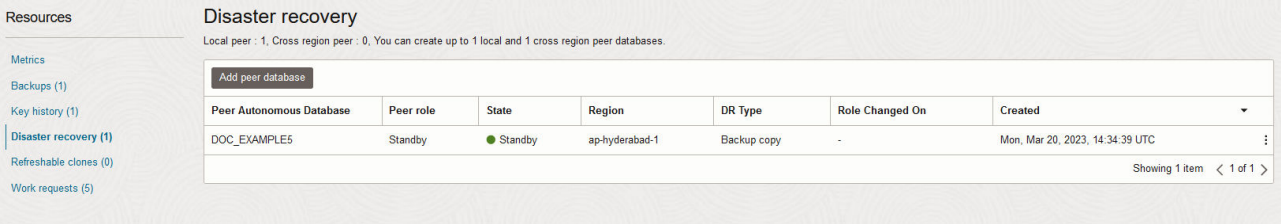
Update Disaster Recovery Type

Describes the steps to change to an alternative Disaster Recovery option.

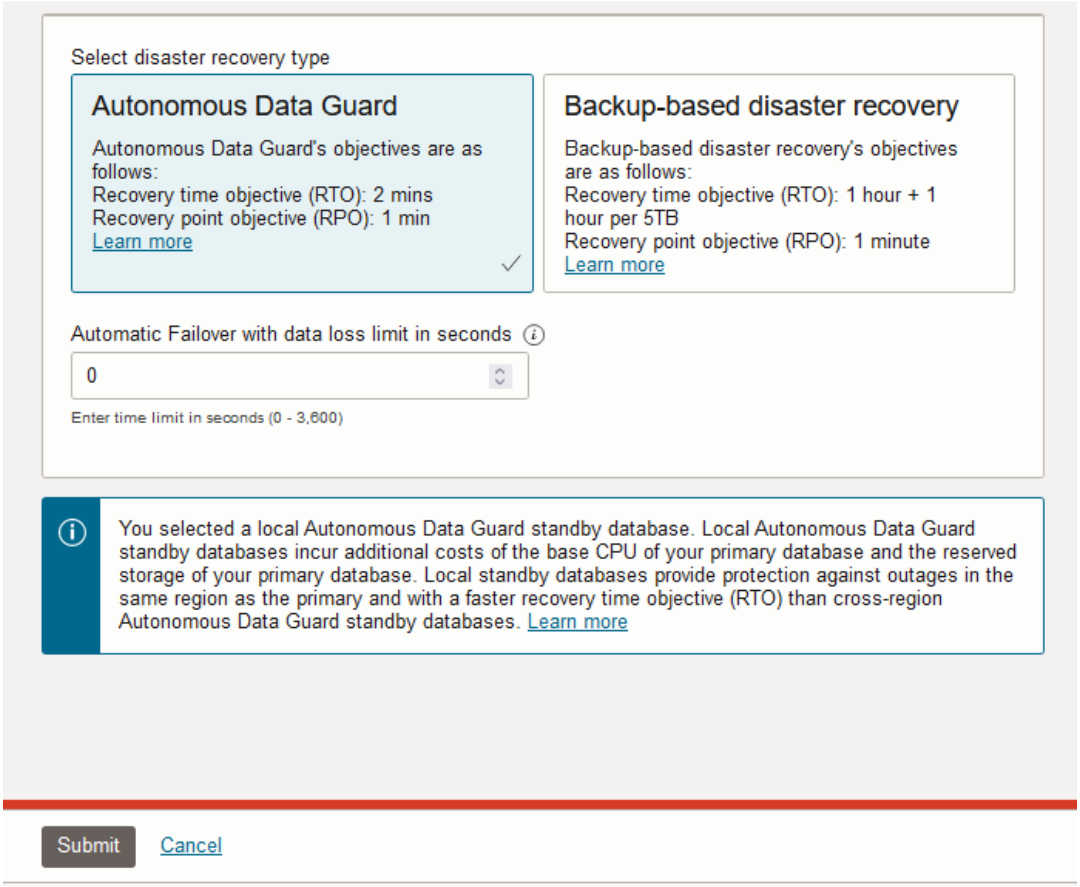
Backup-Based Disaster Recovery is enabled by default for an Autonomous Database with one local peer. Disaster Recovery cannot be disabled for an Autonomous Database instance. However, you can choose to update your Disaster recovery type to Autonomous Data Guard. See Use Standby Databases with Autonomous Data Guard for Disaster Recovery for more information about Autonomous Data Guard.

To update the Disaster recovery type:

1. On the Primary Autonomous Database, on the Autonomous Database Details page under **Resources**, select **Disaster recovery**.



2. In the row showing the database's disaster recovery details, click  at the end of the row and select **Update disaster recovery**.
3. On the Update disaster recovery page, choose **Autonomous Data Guard**.



4. Click **Submit**. This starts provisioning an Autonomous Data Guard standby database.
5. The disaster recovery type is changed to, Autonomous Data Guard, as shown in the **DR type** column.

Backup-Based Disaster Recovery Switchover and Failover Operations

Describes switchover and failover operations with an Backup-Based Disaster Recovery peer database.

- About Switchover to a Backup Copy Peer
  When you perform a switchover, the primary database becomes the backup copy and the backup copy becomes the primary database, with no data loss.
- Perform a Switchover to a Local Backup Copy Peer
  When you perform a switchover the primary database becomes the peer, and the peer becomes the primary database, with no data loss.
- Perform a Switchover to a Cross-Region Backup Copy Peer
- Perform a Failover
- Notes for Performing a Switchover to a Backup Copy Peer
  Provides notes for Backup-Based Disaster Recovery switchover:

About Switchover to a Backup Copy Peer

When you perform a switchover, the primary database becomes the backup copy and the backup copy becomes the primary database, with no data loss.

A switchover is typically done to test failover to the backup copy for audit or certification reasons, or to test your application's failover procedures when you have added a backup copy peer.

For switchover to a backup copy peer, the **Autonomous Database Details** page, shows the **Switchover** link under **Disaster recovery** and the Oracle Cloud Infrastructure Console on the primary database also shows a **Switchover** link in the **Role** field when both the primary database and a backup copy peer are available. You can perform a switchover when the primary database **Lifecycle state** shows **Available** or **Stopped** and a backup copy is available (the **State** field shows **Standby**).

To see the peer state, under **Resources** click **Disaster recovery** and for the peer listed in the **Peer Autonomous Database** column, check that the **State** shows **Standby**.

Using the Autonomous Database API, you can initiate the switchover operation at any time. See Use the API for more information.

Perform a Switchover to a Local Backup Copy Peer

When you perform a switchover the primary database becomes the peer, and the peer becomes the primary database, with no data loss.

A switchover is typically done to test failover to the peer for audit or certification reasons or to test your application's failover procedures with Backup-Based Disaster Recovery.

For a switchover to a backup copy, the **Autonomous Database Details** page, the Oracle Cloud Infrastructure Console on the database with the **Primary** role shows a **Switchover** link in the **Role** field when both the primary database and a peer are available. You can perform a switchover when the primary database **Lifecycle state** shows **Available** or **Stopped** and a peer is available (the **State** field shows **Standby**).

To see the peer status, under **Resources** click Disaster recovery and for the peer listed in the **Peer Autonomous Database** column, check that the **State** field shows **Standby**.

Using the Autonomous Database API, you can initiate the switchover operation at any time. See Use the API for more information.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To perform a switchover, do the following:

1. On the Autonomous Database Details page, under **Disaster recovery**, in the **Role** field, click **Switchover**.

   As an alternative, to initiate a switchover you can select **More actions** and **Switchover** or select **Disaster recovery** under **Resources** and **Switchover**.
2. In the **Confirm switchover to peer** dialog, in the **Select peer** list, choose the peer to switchover.
3. Enter the database name to confirm that you want to switch over.
4. Click **Confirm switchover to peer**.

   When concurrent operations such as scaling are active, the confirmation also confirms either pausing or canceling the concurrent operation. See Concurrent Operations on Autonomous Database for more information.

   The database **Lifecycle state** changes to **Updating**. To see the state of the peer, under **Resources**, click **Disaster recovery** where the **State** column shows **Role Change in Progress**.

When the switchover completes, Backup-Based Disaster Recovery does the following:

- The Backup-Based Disaster Recovery resource information is updated to reflect the switchover. Select **Disaster recovery** under **Resources** to see the updated information.
- Autonomous Database reports the time in the **Role Changed On** the field.

Perform a Switchover to a Cross-Region Backup Copy Peer

When you perform a switchover, the primary database becomes the peer database and the peer database becomes the primary database, with no data loss.

> **Note:**
>
> For a cross-region switchover you must initiate the switchover from the cross-region peer.

You have several options to access the cross-region peer:

- Select the remote region in Oracle Cloud Infrastructure Console and then access the peer directly.

- Access the primary, and from the primary database you can access the peer from the Autonomous Database Details page by selecting **Disaster recovery**, under **Resources** and clicking the link for the backup copy peer in the **Peer Autonomous Database** column.

To perform a switchover:

1. On the cross-region peer, on the Autonomous Database Details page, under **Disaster recovery**, in the **Role** field, click **Switchover**.

   As an alternative, in the row showing the database's disaster recovery details, click  at the end of the row and select **Switchover**.
2. In the **Confirm switchover to peer** dialog, enter the peer database name to confirm that you want to switch over.
3. In the **Confirm switchover to peer** dialog, click **Confirm switchover to peer**.

   When concurrent operations such as scaling are active, the confirmation also confirms either pausing or canceling the concurrent operation. See Concurrent Operations on Autonomous Database for more information.

   The database **Lifecycle state** changes to **Role change in progress**. To see the state of the peer database, under **Resources** click **Disaster recovery**. The state **State** shows **Updating**.

When the switchover completes, Autonomous Database does the following:

- The display name shows the Primary indicator.
- The **Disaster recovery** resource information is updated to reflect the switchover. Under **Resources** select **Disaster recovery** to see the updated information.
- Autonomous Database reports the time of the last switchover when you hover over the  in the **Role** field.

See Notes for Performing a Switchover for more information.

Perform a Failover

When the primary database goes down, with Backup-Based Disaster Recovery you can perform a manual failover to make the local peer the primary database.

Backup-Based Disaster Recovery does not provide an automatic failover option. If you want to provide automatic failover, where the system monitors the primary instance and automatically fails over to a local standby database in certain scenarios, you need to change the disaster recovery type for the local instance to use Autonomous Data Guard.

With both a local Backup-Based Disaster Recovery peer and a cross-region Backup-Based Disaster Recovery peer, Oracle recommends that you attempt a manual failover to the local peer first (not to the cross-region peer).

Depending on how you enable Backup-Based Disaster Recovery, there are different steps to perform a manual failover to a peer:

- **When you configure Backup-Based Disaster Recovery with only a local peer**:

  When you have a local peer and the switchover is not successful, the Oracle Cloud Infrastructure console shows a banner with information about why the switchover was not successful and the Oracle Cloud Infrastructure console shows a **failover** link in the **Role** field that you can click to initiate a failover to the local peer. The failover link only shows when the Primary database is unavailable and a peer is available. That is, the Primary database **Lifecycle state** field shows **Unavailable** and the local peer is available. Using

the API, you can initiate manual failover at any time. See Use the API for information on using the API.

To see the peer status, under **Resources** click **Disaster recovery** and for the peer listed in the Peer Autonomous Database column check that the **State** field shows **Available** or **Stopped**.

- **When you use Backup-Based Disaster Recovery with both a local peer and a cross-region (remote) peer:**

  With Backup-Based Disaster Recovery enabled with both a local peer and a cross-region peer, and the local peer is available, Oracle recommends that you attempt a manual failover to the local peer first (not to the cross-region peer).

  If a local peer is unavailable or a manual failover to the local peer fails, you can perform a manual switchover to the cross-region peer. If the switchover to the cross-region peer fails, on the peer the Oracle Cloud Infrastructure console shows a **failover** link in the **Role** field that you can click to initiate a manual failover to the peer.

When you initiate a manual failover it is failed over to the peer based on the Recovery Time Objective (RTO) and Recovery Point Objective (RPO) targets. See Backup-Based Disaster Recovery Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for more information.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To initiate a manual failover to a cross-region backup copy, do the following:

1. On the peer, perform a switchover. See Perform a Switchover to a Local Backup Copy Peer for details.
2. If the switchover attempt in Step 1 fails, on the peer the **Role** field shows a **Failover** link. On the peer, click the **Failover** link.

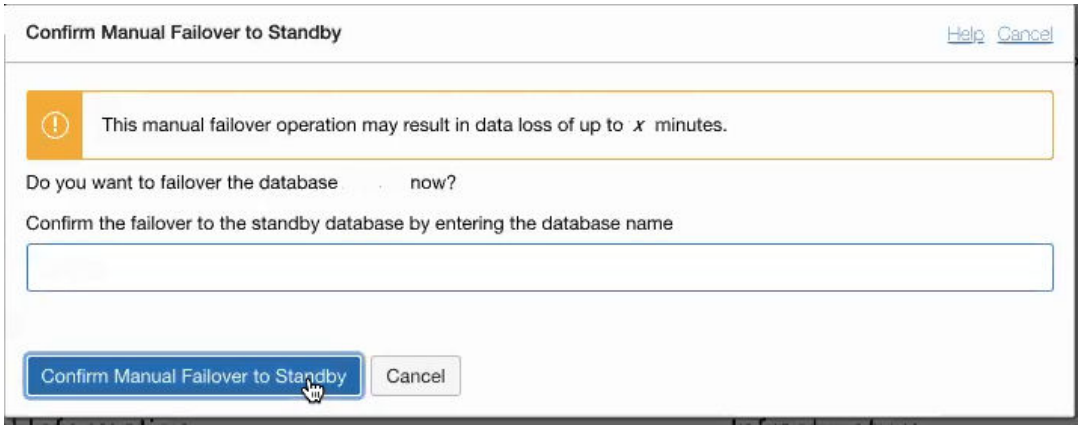
   This shows the **Confirm manual failover to peer** dialog, along with information on possible data loss that may result if you perform the manual failover to the peer.
3. In the **Confirm manual failover to peer** dialog, enter the Autonomous Database name to confirm that you want to failover.
4. In the **Confirm manual failover to peer** dialog, click **Confirm manual failover to peer**.

   When concurrent operations such as scaling are active, the confirmation also confirms either pausing or canceling the concurrent operation. See Concurrent Operations on Autonomous Database for more information.

To initiate a manual failover when the Primary database is unavailable and the local Peer is available, do the following:

1. On the **Details** page, under **Disaster recovery**, in the **Role** field, click **Failover**.

   This shows the **Confirm manual failover to peer** dialog, along with information on possible data loss that may result if you perform the manual failover to peer.



2. In the **Confirm manual failover to peer** dialog, enter the Autonomous Database name to confirm that you want to failover.
3. In the **Confirm manual failover to peer** dialog, click **Confirm manual failover to peer**.

When the failover completes, Backup-Based Disaster Recovery does the following:

- After a manual failover operation completes you can see any data loss associated with the manual failover in the message on the Oracle Cloud Infrastructure console banner and if you hover over the ⓘ in the **Role** field. The manual failover data loss is specified in minutes.
- After a manual failover with Backup-Based Disaster Recovery, if there was a regional failure, when the region comes back online the peer is automatically reconnected or if required reprovisioned.
- After you perform a manual failover to the cross-region peer, the cross-region peer becomes the primary database. In this case, if a local Autonomous Data Guard standby was enabled, a local Standby will be created and attached. If a local Autonomous Data Guard was not enabled before the failover in the primary database, as is the default, you have a local backup copy.

Notes for Performing a Switchover to a Backup Copy Peer

Provides notes for Backup-Based Disaster Recovery switchover:

- For a cross-region switchover, you must initiate the switchover from the cross-region peer.
- During the switchover, most of the actions on the Oracle Cloud Infrastructure Console are not available and the Autonomous Database Information page shows the **Lifecycle state** with the value **Updating**.
- The switchover operation keeps the original state of the Primary database. If the Primary database is stopped when you perform a switchover, the Primary database is stopped after the switchover.
- Autonomous Database generates the Switchover Autonomous Database work request. To view the request, under **Resources** click **Work Requests**.
- After a switchover or failover to the peer, the peer becomes the Primary and the graphs on the **Database Dashboard** card in Database Actions and the Oracle Cloud Infrastructure Metrics display information about the Primary database. The graphs and metrics do not contain information about the database that was the Primary before the switchover or failover operation.
- You cannot cancel a cross-region switchover operation after the switchover begins and the **State** shows **Role change in progress**. Your options to cancel the switchover are:

- Try or retry a switchover or a failover until the operation succeeds.
- File a service request at Oracle Cloud Support or contact your support representative.

Backup-Based Disaster Recovery Cross-Region Operations

Describes cross-region operations with Backup-Based Disaster Recovery.

- Add a Cross-Region Disaster Recovery Peer
  In addition to a local Backup-Based Disaster Recovery peer, you can add one or more remote region (cross-region) Backup-Based Disaster Recovery peers.
- Enable or Disable Backup Replication for an Existing Cross-Region Peer
- Convert Cross-Region Backup-Based Disaster Recovery Peer to Snapshot Standby
  A cross-region Backup-Based Disaster Recovery peer can be converted to a snapshot standby. This converts the peer to a read-write database for up to two days.
- Disconnect a Cross-Region Backup-Based Disaster Recovery Peer
  Shows you the steps to disconnect a cross-region Backup-Based Disaster Recovery peer from the Primary database.
- Disable a Cross-Region (Remote) Peer
  Describes the steps to terminate a cross-region (remote) peer.
- Update Remote Peer Network ACLs for a Backup-Based Disaster Recovery Peer
- Notes for Cross-Region Backup-Based Disaster Recovery Peer Operations
  Lists notes and restrictions for adding and managing a cross-region Backup-Based Disaster Recovery peer.

Add a Cross-Region Disaster Recovery Peer

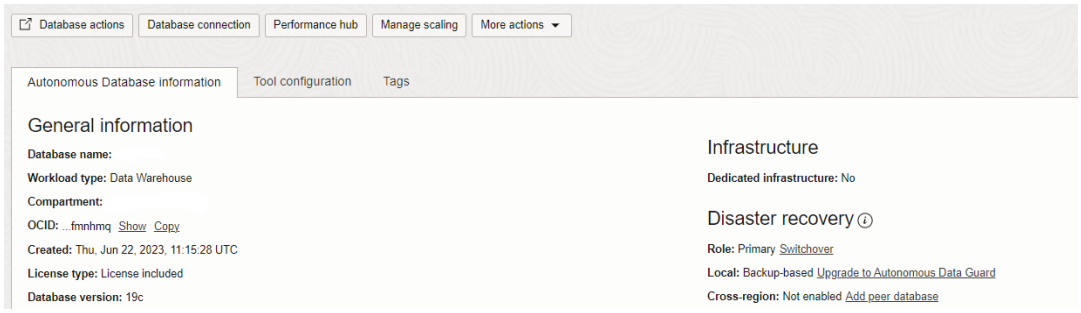
In addition to a local Backup-Based Disaster Recovery peer, you can add one or more remote region (cross-region) Backup-Based Disaster Recovery peers.

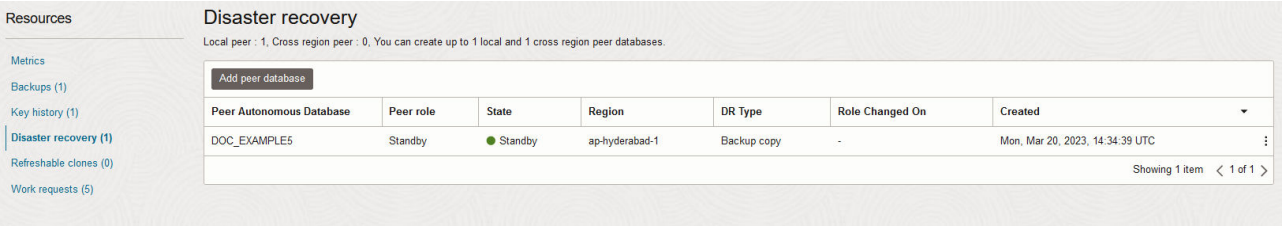
Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the  next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To add a cross-region Backup-Based Disaster Recovery peer, do the following:

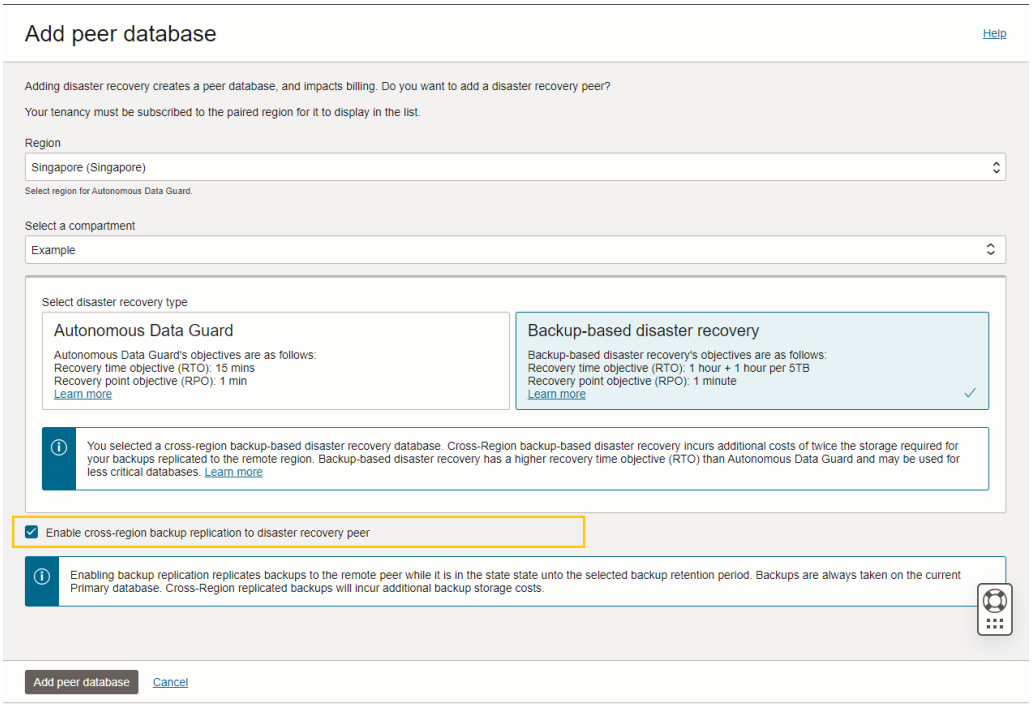
1. On the Autonomous Database Details page, under **Disaster recovery**, in the **Cross-region** field, click **Add peer database**, alternatively in the **Resources** area click **Disaster recovery**.



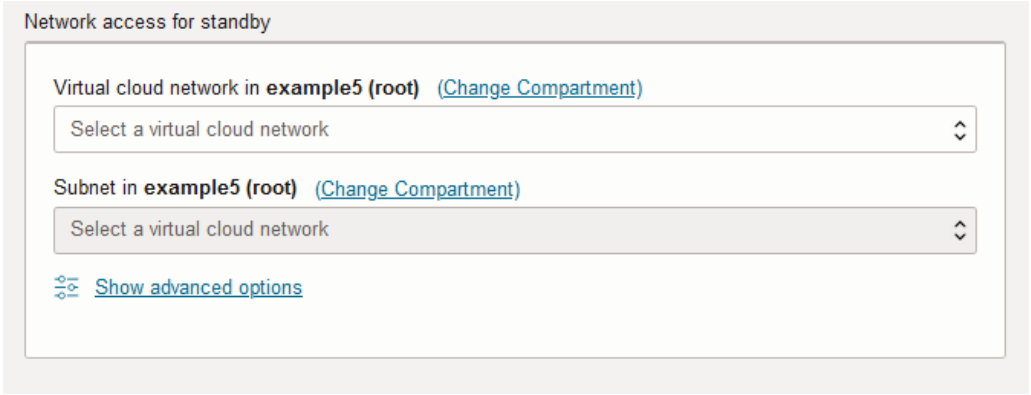


2. Click **Add peer database**.
3. In the **Region** drop-down list, select a region.

   The region list shows the available remote regions where you can create a cross-region peer. When you add a peer database, the list of available regions only shows a remote region if your tenancy is subscribed to the remote region (you must be subscribed to the paired remote region). See Autonomous Database Cross-Region Paired Regions for more information.
4. In the **Select a compartment** drop-down list, select a compartment.
5. Select the disaster recovery type. In addition, when the source database is configured with a private endpoint, enter the private endpoint information for the peer.
   a. Select the disaster recovery type: **Backup-based disaster recovery**.
   b. If you want to enable Cross-Region Backup Replication select the **Enable cross-region backup replication to disaster recovery peer** check-box. See Replicating Backups to a Cross-Region Backup Based Disaster Recovery Peer for more information.



c. When the source database is configured with a private endpoint, in the **Network access for standby** area enter the **Virtual cloud network** and the **Subnet**.



In these Network access for standby fields you specify the private endpoint's VCN and Subnet on the remote region where the standby is created. Configure Private Endpoints.

> **Note:**
>
> If you change your network access on the source database to enable a private endpoint after the standby is created, you must manually access the standby to enable a private endpoint on the peer.

6. Click **Add peer database**.

   The Autonomous Database Lifecycle state changes to **Updating**. In the **Resources** area, the number next to **Disaster recovery** increments to show that you have another disaster recovery peer, and the **State** field shows **Provisioning** for the new cross-region peer.

> **Note:**
>
> While you add a new peer, the primary database is available for read/write operations. There is no downtime on the primary database.

Autonomous Database generates a work request when you add a cross-region peer To view the work request, under **Resources** click **Work Requests**.

See Notes for Cross-Region Backup-Based Disaster Recovery Peer Operations for more information.

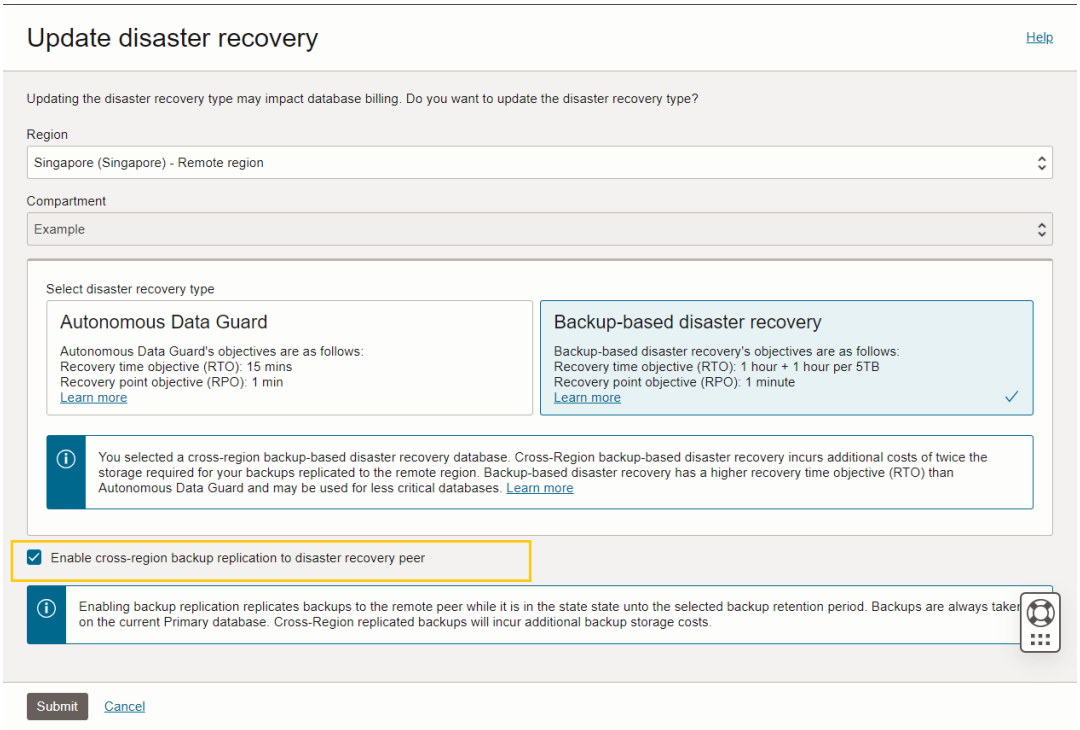
Enable or Disable Backup Replication for an Existing Cross-Region Peer

You can enable or disable backup replication on a Backup-Based Disaster Recovery cross-region peer.

To enable or disable backup replication for an existing cross-region Autonomous Data Guard standby:

1. On the **Autonomous Database Details** page, in the **Resources** area select **Disaster recovery**.
2. In a row that lists a cross-region standby, click ⋮ at the end of the row and select **Update disaster recovery**.

   This shows the Update disaster recovery page.



3. Enable or disable backup replication.
   a. If cross-region backup replication is disabled, select **Enable cross-region backup replication to disaster recovery peer** to enable the option.

b. If cross-region backup replication is enabled, deselect **Enable cross-region backup replication to disaster recovery peer** to disable the option.

4. Click **Submit**.

   The Autonomous Database Lifecycle state changes to **Updating**.

If you select **Enable cross-region backup replication to disaster recovery peer**, it can take between several minutes and several hours to replicate the backups to the remote region, depending on the size of the backups. After backups are replicated, when you select **Backups** under **Resources** on the peer database's Oracle Cloud Infrastructure Console, you will see the list of replicated backups.

Convert Cross-Region Backup-Based Disaster Recovery Peer to Snapshot Standby

A cross-region Backup-Based Disaster Recovery peer can be converted to a snapshot standby. This converts the peer to a read-write database for up to two days.

See Convert Cross-Region Peer to Snapshot Standby for more information.

Disconnect a Cross-Region Backup-Based Disaster Recovery Peer

Shows you the steps to disconnect a cross-region Backup-Based Disaster Recovery peer from the Primary database.

> **Note:**
>
> The disconnect operation for a cross-region Backup-Based Disaster Recovery peer can only be performed on an Autonomous Database instance that uses the ECPU compute model.

When you disconnect a cross-region peer, the peer database is disassociated from the primary database. This converts the database from a peer database to a standalone database. Following the disconnect operation you are not allowed to reconnect to the Primary.

The steps to disconnect a Backup-Based Disaster Recovery peer standby are the same as those to disconnect a standby database. See Disconnect a Peer Database for more information.

Notes for disconnecting a remote peer.

- The disconnect operation for a remote peer can only be performed on an Autonomous Database instance that uses the ECPU compute model.
- After the disconnect operation, the standalone database is no longer associated with the database that was the Primary database. To use the database as a standalone database, you must know the name of the database that was disconnected from the Primary database.
- There is no reconnect operation. After you disconnect a snapshot standby you are not allowed to reconnect to the Primary.
- After the disconnect operation, the standalone database is no longer associated with the database that was the Primary database. To use the database as a standalone database, you must know the name of the database that was disconnected from the Primary database.
- After the disconnect operation, the standalone database starts taking new backups as a standalone database.

Disable a Cross-Region (Remote) Peer

Describes the steps to terminate a cross-region (remote) peer.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To terminate a cross-region (remote) peer:

1. On the Primary database, on the Autonomous Database Details page under **Resources**, select Disaster recovery.
2. Access the Oracle Cloud Infrastructure Console for the remote region peer.

   The **Disaster recovery** information area shows the **Peer Autonomous Database**. The remote peer has the same display name as the primary database, with an "*_region*" extension. Where *region* is the region name, such as `IAD` or `BOM`.

   If you created the cross-region peer before the introduction of support for multiple cross-region peers, the remote peer has the same display name as the primary database, with an "`_Remote`" extension.

   Under **Peer Autonomous Database**, click the link for the remote peer to access the cross-region peer.
3. On the Oracle Cloud Infrastructure Console for the remote peer, on the **Details** page, from the **More actions** drop-down list, select **Terminate**.
4. On the **Terminate Autonomous Database** page, enter the database name to confirm that you want to terminate the cross-region peer.
5. Click **Terminate Autonomous Database**.

   While the peer is terminating, the **Lifecycle state** changes to **Terminating**.

There are limitations for disabling when an instance has a cross-region Backup-Based Disaster Recovery peer, as follows:

- A peer in the remote region cannot be disabled from the primary database.
- When Backup-Based Disaster Recovery is enabled with a cross-region peer, you must terminate all cross-region disaster recovery peers before you terminate the primary role database. If you attempt to terminate the primary, the system gives you an error.

  In this case, after you terminate all cross-region (remote) peers, you can terminate the primary database.

Update Remote Peer Network ACLs for a Backup-Based Disaster Recovery Peer

You can independently modify network ACLs on a remote disaster recovery peer database.

By default the disaster recovery primary and remote peer databases use the same network Access Control Lists (ACLs). Optionally, you can configure ACLs independently on remote peer databases. This provides an option to use different ACLs on remote peer databases.

If you modify the ACLs on a remote peer, Autonomous Database no longer keeps the ACL configuration synchronized between the primary and the remote peer. After you modify the ACLs on a remote peer, Autonomous Database manages the ACLs on the remote peer independently.

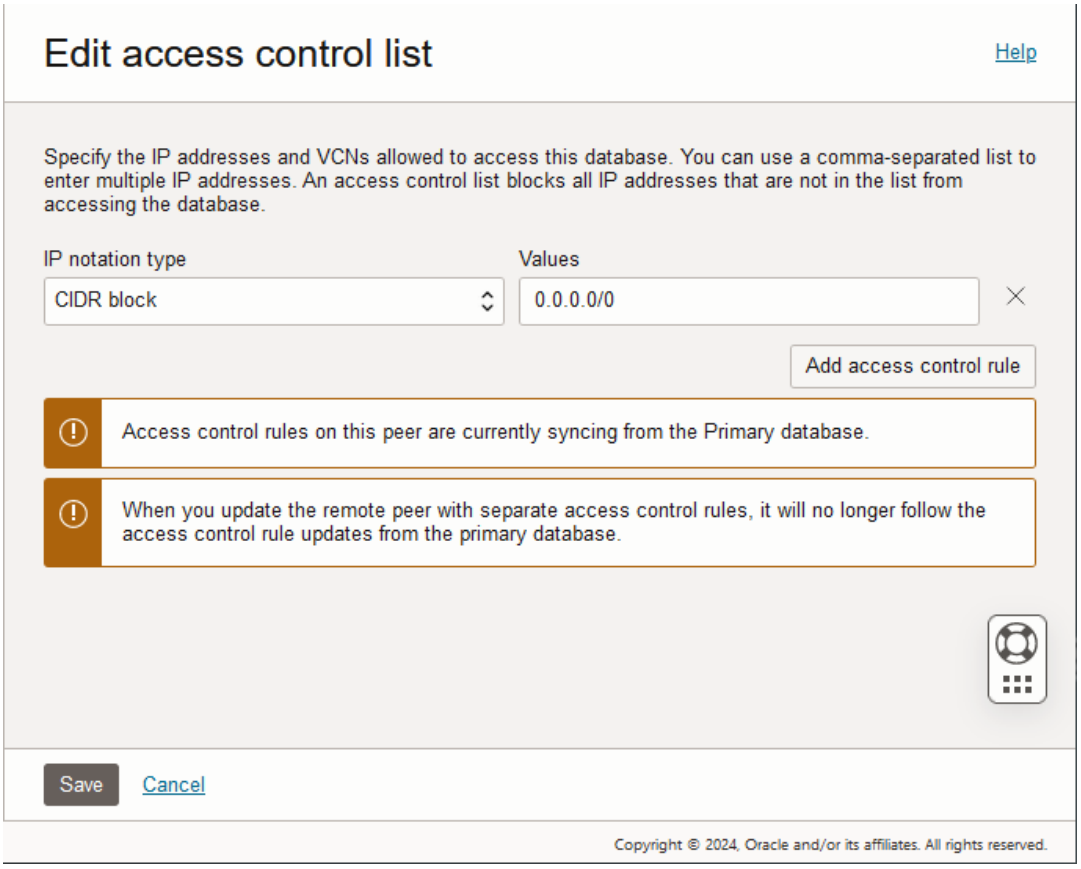
To use different network ACLs on a remote Autonomous Database peer:

1. On the primary database, on the Autonomous Database Details page, under **Resources** select **Disaster recovery**.
2. Access the remote peer.

   The **Disaster recovery** information area shows the **Peer Autonomous Database**. The remote peer database by default has the same display name as the primary database, with an extension. For example, `DBNAME_remote`.

   Under **Peer Autonomous Database**, click the link to access a cross-region peer.
3. On the remote peer database, edit the access control list.

   Before you change ACL the dialog shows a message indicating that ACLs on the peer database are syncing from the Primary database. For example:



   See Configure Access Control Lists for an Existing Autonomous Database Instance for more information.
4. Add, remove, or modify one or more ACLs.
5. Click **Save**.

After you modify ACLs, the ACLs on the primary and on the remote peer are managed separately.

If you want to restart the synchronization of ACLs between the primary and the remote peer, you have two options:

- Terminate the peer Autonomous Database and create a new cross-region disaster recovery peer database.

  See Disable a Cross-Region (Remote) Peer for details on terminating a remote peer database.
- Contact Oracle Cloud Support and file a service request or contact your support representative.

Notes for Cross-Region Backup-Based Disaster Recovery Peer Operations

Lists notes and restrictions for adding and managing a cross-region Backup-Based Disaster Recovery peer.

- After you add a cross-region (remote) peer, the wallet and connection string from the primary database will contain only the hostname of the primary database, and the wallet and connection string from the remote database will contain only the hostname of the remote database. This applies to both instance and regional wallets.

  See Cross-Region Disaster Recovery Connection Strings and Wallets for more information.
- If you enable the option **Enable cross-region backup replication to disaster recovery peer**, it can take between several minutes and several hours to replicate the backups to the remote region, depending on the size of the backups. After backups are replicated, when you select **Backups** under **Resources** on the peer database's Oracle Cloud Infrastructure Console, you will see the list of replicated backups.

  See Enable or Disable Backup Replication for an Existing Cross-Region Peer for more information.
- When you add a cross region peer there are special considerations when the primary instance is using customer-managed keys, or if you want to switch to use customer-managed keys. See Autonomous Data Guard with Customer Managed Keys for more information.
- When you add a Backup-Based Disaster Recovery peer, while the **Lifecycle state** field shows **Updating**, the following actions are disabled on the Primary database:
  - Move Resource. See Move an Autonomous Database to a Different Compartment for information on moving an instance.
  - Stop. See Stop Autonomous Database for information on stopping an instance.
  - Restart. See Restart Autonomous Database for information on restarting an instance.
  - Restore. See Restore and Recover your Autonomous Database for information on restoring.

Backup-Based Disaster Recovery Events

You can use Oracle Cloud Infrastructure events to respond when Autonomous Database changes its state due to a Backup-Based Disaster Recovery related event.

Autonomous Database events include the following:

- Begin manual failover
- End manual failover with failover result of success or failure.
- Begin switchover

- End switchover with switchover result of success or failure.

Based on events you can perform actions or send notifications. See Events and Notifications for a Standby Database for more information on using events and producing notifications.

Use the API

Provides links for details on using API operations to manage Backup-Based Disaster Recovery.

For information about using the API and signing requests, see REST APIs and Security Credentials.

For information about SDKs, see Software Development Kits and Command Line Interface.

Use these API operations to manage Disaster Recovery:

- To enable or disable Backup-Based Disaster Recovery, use UpdateAutonomousDatabase.
- To initiate a manual failover operation, use FailOverAutonomousDatabase.
- To initiate a switchover operation, use SwitchOverAutonomousDatabase.

Use these Terraform APIs to manage Autonomous Database resources:

For information about Terraform, see Terraform Provider and for information about Terraform APIs, see Data Source: oci_database_autonomous_database.

Track Table Changes with Flashback Time Travel

Use Flashback Time Travel to view past states of database objects or to return database objects to a previous state without using point-in-time media recovery.

- About Flashback Time Travel
  Flashback Time Travel lets you track and store transactional changes to a table over its lifetime. Flashback Time Travel is useful for compliance with record stage policies and audit reports. You can also use Flashback Time Travel in various scenarios such as enforcing digital shredding, accessing historical data, and selective data recovery.
- Enable Flashback Time Travel for a Table
  Describes steps to enable Flashback Time Travel for a table in Autonomous Database.
- Disable Flashback Time Travel for a Table
  Describes procedure to disable Flashback Time Travel for a table in Autonomous Database.
- Modify the Retention Time for Flashback Time Travel
  Describes procedure to modify the retention time for Flashback Time Travel in Autonomous Database.
- Purge Historical Data for Flashback Time Travel
  Describes procedure to purge historical data for Flashback Time Travel in Autonomous Database.
- View Flashback Time Travel Information
  Describes data dictionary views to view information about Flashback Time Travel files in Autonomous Database.
- Notes for Flashback Time Travel
  Provides notes and restrictions for using Flashback Time Travel on Autonomous Database.

About Flashback Time Travel

Flashback Time Travel lets you track and store transactional changes to a table over its lifetime. Flashback Time Travel is useful for compliance with record stage policies and audit reports. You can also use Flashback Time Travel in various scenarios such as enforcing digital shredding, accessing historical data, and selective data recovery.

Flashback Time Travel was called Flashback Data Archive in previous Oracle Database versions. The Flashback Time Travel APIs and some Flashback Time Travel terms such as the `FLASHBACK ARCHIVE` privilege retain the Flashback Data Archive naming.

Each Autonomous Database instance has one Flashback Archive named `flashback_archive` that supports Flashback Time Travel operations. The default retention time for `flashback_archive` is 60 days. See Modify the Retention Time for Flashback Time Travel for information on changing the default retention time.

There are restrictions to enable Flashback Time Travel for a table, including the following:

- You must have the `FLASHBACK ARCHIVE` object privilege for `flashback_archive`. By default the `ADMIN` user has this privilege.
- The table must not be any of the following:
  - A nested table
  - A temporary table
  - A remote table
  - An external table
- The table cannot contain `LONG` columns.
- The table cannot contain nested columns.

See Notes for Flashback Time Travel for a list of additional Flashback Time Travel restrictions.

See Using Oracle Flashback Technology for more information on Flashback Technology.

Enable Flashback Time Travel for a Table

Describes steps to enable Flashback Time Travel for a table in Autonomous Database.

By default, Flashback Time Travel is disabled for new and existing tables.

The following is the prerequisite to enable Flashback Time Travel:

- You must be logged in as the `ADMIN` user or have the `FLASHBACK ARCHIVE` object privilege.

  The `ADMIN` user or the user having the `FLASHBACK ARCHIVE` privilege `WITH GRANT OPTION` can grant the `FLASHBACK ARCHIVE` object privilege to another user.

  For example:

  ```
  GRANT FLASHBACK ARCHIVE ON FLASHBACK_ARCHIVE TO <username>
  ```

  ```
  GRANT FLASHBACK ARCHIVE ON FLASHBACK_ARCHIVE TO <username> WITH GRANT
  OPTION
  ```

  See Manage User Roles and Privileges on Autonomous Database for more information.

You can enable Flashback Time Travel for an existing table or a new table that you are creating.

1. To enable Flashback Time Travel when you create a table:

```
CREATE TABLE employee (EMPNO NUMBER(4) NOT NULL, ENAME VARCHAR2(10), JOB
VARCHAR2(9), MGR NUMBER(4)) FLASHBACK ARCHIVE;
```

2. To enable the Flashback Time Travel for an existing table:

```
ALTER TABLE departments FLASHBACK ARCHIVE;
```

After you enable Flashback Time Travel for a table you can disable Flashback Time Travel. See Disable Flashback Time Travel for a Table for more information.

Disable Flashback Time Travel for a Table

Describes procedure to disable Flashback Time Travel for a table in Autonomous Database.

After Flashback Time Travel is enabled for a table you can disable it as the `ADMIN` user or if you have the `FLASHBACK ARCHIVE` object privilege.

For example, to disable Flashback Time Travel for the `employee` table:

```
ALTER TABLE employee NO FLASHBACK ARCHIVE;
```

Modify the Retention Time for Flashback Time Travel

Describes procedure to modify the retention time for Flashback Time Travel in Autonomous Database.

You can modify the retention time for Flashback Time Travel if you are logged in as `ADMIN` user or if you have `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.

For example, to set the retention time to 365 days:

```
BEGIN
  DBMS_CLOUD_ADMIN.SET_FLASHBACK_ARCHIVE_RETENTION
   (retention_days => 365);
END;
/
```

See SET_FLASHBACK_ARCHIVE_RETENTION Procedure for detailed information about the procedure.

Purge Historical Data for Flashback Time Travel

Describes procedure to purge historical data for Flashback Time Travel in Autonomous Database.

You can purge historical data for Flashback Time Travel if you are logged in as the `ADMIN` user or if you have `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.

For example, to purge all Flashback Time Travel historical data:

```
BEGIN
  DBMS_CLOUD_ADMIN.PURGE_FLASHBACK_ARCHIVE
   (scope => 'ALL');
END;
/
```

Example to purge the historical Flashback Time Travel data before a specified timestamp:

```
BEGIN
  DBMS_CLOUD_ADMIN.PURGE_FLASHBACK_ARCHIVE
   (scope => 'TIMESTAMP', before_timestamp => '12-JUL-2023 10:24:00');
END;
/
```

Example to purge Flashback Time Travel historical data older than 1 day:

```
BEGIN
  DBMS_CLOUD_ADMIN.PURGE_FLASHBACK_ARCHIVE
   (scope => 'TIMESTAMP', before_timestamp => SYSTIMESTAMP - INTERVAL '1'
DAY);
END;
/
```

Example to purge the historical Flashback Time Travel data before the specified scn:

```
BEGIN
  DBMS_CLOUD_ADMIN.PURGE_FLASHBACK_ARCHIVE
   (scope => 'SCN',before_scn=> '5928826');
END;
/
```

See PURGE_FLASHBACK_ARCHIVE Procedure for detailed information about the procedure.

View Flashback Time Travel Information

Describes data dictionary views to view information about Flashback Time Travel files in Autonomous Database.

You can view information about Flashback Time Travel files in static data dictionary views.

View	Description
*_FLASHBACK_ARCHIVE	Displays information about Flashback Time Travel files.
*_FLASHBACK_ARCHIVE_TS	Displays tablespaces of Flashback Time Travel files.
*_FLASHBACK_ARCHIVE_TABLES	Displays information about tables that have Flashback Time Travel enabled.

Notes for Flashback Time Travel

Provides notes and restrictions for using Flashback Time Travel on Autonomous Database.

Note the following Flashback Time Travel restrictions:

- You cannot enable Flashback Time Travel on tables with `LONG` data type or nested table columns.
- You cannot enable Flashback Time Travel on a nested table, temporary table, external table, materialized view, Advanced Query (AQ) table, hybrid partitioned tables, or non-table object.
- Flashback Time Travel does not support DDL statements that move, split, merge, or coalesce partitions or sub partitions, move tables, or convert `LONG` columns to `LOB` columns.
- Adding or enabling a Constraint, including Foreign Key Constraint, on a table that has been enabled for Flashback Time Travel fails with ORA-55610.

  Dropping or disabling a Constraint (including Foreign Key Constraint) on a table that has been enabled for Flashback Time Travel is supported.
- After enabling Flashback Time Travel on a table, Oracle recommends initially waiting at least 20 seconds before inserting data into the table and waiting up to 5 minutes before using Flashback Query on the table.
- There is one Flashback Data Archive per Autonomous Database instance, named `flashback_archive`, and you cannot create additional Flashback Data Archives.
- You cannot drop the Flashback Data Archive, `flashback_archive`, within an Autonomous Database instance.
- You cannot create, modify, or drop tablespaces for the Flashback Data Archive. Hence, you cannot run these statements:
  - `ALTER FLASHBACK ARCHIVE flashback_archive ADD TABLESPACE;`
  - `ALTER FLASHBACK ARCHIVE flashback_archive MODIFY TABLESPACE;`
  - `ALTER FLASHBACK ARCHIVE flashback_archive REMOVE TABLESPACE;`
- If you enable Flashback Time Travel on a table, but Automatic Undo Management (AUM) is disabled, error ORA-55614 occurs when you try to modify the table.
- To enable Flashback Time Travel on a table, the table cannot use any of the following Flashback Time Travel reserved words as column names:
  - `STARTSCN`
  - `ENDSCN`
  - `RID`
  - `XID`
  - `OP`
  - `OPERATION`

Use OCI Full Stack Disaster Recovery with Autonomous Database

You can enable the Oracle Cloud Infrastructure Full Stack Disaster Recovery on Autonomous Database.

When you enable the OCI Full Stack Disaster Recovery (Full Stack Disaster Recovery) you can use Full Stack Disaster Recovery to perform switchover/failover operations or optionally perform database only cross-region Autonomous Database switchover/failover operations.

> **Note:**
>
> Full Stack Disaster Recovery can only be enabled on an Autonomous Database instance that uses the ECPU compute model.

See Full Stack Disaster Recovery for more information.

Topics

- Prerequisite Steps for Autonomous Database with Full Stack Disaster Recovery
  If you enable Full Stack Disaster Recovery and you want your disaster recovery topology to contain an Autonomous Database instance, first perform these prerequisite steps.
- Enable Full Stack Disaster Recovery with Autonomous Database
  Enable Full Stack Disaster Recovery by adding an Autonomous Database instance to a Disaster Recovery Protection Group.
- Use Full Stack Disaster Recovery to Perform Switchover or Failover Operations
  After you enable Full Stack Disaster Recovery, when you perform a cross-region switchover or failover you have the option to use Full Stack Disaster Recovery.
- Notes for Full Stack Disaster Recovery
  Note the following for using Autonomous Database with Full Stack Disaster Recovery enabled.

Prerequisite Steps for Autonomous Database with Full Stack Disaster Recovery

If you enable Full Stack Disaster Recovery and you want your disaster recovery topology to contain an Autonomous Database instance, first perform these prerequisite steps.

1. Add one or more cross-region Autonomous Data Guard peers to your Autonomous Database instance.

   See Add a Cross-Region Standby Database for more information.

2. Set up Full Stack Disaster Recovery and configure the required OCI Identity and Access Management policies.

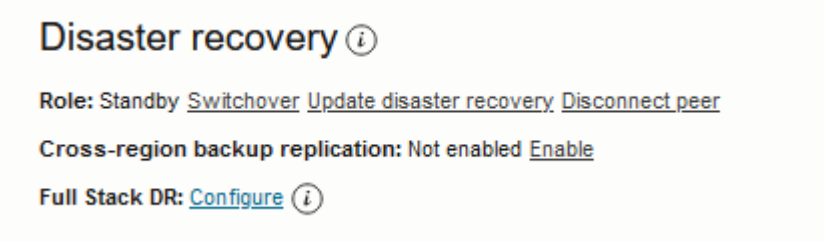
   See Prerequisites for Full Stack Disaster Recovery and Policies for Full Stack Disaster Recovery for more information.

Enable Full Stack Disaster Recovery with Autonomous Database

Enable Full Stack Disaster Recovery by adding an Autonomous Database instance to a Disaster Recovery Protection Group.

1. Create Disaster Recovery Protection Groups in a primary region and in one or more standby regions.

You can create a Disaster Recovery Protection Group from the Oracle Cloud Infrastructure Console. To access Full Stack Disaster Recovery, on the Autonomous Database details page click **Configure** in the **Full Stack DR** field under Disaster recovery.



See Manage Disaster Recovery Protection Groups for more information.

2. Add an Autonomous Database instance to the Disaster Recovery Protection Group in primary region.

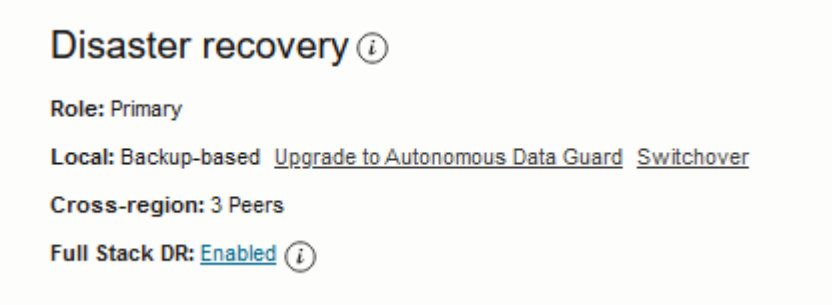
   See Add Members to a Disaster Recovery Protection Group for more information.

3. Add at least one Autonomous Data Guard cross-region standby database to a Disaster Recovery Protection Group in a remote (standby) region.

   > **Note:**
   >
   > Full Stack Disaster Recovery does not support using a local, same region, Autonomous Data Guard standby.

   See Add Members to a Disaster Recovery Protection Group for more information.

After you create a Disaster Recovery Protection Group that includes an Autonomous Database, the Oracle Cloud Infrastructure Console shows that Full Stack Disaster Recovery is enabled.



Click **Enabled** in the **Full Stack DR** field to view the Disaster Recovery Protection Group details.

Use Full Stack Disaster Recovery to Perform Switchover or Failover Operations

After you enable Full Stack Disaster Recovery, when you perform a cross-region switchover or failover you have the option to use Full Stack Disaster Recovery.

For a cross-region switchover or failover you must initiate the operation from a Standby database. You have several options to access a Standby database:

- Select the remote region in Oracle Cloud Infrastructure Console and then access the Standby database.
- On the Primary database you can access a standby database from the Oracle Cloud Infrastructure Console. To access a standby database, from the Autonomous Database details page select **Disaster recovery** under **Resources** and click a link for a standby database in the **Peer Autonomous Database** column.

To perform a switchover or a failover:

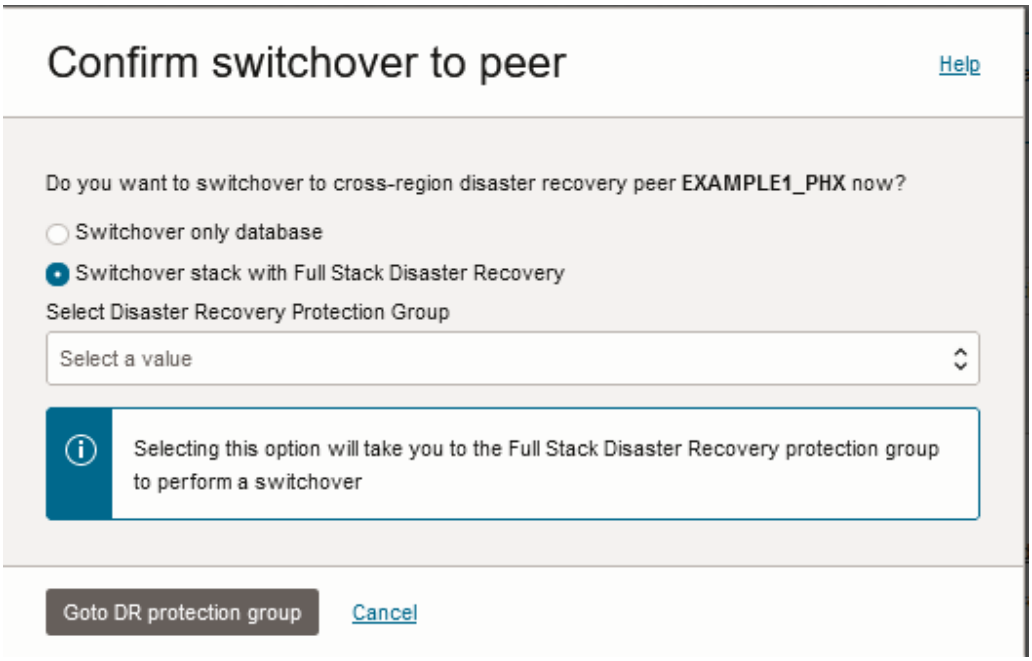
1. Access the cross-region standby database and from the Autonomous Database details page, under **Disaster recovery** click **Switchover**(or **Failover** when the Primary is not available) in the **Role** field.

   As an alternative, to initiate a switchover you can select **More actions** and **Switchover**(or **Failover** when the Primary is not available).

   This shows the Confirm switchover to peer dialog.

2. Select an option and confirm the operation.

   a. Select **Switchover stack with Full Stack Disaster Recovery**.



   b. In the **Select Disaster Recovery Protection Group** field, select a value.

   c. Click **Goto DR protection group**.

   This brings you to the Disaster Recovery DR protection group details page. From this page you can perform disaster recovery actions. See Manage Disaster Recovery Plan Executions for more information.

Notes for Full Stack Disaster Recovery

Note the following for using Autonomous Database with Full Stack Disaster Recovery enabled.

- When Full Stack Disaster Recovery is enabled, depending on the configuration of the Disaster Recovery Protection Group, a cross-region Standby database should be part of a larger disaster recovery configuration. If you perform disaster recovery operations just on

the database, this can impact other Oracle Cloud Infrastructure resources that are part of the Disaster Recovery Protection Group. For example, if you select the **Switchover only database option**, Full Stack Disaster Recovery receives an event indicating that a database switchover has occurred. In this case Full Stack Disaster Recovery can automatically trigger a switchover of the remaining stack. The action that Full Stack Disaster Recovery takes when it receives the database switchover event is configurable with Full Stack Disaster Recovery.

- Full Stack Disaster Recovery can only be enabled on an Autonomous Database instance that uses the ECPU compute model.

  See Compute Models in Autonomous Database for more information.

- Full Stack Disaster Recovery does not support disaster recovery using Autonomous Database Backup-Based Disaster Recovery peers.

  See Use Backup-Based Disaster Recovery for more information.

Security

Autonomous Database stores all data with inbuilt security. Only authenticated users and applications can access the data when they connect to the database.

- Security Features Overview
  Describes some of the robust security features that Autonomous Database provides.
- Security and Authentication in Oracle Autonomous Database
  Oracle Autonomous Database stores all data in encrypted format in the database. Only authenticated users and applications can access the data when they connect to the database.
- Security Testing Policies on Autonomous Database
  Autonomous Database is subject to the Oracle Cloud Security Testing policy which describes when and how you may conduct certain types of security testing of Oracle Cloud Infrastructure services, including vulnerability and penetration tests, as well as tests involving data scraping tools.
- Manage Encryption Keys on Autonomous Database
  Describes how to use customer-managed encryption keys with Autonomous Database, and if you are using customer-managed encryption keys, shows how to rotate the keys, switch to Oracle-managed encryption keys, or view the encryption key history.
- Configure Network Access with Access Control Rules (ACLs) and Private Endpoints
  Provides details on how to configure network access with access control rules or using a private endpoint and describes the secure client connection options. Also describes how to enable support for TLS connections (require mutual TLS only or allow both mutual TLS and TLS authentication).
- Audit Autonomous Database
  Autonomous Database provides auditing that allows you to monitor Oracle database activities.
- Rotate Wallets for Autonomous Database
- Use Oracle Database Vault with Autonomous Database
  Oracle Database Vault implements powerful security controls for your database. These unique security controls restrict access to application data by privileged database users, reducing the risk of insider and outside threats and addressing common compliance requirements.

- IAM Policies for Autonomous Database
  Provides information on IAM policies required for API operations on Autonomous Database.
- Use Oracle Data Safe with Autonomous Database
  Provides information on using Oracle Data Safe on Autonomous Database.
- Break Glass Access for SaaS on Autonomous Database
  Autonomous Database supports break glass access for SaaS providers. Break glass access allows a SaaS operations team, when explicitly authorized by a SaaS customer, to access a customer's database to perform critical or emergency operations.
- Encrypt Data While Exporting or Decrypt Data While Importing
  For greater security you can encrypt data that you export to Object Storage. When data on Object Storage is encrypted you can decrypt the data you import or when you use the data in an external table.
- Manage Oracle Cloud Operator Access
  Describes how to grant temporary access to your database for an Oracle Cloud Operator.

Security Features Overview

Describes some of the robust security features that Autonomous Database provides.

Autonomous Database security features include:

- Autonomous Database meets a broad set of international and industry-specific compliance standards, and as part of Oracle Cloud Infrastructure Autonomous Database has achieved attestations for the common compliance frameworks providing an independent assessment of the service's security, privacy, and compliance controls.

  See Regulatory Compliance Certification for more information.
- Autonomous Database applies security patches automatically as soon as they become available.

  See Configuration Management for more information.
- Autonomous Database is governed by the Oracle Cloud Hosting and Delivery Policies which explain the Oracle Cloud Security Policy. See the Delivery Policies area on Oracle Contracts for more information.
- Autonomous Database is subject to the Oracle Cloud Security Testing policy, which describes when and how you may conduct certain types of security testing of Oracle Cloud Infrastructure services, including vulnerability and penetration tests and tests involving data scraping tools.

  See Security Testing Policies on Autonomous Database for more information.
- Autonomous Database provides end-to-end encryption out of the box for the database, backups, and all network communication. All your data, including backups, are encrypted with AES256. You can use Oracle-managed or customer-managed keys to encrypt your data.

  See Manage Encryption Keys on Autonomous Database for more information.
- All network connections are encrypted using TLS 1.2. You can use mutual TLS or one-way TLS connections.
- Autonomous Database provides several options to control client access to your database. You can use public endpoints with access control lists to specify which clients can connect. You can also use private endpoints to place the database in your VCN and use security lists and network security groups to control access to the database.

See Client Access Control for more information.

- Autonomous Database provides fully automated immutable backups that cannot be tampered with by the users in your tenancy.

  See About Backup and Recovery on Autonomous Database for more information.

- Autonomous Database provides several user authentication methods. You can use local database user names and passwords or external authentication methods, including:
  - Oracle Cloud Infrastructure Identity and Access Management
  - Microsoft Active Directory
  - Azure Active Directory
  - Kerberos

  See Manage Users for more information.

- You can configure Oracle Database Vault to control access to specific schemas from privileged database users such as the ADMIN user.

  See Use Oracle Database Vault with Autonomous Database for more information,

- Autonomous Database provides robust auditing capabilities that enable you to track who did what on the service and on specific databases. You can configure database auditing to audit all actions, such as access to specific objects, schema changes, logons by specific users, and much more.

  See Auditing Overview on Autonomous Database for more information.

- Oracle Cloud Operators do not have authorization to access your data or any other information in your database schemas. When access to your database schemas is required to troubleshoot or mitigate an issue, you can allow a cloud operator to access your Autonomous Database schemas for a limited time.

  See Manage Oracle Cloud Operator Access for more information.

See Security and Authentication in Oracle Autonomous Database for more information.

Security and Authentication in Oracle Autonomous Database

Oracle Autonomous Database stores all data in encrypted format in the database. Only authenticated users and applications can access the data when they connect to the database.

> **Note:**
>
> Oracle Autonomous Database supports the standard security features of the Oracle Database including privilege analysis, network encryption, centrally managed users, secure application roles, transparent sensitive data protection, and others. Additionally, Oracle Autonomous Database adds Label Security, Database Vault, Data Safe, and other advanced security features at no additional cost.

- Configuration Management
  Oracle Autonomous Database provides standard, hardened security configurations that reduce the time and money managing configurations across your databases.
- Data Encryption
  Oracle Autonomous Database uses always-on encryption that protects data at rest and in transit. Data at rest and in motion is encrypted by default. Encryption cannot be turned off.

- Data Access Control
  Securing access to your Oracle Autonomous Database and your data uses several different kinds of access control:
- Auditing Overview on Autonomous Database
  Oracle Autonomous Database provides robust auditing capabilities that enable you to track who did what on the service and on specific databases. Comprehensive log data allows you to audit and monitor actions on your resources, which helps you to meet your audit requirements while reducing security and operational risk.
- Assessing the Security of Your Database and its Data
  Oracle Autonomous Database integrates with Oracle Data Safe to help you assess and secure your databases.
- Regulatory Compliance Certification
  Oracle Autonomous Database meets a broad set of international and industry-specific compliance standards.

Configuration Management

Oracle Autonomous Database provides standard, hardened security configurations that reduce the time and money managing configurations across your databases.

Security patches and updates are done automatically, so you don't spend time, money, or attention to keeping security up to date. These capabilities protect your databases and data from costly and potentially disastrous security vulnerabilities and breaches.

Data Encryption

Oracle Autonomous Database uses always-on encryption that protects data at rest and in transit. Data at rest and in motion is encrypted by default. Encryption cannot be turned off.

Encryption of Data at Rest

Data at rest is encrypted using TDE (Transparent Data Encryption), a cryptographic solution that protects the processing, transmission, and storage of data. Using AES256 tablespace encryption, each database has its own encryption key, and any backups have their own different encryption keys.

By default, Oracle Autonomous Database creates and manages all the master encryption keys used to protect your data, storing them in a secure PKCS 12 keystore on the same systems where the database resides. If your company security policies require, Oracle Autonomous Database can instead use keys you create and manage in the Oracle Cloud Infrastructure Vault service. For more information, see About Master Encryption Key Management on Autonomous Database.

Additionally, customer-managed keys can be rotated when needed in order to meet your organization's security policies.

Note: When you clone a database, the new database gets its own new set of encryption keys.

Encryption of Data in Transit

Clients (applications and tools) connect to Oracle Autonomous Database using supported protocols including SQL*Net, JDBC, and ODBC.

TCPS (Secure TCP) database connection services use the industry-standard TLS 1.2 (Transport Layer Security) protocol for connections and symmetric-key data encryption.

- With mTLS connections, Oracle Autonomous Database users download a connection wallet that contains all necessary files for a client to connect using TCPS. Distribute this wallet only to those users who need and are permitted to have database access. The client-side configuration uses information in the wallet to perform symmetric-key data encryption.
- Autonomous Database by default supports Mutual TLS (mTLS) connections. You have the option to configure an Autonomous Database instance to allow both mTLS and TLS connections. Using TLS connections, some clients, such as JDBC Thin Driver clients, do not need to download a wallet if you use a TLS connection string and TLS is enabled for the Autonomous Database instance.

See Secure Connections to Autonomous Database with mTLS or with TLS for more information.

Data Access Control

Securing access to your Oracle Autonomous Database and your data uses several different kinds of access control:

- Client Access Control
  Client access control for an Autonomous Database instance is enforced by network access control policies, through client connection protocols, and by the access rights of the database user the client uses to connect.
- Database User Access Control
  The Oracle Autonomous Database is configured with an administrative account, ADMIN, that is used to create and manage other database accounts. Oracle Autonomous Database provides a robust set of features and controls including system and object privileges and roles. User profiles allows you to customize password policies to define and implement a secure database user access strategy.
- Oracle Cloud User Access Control
  You use Identity and Access Management (IAM) services to control the privileges of your Oracle Cloud users by specifying the actions those users can perform on your Oracle Autonomous Database.
- Authorized Access on Autonomous Database
  Only authorized users are allowed access to an Autonomous Database instance.
- Autonomous Database Fully Managed Service
  Autonomous Database is a fully managed service and Oracle uses its own Oracle Cloud Infrastructure tenancies to run the Autonomous Database service.

Client Access Control

Client access control for an Autonomous Database instance is enforced by network access control policies, through client connection protocols, and by the access rights of the database user the client uses to connect.

Network Access Control

You define network access control when you set up and configure your Oracle Autonomous Database. There are two options to consider.

- **Private Endpoints and Security Lists:** This is the recommended option. Create your Oracle Autonomous Database in your virtual cloud network (VCN) using private endpoints. You control access to your database using security lists and network security groups allowing you to specify who can create connections to your database.

For detailed information about creating these resources, see Configure Network Access with Private Endpoints.

- **Public Endpoints and Access Control Lists:** Create your Oracle Autonomous Database using public endpoints allowing access from any client with client credentials. You control access to your database using network access control lists (ACLs) allowing you to specify IP addresses, CIDR blocks, or VCNs that can connect to your database. Public IPs are easier to discover and attack, and Private Endpoints are recommended where possible.

  For detailed information about setting up an ACL, see Configure Access Control Lists for an Existing Autonomous Database Instance.

Client Connection Control

Clients connect through a TCPS (Secure TCP) database connection using standard TLS 1.2 to secure the connection. Oracle Autonomous Database uses self-signed certificates. You can rotate the self-signed certificates from the Oracle Cloud Infrastructure console to meet your organization's security compliance needs. See Rotate Wallets with Immediate Rotation.

The access the client has to the database is restricted by the access rights of the database user the client uses to connect.

Database User Access Control

The Oracle Autonomous Database is configured with an administrative account, ADMIN, that is used to create and manage other database accounts. Oracle Autonomous Database provides a robust set of features and controls including system and object privileges and roles. User profiles allows you to customize password policies to define and implement a secure database user access strategy.

For basic information about standard user management, see User Accounts in *Oracle Database Concepts*. For detailed information and guidance, see Managing Security for Oracle Database Users in *Oracle Database Security Guide*.

If your database user access strategy demands more controls than are provided by standard user management, you can configure your Oracle Autonomous Database to use Database Vault to meet more rigorous requirements.

Using Microsoft Active Directory to Manage Database Users

If you use Microsoft Active Directory as a user repository, you can configure your database to authenticate and authorize Microsoft Active Directory users. This integration enables you to consolidate your user repository while still implementing a rigorous database user access strategy, regardless of whether you use standard user management or Database Vault.

For more information about integrating Microsoft Active Directory with your databases, see Use Microsoft Active Directory with Autonomous Database.

Database Vault

Oracle Database Vault comes preconfigured and ready to use. You can use its powerful security controls to restrict access to application data by privileged database users, reducing the risk threats, and addressing common compliance requirements.

You configure controls to block privileged account access to application data and control sensitive operations inside the database. You configure trusted paths to add additional security controls to authorized data access, database objects, and database commands. Database Vault secures existing database environments transparently, eliminating costly and time consuming application changes.

Before using Database Vault, be sure to review Use Oracle Database Vault with Autonomous Database to gain an understanding of the impact of configuring and enabling Database Vault.

For detailed information on implementing Database Vault features, refer to *Oracle Database Vault Administrator's Guide*.

Oracle Cloud User Access Control

You use Identity and Access Management (IAM) services to control the privileges of your Oracle Cloud users by specifying the actions those users can perform on your Oracle Autonomous Database.

The IAM service provides several kinds of components to help you define and implement a secure cloud user access strategy:

- **Compartment**: A collection of related resources. Compartments are a fundamental component of Oracle Cloud Infrastructure for organizing and isolating your cloud resources.
- **Group**: A collection of users who all need the same type of access to a particular set of resources or compartment.
- **Dynamic Group**: A special type of group that contains resources that match rules that you define. Thus, the membership can change dynamically as matching resources are created or deleted.
- **Policy**: A group of statements that specify who can access which resources, and how. Access is granted at the group and compartment level, which means you write a policy statement that gives a specific group a specific type of access to a specific type of resource within a specific compartment.

Of these, the policy is the primary tool you use to control access because it provides the "Who", "How", "What" and "Where" of a single access constraint. A policy statement has this format:

The format of a policy statement is:

```
Allow
  group <group-name>
  to <control-verb>
  <resource-type>
  in compartment <compartment-name>
```

- **`group <group-name>`** specifies the "Who" by providing the name of an existing IAM group.
- **`to <control-verb>`** specifies the "How" using one of these predefined control verbs:
  - **`inspect`**: the ability to list resources of the given type, without access to any confidential information or user-specified metadata that may be part of that resource.
  - **`read`**: `inspect` plus the ability to get user-specified metadata and the actual resource itself.
  - **`use`**: `read` plus the ability to work with existing resources, but not to create or delete them. Additionally, "work with" means different operations for different resource types.
  - **`manage`**: all permissions for the resource type, including creation and deletion.
- ***`<resource-type>`*** specifies the "What" using a predefined resource-type. The resource-type values for infrastructure resources are:
  - **`autonomous-databases`**

  - **autonomous-backups**

    You may create policy statements that refer to the **tag-namespaces** resource-type value if tagging is used in your tenancy.

- **in compartment *<compartment-name>*** specifies the "Where" by providing the name of an existing IAM compartment.

For information about how the IAM service and its components work and how to use them, see Overview of Oracle Cloud Infrastructure Identity and Access Management. For quick answers to common questions about IAM, see the Identity and Access Management FAQ.

Authorized Access on Autonomous Database

Only authorized users are allowed access to an Autonomous Database instance.

Oracle Cloud Operators do not have authorization to access your Autonomous Database. When access to your database is required to troubleshoot or mitigate an issue, you can allow a Cloud Operator to access a database for a limited time.

You allow a Cloud Operator to access the database by running the procedure `DBMS_CLOUD_ADMIN.ENABLE_OPERATOR_ACCESS`. This means if you file a service request with Oracle Cloud Support or by contacting your support representative and Oracle Cloud Operators need to access your database, you must also enable operator access by running `DBMS_CLOUD_ADMIN.ENABLE_OPERATOR_ACCESS`.

Each database access by Oracle Cloud Operators is logged with a request ID and reason.

See Manage Oracle Cloud Operator Access and View Oracle Cloud Infrastructure Operations Actions for more information.

Autonomous Database Fully Managed Service

Autonomous Database is a fully managed service and Oracle uses its own Oracle Cloud Infrastructure tenancies to run the Autonomous Database service.

Oracle Cloud Operators do not have access to customer tenancies and cloud operators cannot access the network.

Auditing Overview on Autonomous Database

Oracle Autonomous Database provides robust auditing capabilities that enable you to track who did what on the service and on specific databases. Comprehensive log data allows you to audit and monitor actions on your resources, which helps you to meet your audit requirements while reducing security and operational risk.

- Auditing Service Level Activities
  All actions Oracle Cloud users perform on the resources that make up your deployment of Oracle Autonomous Database are logged by the Audit service, regardless of the interface used: the Oracle Cloud Infrastructure Console, REST API, Command Line Interface (CLI), Software Development Kits (SDK) and so on.
- Auditing Database Activities
  Oracle Autonomous Database configures the autonomous databases you create to use the unified auditing feature of Oracle Database.

Auditing Service Level Activities

All actions Oracle Cloud users perform on the resources that make up your deployment of Oracle Autonomous Database are logged by the Audit service, regardless of the interface used: the Oracle Cloud Infrastructure Console, REST API, Command Line Interface (CLI), Software Development Kits (SDK) and so on.

You can use the Audit service to perform diagnostics, track resource usage, monitor compliance, and collect security-related events. For more information about the Audit service, see Overview of Audit in *Oracle Cloud Infrastructure Documentation*.

Additionally, when users perform operations on your Oracle Autonomous Database, the database publishes *events* to the Oracle Cloud Events service. The Oracle Cloud Events service allows you to create rules to capture these events and perform actions.

For more information about how the Events service works and how to set up the rules and actions it uses, see Overview of Events. For listings of the Oracle Autonomous Database operations that generate events, see Autonomous Database Event Types.

Auditing Database Activities

Oracle Autonomous Database configures the autonomous databases you create to use the unified auditing feature of Oracle Database.

This feature captures audit records from the following sources and gathers them in a single audit trail in a uniform format:

- Audit records (including `SYS` audit records) from unified audit policies and `AUDIT` settings
- Fine-grained audit records from the `DBMS_FGA` PL/SQL package
- Oracle Database Real Application Security audit records
- Oracle Recovery Manager audit records
- Oracle Database Vault audit records
- Oracle Label Security audit records
- Oracle Data Mining records
- Oracle Data Pump
- Oracle SQL*Loader Direct Load

Audit information is retained for up to 14 days, after which it is automatically purged. To retain audit information for longer, and to easily analyze and report on database activity, use Oracle Data Safe (included with your Oracle Autonomous Database subscription).

See About Auditing Autonomous Database for more information.

Assessing the Security of Your Database and its Data

Oracle Autonomous Database integrates with Oracle Data Safe to help you assess and secure your databases.

Oracle Data Safe helps you understand the sensitivity of your data, evaluate risks to data, mask sensitive data, implement and monitor security controls, assess user security, monitor user activity, and address data security compliance requirements in your databases.

You use Oracle Data Safe to identify and protect sensitive and regulated data your Oracle Autonomous Database by registering your database with Data Safe. Then, you use the Data Safe console directly from the Details page of your database.

For more information about using Data Safe, see Use Oracle Data Safe Features.

Regulatory Compliance Certification

Oracle Autonomous Database meets a broad set of international and industry-specific compliance standards.

Certification	Description
C5	The Cloud Computing Compliance Controls Catalog (C5)
CSA STAR	The Cloud Security Alliance (CSA) Security Trust, Assurance and Risk (STAR)
Cyber Essentials (UK) Cyber Essentials Plus (UK)	Oracle Cloud Infrastructure has achieved Cyber Essentials and Cyber Essentials Plus certification in these regions: • UK Gov South (London) • UK Gov West (Newport)
DESC (UAE)	Dubai Electronic Security Center CSP Security Standard
DoD IL4/IL5	DISA Impact Level 5 authorization in these regions: • US DoD East (Ashburn) • US DoD North (Chicago) • US DoD West (Phoenix)
ENS High (Spain)	The Esquema Nacional de Seguridad with the level of accreditation High.
FedRAMP High	Federal Risk and Authorization Management Program (U.S. Government Regions only)
FSI (S. Korea)	The Financial Security Institute
HDS	The French Public Health Code requires healthcare organizations that control, process, or store health or medical data to use infrastructure, hosting, and platform service providers that are Hébergeur de Données de Santé (HDS) accredited and certified
HIPAA	Health Insurance Portability and Accountability Act
HITRUST	The Health Information Trust Alliance
IRAP (Australia)	The Infosec Registered Assessors Program. Sydney and Melbourne regions
ISMS (S. Korea)	The Information Security Management System
ISO/IEC 27001:2013	International Organization for Standardization 27001
ISO/IEC 27017:2015	Code of Practice for Information Security Controls Based on ISO/IEC 27002 for Cloud Services
ISO/IEC 27018:2014	Code of Practice for Protection of Personally Identifiable Information (PII) In Public Clouds Acting as PII Processors
MeitY (India)	The Ministry of Electronics and Information Technology
MTCS (Singapore)	Multi-Tier Cloud Service (MTCS) Level 3
PASF (UK OC4)	Police Assured Secure Facilities (PASF) in these regions: • UK Gov South (London) • UK Gov West (Newport)

Certification	Description
PCI DSS	Payment Card Industry Data Security Standard is a set of requirements intended to ensure that all companies that process, store, or transmit credit card information maintain a secure environment
SOC 1	System and Organization Controls 1
SOC 2	System and Organization Controls 2

For more information and a complete list of certifications, see Oracle Cloud Compliance.

Security Testing Policies on Autonomous Database

Autonomous Database is subject to the Oracle Cloud Security Testing policy which describes when and how you may conduct certain types of security testing of Oracle Cloud Infrastructure services, including vulnerability and penetration tests, as well as tests involving data scraping tools.

Oracle regularly performs penetration and vulnerability testing and security assessments against Oracle Cloud Infrastructure, platforms, and applications. These tests are intended to validate and improve the overall security of Oracle Cloud services. You can also conduct your own testing as per the rules explained in the policy.

See Oracle Cloud Security Testing Policies for more information.

Manage Encryption Keys on Autonomous Database

Describes how to use customer-managed encryption keys with Autonomous Database, and if you are using customer-managed encryption keys, shows how to rotate the keys, switch to Oracle-managed encryption keys, or view the encryption key history.

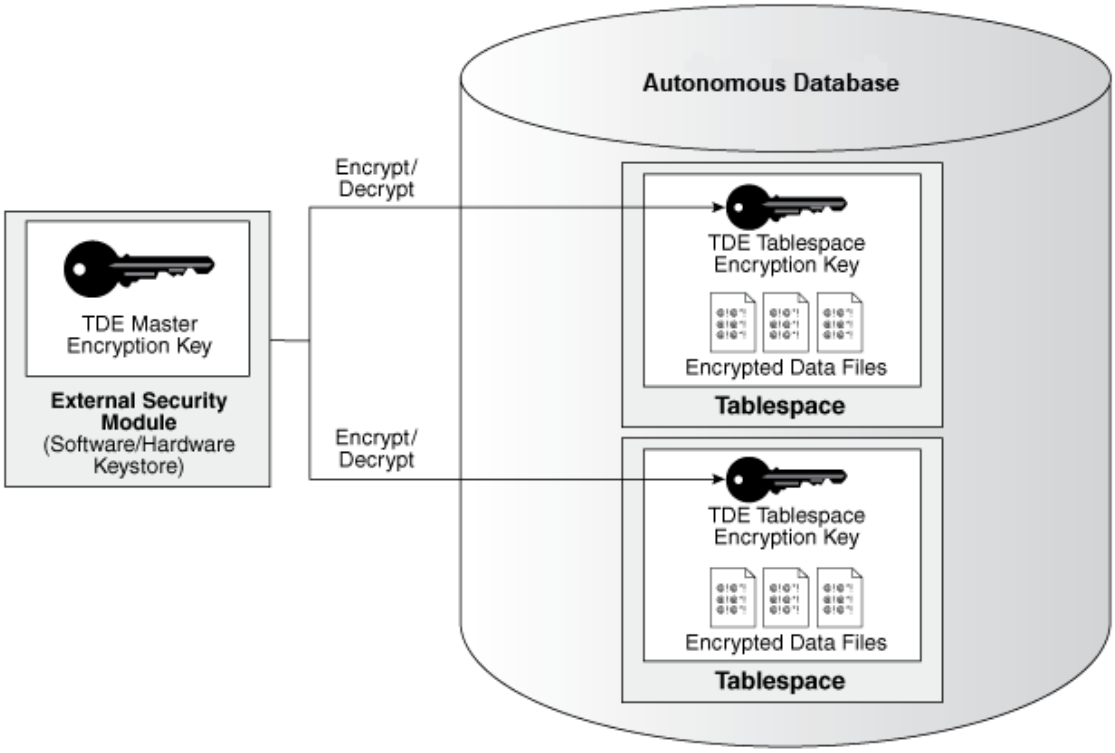
- About Master Encryption Key Management on Autonomous Database
  Autonomous Database provides two options for Transparent Data Encryption (TDE) to encrypt your database: Oracle-managed encryption keys and Customer-managed encryption keys.
- Manage Master Encryption Keys in OCI Vault
  If your organization's security policies require customer-managed encryption keys, you can configure Autonomous Database to use an Oracle Cloud Infrastructure Vault master encryption key. With customer-managed master encryption keys, Autonomous Database uses the master encryption key to generate the TDE master key.
- Manage Master Encryption Keys in Azure Key Vault
  Autonomous Database supports customer-managed Transparent Data Encryption (TDE) keys that reside in Azure Key Vault.
- Manage Master Encryption Keys in AWS Key Management Service
  Autonomous Database supports customer-managed Transparent Data Encryption (TDE) keys that reside in AWS Key Management Service (KMS).
- Manage Master Encryption Keys in Oracle Key Vault
  Autonomous Database supports customer-managed Transparent Data Encryption (TDE) keys that reside in Oracle Key Vault (OKV).
- Switch to Oracle-Managed Encryption Keys on Autonomous Database
  Shows the steps to switch to Oracle-managed master encryption keys on Autonomous Database if you are using customer-managed encryption keys.

- View History for Customer-Managed Encryption Keys on Autonomous Database
  You can view the key history from the Oracle Cloud Infrastructure Console or by selecting from the `V$ENCRYPTION_KEYS` view.
- Notes for Using Customer-Managed Keys with Autonomous Database
  Provides additional information and notes for using customer-managed keys with Autonomous Database

About Master Encryption Key Management on Autonomous Database

Autonomous Database provides two options for Transparent Data Encryption (TDE) to encrypt your database: Oracle-managed encryption keys and Customer-managed encryption keys.

Autonomous Database uses Transparent Data Encryption, including a TDE master key and TDE tablespace keys to encrypt data in the database. As shown in the following figure, the TDE master key generates and encrypts/decrypts the TDE tablespace keys, and the TDE tablespace keys encrypt the data files.



Oracle-Managed Master Encryption Keys on Autonomous Database

By default, Autonomous Database uses Oracle-managed encryption keys.

Using Oracle-managed keys, Autonomous Database creates and manages the encryption keys that protect your data and Oracle handles rotation of the TDE master key.

Customer-Managed Master Encryption Keys on Autonomous Database

With customer-managed master encryption keys, Autonomous Database uses the master encryption key in a customer-managed key vault to generate the TDE master key. If your organization's security policies require customer-managed encryption keys, you can configure Autonomous Database to use a master encryption key in the following key management systems:

- Oracle Cloud Infrastructure (OCI) Vault

See Manage Master Encryption Keys in OCI Vault for more information.

- Microsoft Azure Key Vault
  See Manage Master Encryption Keys in Azure Key Vault for more information.
- Amazon Web Services (AWS) Key Management Service (KMS)
  See Manage Master Encryption Keys in AWS Key Management Service for more information.
- Oracle Key Vault (OKV)
  See Manage Master Encryption Keys in Oracle Key Vault for more information.

Manage Master Encryption Keys in OCI Vault

If your organization's security policies require customer-managed encryption keys, you can configure Autonomous Database to use an Oracle Cloud Infrastructure Vault master encryption key. With customer-managed master encryption keys, Autonomous Database uses the master encryption key to generate the TDE master key.

- About Customer-Managed Keys on Autonomous Database in OCI Vault
  Using customer-managed encryption keys on Autonomous Database in Oracle Cloud Infrastructure (OCI) Vault involves creating a master key in your OCI Vault and configuring your Autonomous Database instance to use encryption keys in the OCI Vault.
- Prerequisites to Use Customer-Managed Encryption Keys on Autonomous Database in OCI Vault
  Perform these prerequisite steps to use customer-managed keys on Autonomous Database in OCI Vault:
- Use Customer-Managed Encryption Keys with Vault Located in Local Tenancy
  Shows the steps to select customer-managed master encryption keys on Autonomous Database. If you are using customer-managed master encryption keys, follow these steps to rotate the master keys.
- Use Customer-Managed Encryption Key Located in a Remote Tenancy
  Shows the steps to select customer-managed master encryption keys from a Vault on a remote tenancy.
- Rotate Customer-Managed Encryption Keys on Autonomous Database in OCI Vault
  Describes how to rotate customer-managed encryption keys on Autonomous Database in OCI Vault.

About Customer-Managed Keys on Autonomous Database in OCI Vault

Using customer-managed encryption keys on Autonomous Database in Oracle Cloud Infrastructure (OCI) Vault involves creating a master key in your OCI Vault and configuring your Autonomous Database instance to use encryption keys in the OCI Vault.

Follow these general steps:

1. Create a master encryption key in your OCI Vault.

   See Prerequisites to Use Customer-Managed Encryption Keys on Autonomous Database in OCI Vault for more information.

2. Select customer-managed encryption keys from the Oracle Cloud Infrastructure Console:
   - For an existing database, select **Manage Encryption Key** on the Oracle Cloud Infrastructure Console.

- While provisioning, under **Advanced Options**, on the **Encryption Key** tab select either **Encrypt using customer-managed key in this tenancy** or **Encrypt using a customer-managed key located in a remote tenancy**.
- While cloning, under **Advanced Options**, on the **Encryption Key** tab select either **Encrypt using customer-managed key in this tenancy** or **Encrypt using a customer-managed key located in a remote tenancy**.

See Use Customer-Managed Encryption Keys with Vault Located in Local Tenancy, Use Customer-Managed Encryption Key Located in a Remote Tenancy and Notes for Using Customer-Managed Keys with Autonomous Database for more information.

Prerequisites to Use Customer-Managed Encryption Keys on Autonomous Database in OCI Vault

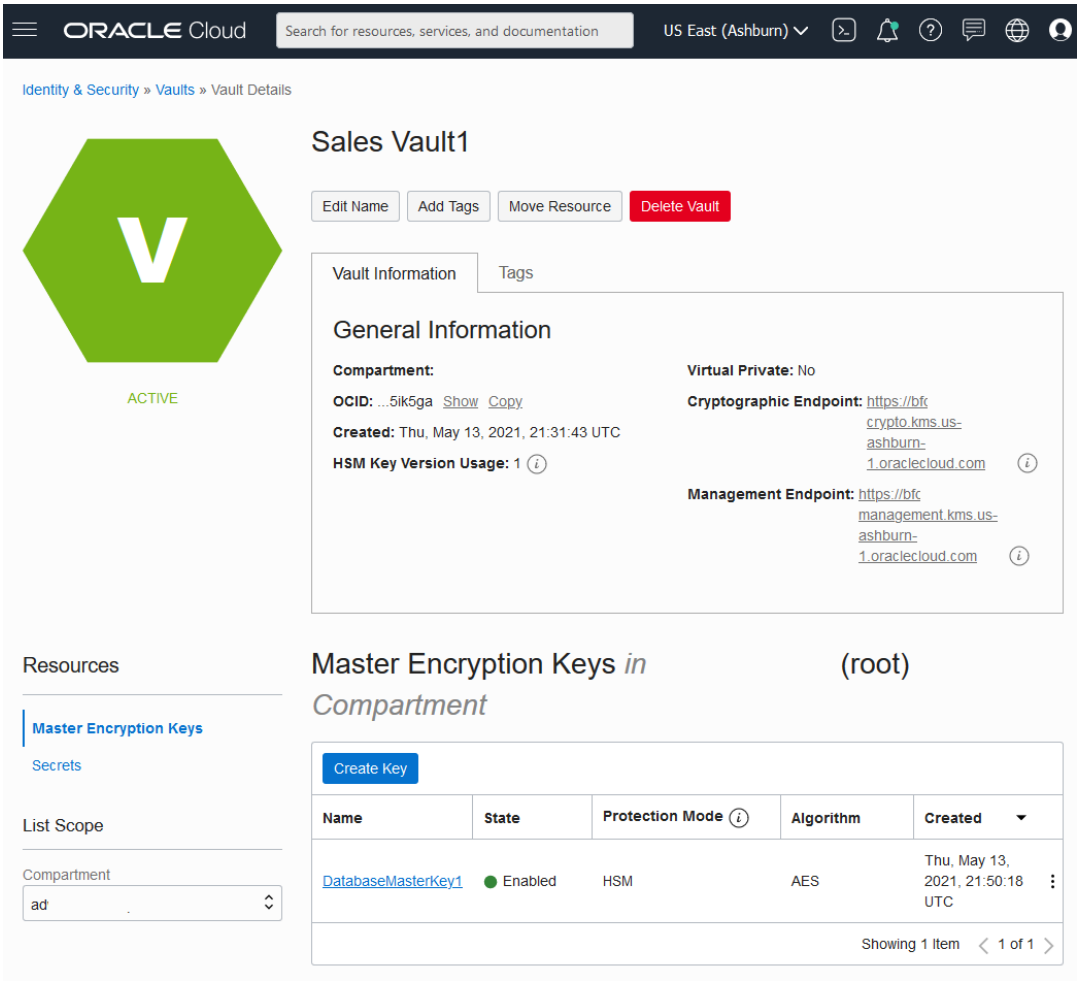
Perform these prerequisite steps to use customer-managed keys on Autonomous Database in OCI Vault:

1. Create an Oracle Cloud Infrastructure Vault.
   a. Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
   b. From the Oracle Cloud Infrastructure left navigation menu click **Identity and Security**.
   c. Under **Key Management & Secret Management** click **Vault**.
   d. Select an existing Vault or create a new Vault.

      For more details, see Creating a Vault.
2. Create a Master Encryption Key in the Vault.

> **Note:**
>
> You must use these options when you create the key:
>
> - **Key Shape: Algorithm**: AES (Symmetric key used for Encrypt and Decrypt)
> - **Key Shape: Length**: 256 bits

For more information, see Creating a Master Encryption Key and Overview of Key Management.



3. Create dynamic group and policy statements for the dynamic group to enable access to Oracle Cloud Infrastructure resources (Vaults and Keys).

   This step depends on whether the vault is on the same tenancy as the Autonomous Database instance or on a different tenancy:

   - Vault and keys are in the same tenancy as the Autonomous Database instance. See Create Dynamic Group and Policies for Customer Managed Keys with Vault in Same Tenancy as Database for more information.
   - Vault and keys are in a different tenancy. See Create Dynamic Group and Policies for Customer Managed Keys with Vault in Different Tenancy than the Database for more information.

You must replicate the vault and keys to use customer-managed encryption keys with Autonomous Data Guard with a remote Standby database. Customer-Managed Encryption Keys are only supported with a single cross-region Autonomous Data Guard standby. Multiple cross-region standbys are not supported because Oracle Cloud Infrastructure Vault only supports replication to one remote region.

See the following for more information:

- To Replicate a Vault and Its Keys
- Autonomous Data Guard with Customer Managed Keys

- Create Dynamic Group and Policies for Customer Managed Keys with Vault in Same Tenancy as Database
  Create dynamic group and policies to provide access to the vault and keys for customer-managed keys when the vault and keys are in the same tenancy as the Autonomous Database instance.
- Create Dynamic Group and Policies for Customer Managed Keys with Vault in Different Tenancy than the Database
  Perform these steps to use customer-managed keys when the Autonomous Database instance and vaults and keys are in different tenancies.

Create Dynamic Group and Policies for Customer Managed Keys with Vault in Same Tenancy as Database

Create dynamic group and policies to provide access to the vault and keys for customer-managed keys when the vault and keys are in the same tenancy as the Autonomous Database instance.

1. Create a dynamic group to make the master encryption key accessible to the Autonomous Database instance.
   a. In the Oracle Cloud Infrastructure console click **Identity & Security**.
   b. Under **Identity** click **Domains** and select an identity domain (or create a new identity domain).
   c. Under **Identity domain**, click **Dynamic groups**.
   d. Click **Create dynamic group** and enter a **Name**, a **Description**, and a rule.
      - **Create Dynamic Group for an existing database:**

        You can specify that an Autonomous Database instance is part of the dynamic group. The dynamic group in the following example includes only the Autonomous Database whose OCID is specified in the `resource.id` parameter:

        ```
        resource.id = '<your_Autonomous_Database_instance_OCID>'
        ```

      - **Create a Dynamic Group for a database that has not been provisioned yet:**

        When you are creating the dynamic group before you provision or clone an Autonomous Database instance, the OCID for the new database is not yet available. For this case, create a dynamic group that specifies the resources in a given compartment:

        ```
        resource.compartment.id = '<your_Compartment_OCID>'
        ```

   e. Click **Create**.
2. Write policy statements for the dynamic group to enable access to Oracle Cloud Infrastructure resources (vaults and keys).
   a. In the Oracle Cloud Infrastructure console click **Identity & Security** and click **Policies**.
   b. To write policies for a dynamic group, click **Create Policy**, and enter a **Name** and a **Description**.
   c. Use the **Policy Builder** to create a policy for vault and keys in the local tenancy.

For example, the following allows the members of the dynamic group `DGKeyCustomer1` to access the vaults and keys in the compartment named `training`:

```
Allow dynamic-group DGKeyCustomer1 to use vaults in compartment training
Allow dynamic-group DGKeyCustomer1 to use keys in compartment training
```

This sample policy applies for a single compartment. You can specify that a policy applies for your tenancy, a compartment, a resource, or a group of resources.

To use customer-managed keys with Autonomous Data Guard with a remote standby, the following policy is also required:

```
Allow dynamic-group DGKeyCustomer1 to manage vaults in compartment
training
Allow dynamic-group DGKeyCustomer1 to manage keys in compartment
training
```

d. Click **Create** to save the policy.

See the following for more information:

- Managing Policies
- Autonomous Data Guard with Customer Managed Keys
- Replicating Vaults and Keys

Create Dynamic Group and Policies for Customer Managed Keys with Vault in Different Tenancy than the Database

Perform these steps to use customer-managed keys when the Autonomous Database instance and vaults and keys are in different tenancies.

In this case, you need to supply OCID values when you change to customer-managed keys. In addition, you need to define dynamic groups and policies that allow the Autonomous Database instance to use vaults and keys in a different tenancy.

1. Copy the master encryption key OCID.
2. Copy the vault OCID.
3. Copy the tenancy OCID (the remote tenancy that contains vaults and keys).
4. On the tenancy with the Autonomous Database instance, create a dynamic group.
   a. In the Oracle Cloud Infrastructure console, on the tenancy with the Autonomous Database instance, click **Identity & Security**.
   b. Under **Identity** click **Domains** and select an identity domain (or create a new identity domain).
   c. Under **Identity domain**, click **Dynamic groups**.
   d. Click **Create dynamic group** and enter a **Name**, a **Description**, and a rule.
      - **Create Dynamic Group for an existing database:**

        You can specify that an Autonomous Database instance is part of the dynamic group. The dynamic group in the following example includes only the Autonomous Database whose OCID is specified in the `resource.id` parameter:

        ```
        resource.id = '<your_Autonomous_Database_instance_OCID>'
        ```

- **Create a Dynamic Group for a database that has not been provisioned yet:**

  When you are creating the dynamic group before you provision or clone an Autonomous Database instance, the OCID for the new database is not yet available. For this case, create a dynamic group that specifies the resources in a given compartment:

  ```
  resource.compartment.id = '<your_Compartment_OCID>'
  ```

e. Click **Create**.

5. On the tenancy with the Autonomous Database instance, define the policies to allow access to vaults and keys (where the vaults and keys are on a different tenancy).

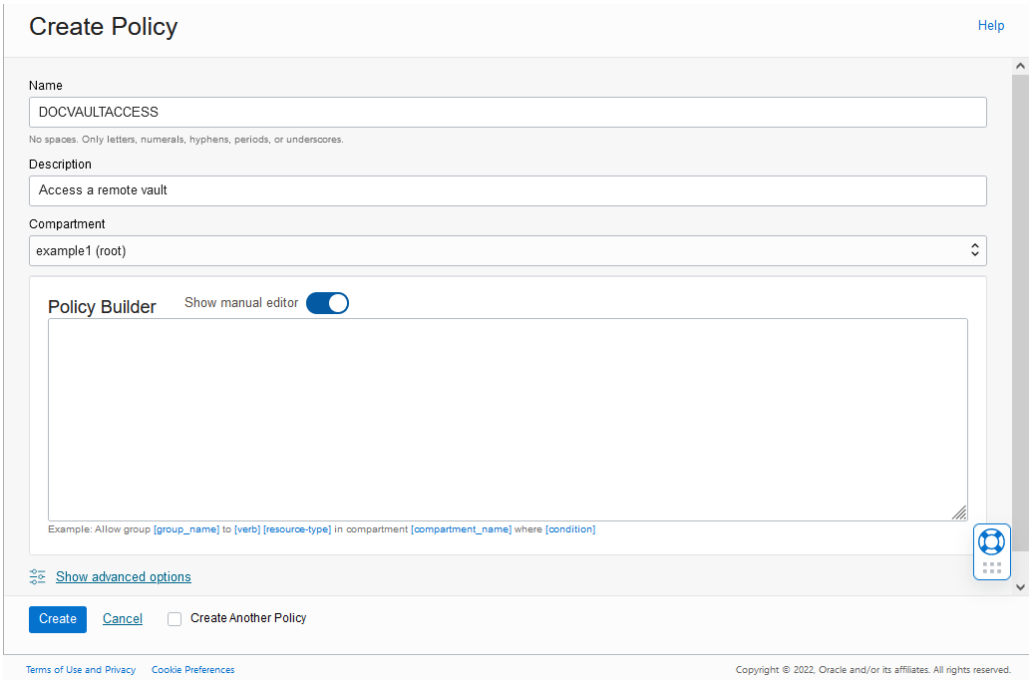
   a. In the Oracle Cloud Infrastructure console click **Identity & Security**.

   b. Under **Identity** click **Policies**.

   c. To write a policy, click **Create Policy**.

   d. On the Create Policy page, enter a Name and a Description.

   e. On the Create Policy page, select **Show manual editor**.



   f. In the policy builder, add policies so that the Autonomous Database instance is able to access vaults and keys located in the different tenancy. Also add policies for the IAM group that the IAM user belongs to so that the Oracle Cloud Infrastructure Console for the Autonomous Database instance can show details about the key that resides in a different tenancy.

   For example, in the generic policy, call the tenancy with the Autonomous Database instance Tenancy-1 and the tenancy with vaults and keys, Tenancy-2:

   Copy the following policy and replace the variables and names with the values you define, where the dynamic group name `ADB-DynamicGroup` is the dynamic group you created in Step 4:

   ```
   define tenancy REMTEN as <ocid of tenancy-2>
   endorse dynamic-group ADB-DynamicGroup to use vaults in tenancy REMTEN
   ```

```
endorse dynamic-group ADB-DynamicGroup to use keys in tenancy REMTEN
endorse group MyUserGroup to use vaults in tenancy REMTEN
endorse group MyUserGroup to use keys in tenancy REMTEN
```

For example, the following allows the members of the dynamic group `DGKeyCustomer1` to access the remote vaults and keys in the tenancy named `training2`:

```
define tenancy training2 as ocid1.tenancy.oc1..aaa_example_rcyx2a
endorse dynamic-group DGKeyCustomer1 to use vaults in tenancy training2
endorse dynamic-group DGKeyCustomer1 to use keys in tenancy training2
endorse group MyUserGroup to use vaults in tenancy training2
endorse group MyUserGroup to use keys in tenancy training2
```

g. Click **Create** to save the policy.

6. Copy the tenancy OCID (the tenancy that contains the Autonomous Database instance).
7. Copy the Dynamic Group OCID (for the Dynamic Group you created in Step 4).
8. On the remote tenancy with vaults and keys, define a dynamic group and policies to allow the Autonomous Database instance to access vaults and keys.
   a. In the Oracle Cloud Infrastructure console, click **Identity & Security**.
   b. Under **Identity** click **Policies**.
   c. To create a policy, click **Create Policy**.
   d. On the Create Policy page, enter a Name and a Description.
   e. On the Create Policy page, select **Show manual editor**.
   f. In the policy builder, add policies and a dynamic group to provide access to the dynamic group on the tenancy with the Autonomous Database instance (Tenancy-1), such that the Autonomous Database instance can use the vaults and keys in Tenancy-2. Also need to add policies to allow the user group to access the vault and keys to display information on the Oracle Cloud Infrastructure Console for the Autonomous Database instance in a different tenancy.

      Use the **Policy Builder** to create a dynamic group and a policy for vaults and keys.

```
define tenancy ADBTEN as <ocid of tenancy-1>
define dynamic-group REM-ADB-DG as <ocid of the Dynamic Group in
tenancy-1>
define group REMGROUP as <group-ocid>
admit dynamic-group REM-ADB-DG of tenancy ADBTEN to use vaults in
tenancy
admit dynamic-group REM-ADB-DG of tenancy ADBTEN to use keys in tenancy
admit group REMGROUP of tenancy ADBTEN to use vaults in tenancy
admit group REMGROUP of tenancy ADBTEN to use keys in tenancy
```

For example define the following on the remote tenancy to allow the members of the dynamic group `DGKeyCustomer1` and the group `REMGROUP` to access the remote vaults and keys in the tenancy named `training2`:

```
define tenancy adbdemo5 as ocid1.tenancy.oc1..aaa_example_4cnl5q
define dynamic-group REM-ADB-DG as
ocid1.dynamicgroup.oc1..aaa_example_526bia
define group REMGROUP as ocid1.group.oc1..aaa_example_6vctn6xsaq
admit dynamic-group REM-ADB-DG of tenancy adbdemo5 to use vaults in
```

```
tenancy
admit dynamic-group REM-ADB-DG of tenancy adbdemo5 to use keys in
tenancy
admit group REMGROUP of tenancy ADBTEN to use vaults in tenancy
admit group REMGROUP of tenancy ADBTEN to use keys in tenancy
```

g. Click **Create** to save the policy.

Use Customer-Managed Encryption Keys with Vault Located in Local Tenancy

Shows the steps to select customer-managed master encryption keys on Autonomous Database. If you are using customer-managed master encryption keys, follow these steps to rotate the master keys.

> **Caution:**
>
> The customer-managed encryption key is stored in Oracle Cloud Infrastructure Vault, external to the database host. If the customer-managed encryption key is disabled or deleted, the database will be inaccessible.

For details on using customer-managed keys where the Vault is located in a remote tenancy, see Use Customer-Managed Encryption Key Located in a Remote Tenancy.

On Autonomous Database you can choose customer-managed keys as follows:

- While provisioning, under **Advanced Options**, on the **Encryption Key** tab.
- While cloning, under **Advanced Options**, on the **Encryption Key** tab

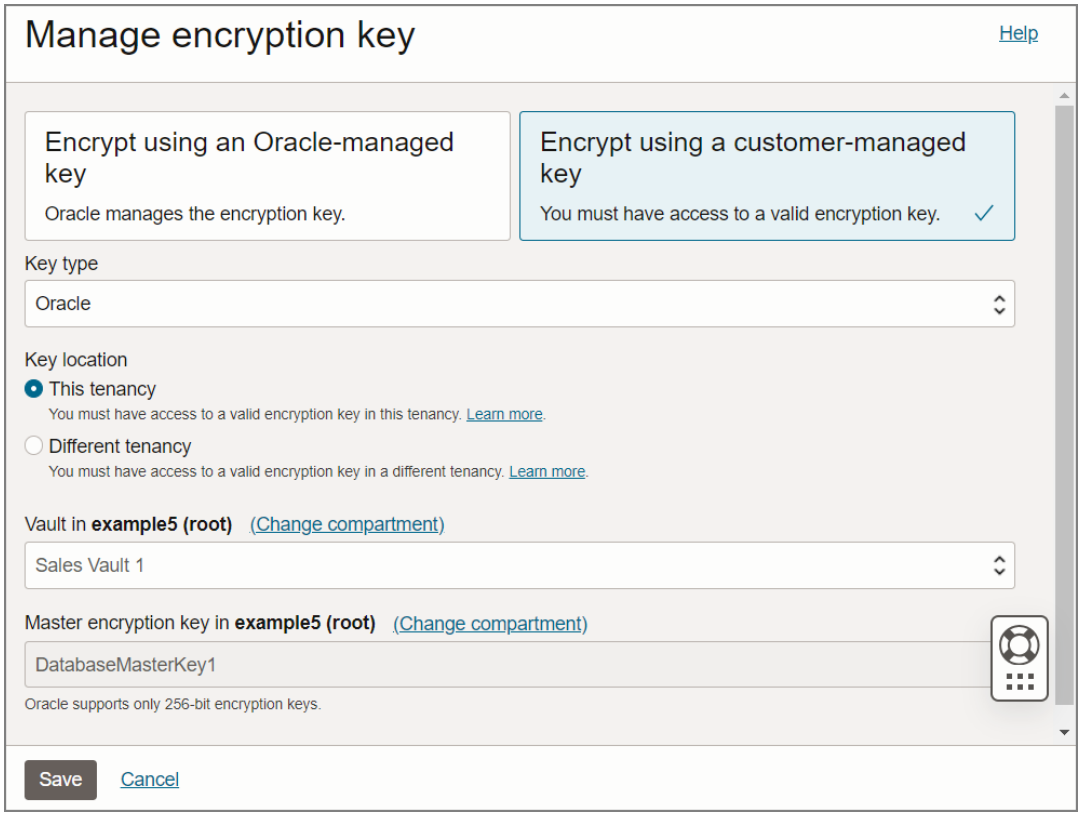
Follow these steps if your Autonomous Database is using Oracle-managed keys and you want to switch to customer-managed encryption keys with the vault in the local tenancy, or if you are using customer-managed encryption keys and you want to rotate the master key.

1. Perform the required customer-managed encryption key prerequisite steps as necessary. See Prerequisites to Use Customer-Managed Encryption Keys on Autonomous Database in OCI Vault for more information.
2. On the **Details** page, from the **More actions** drop-down list, select **Manage encryption key**.
3. On the **Manage encryption key** page, select **Encrypt using customer-managed key**.

   If you are already using customer-managed keys and you want to rotate the TDE keys, follow these steps and select a different key (select a key that is different from the currently selected master encryption key).
4. For **Key type**, select **Oracle**.
5. For **Key location**, click **This tenancy**.
6. Select a Vault.

   Click **Change compartment** to select a vault in a different compartment.
7. Select a Master encryption key.

   Click **Change Compartment** to select a master encryption key in a different compartment.



When cross-region Autonomous Data Guard is enabled the **Vault** and **Master encryption key** values show the keys that are replicated in both the primary and the remote standby region.

8. Click **Save**.

The **Lifecycle state** changes to **Updating**. When the request completes, the **Lifecycle state** shows **Available**.

After the request completes, on the Oracle Cloud Infrastructure Console, the key information shows on the Autonomous Database Information page under the heading **Encryption**. This area shows the **Encryption Key** field with a link to the Master Encryption Key and the **Encryption Key OCID** field with the Master Encryption Key OCID.

See Notes for Using Customer-Managed Keys with Autonomous Database for more information.

Use Customer-Managed Encryption Key Located in a Remote Tenancy

Shows the steps to select customer-managed master encryption keys from a Vault on a remote tenancy.

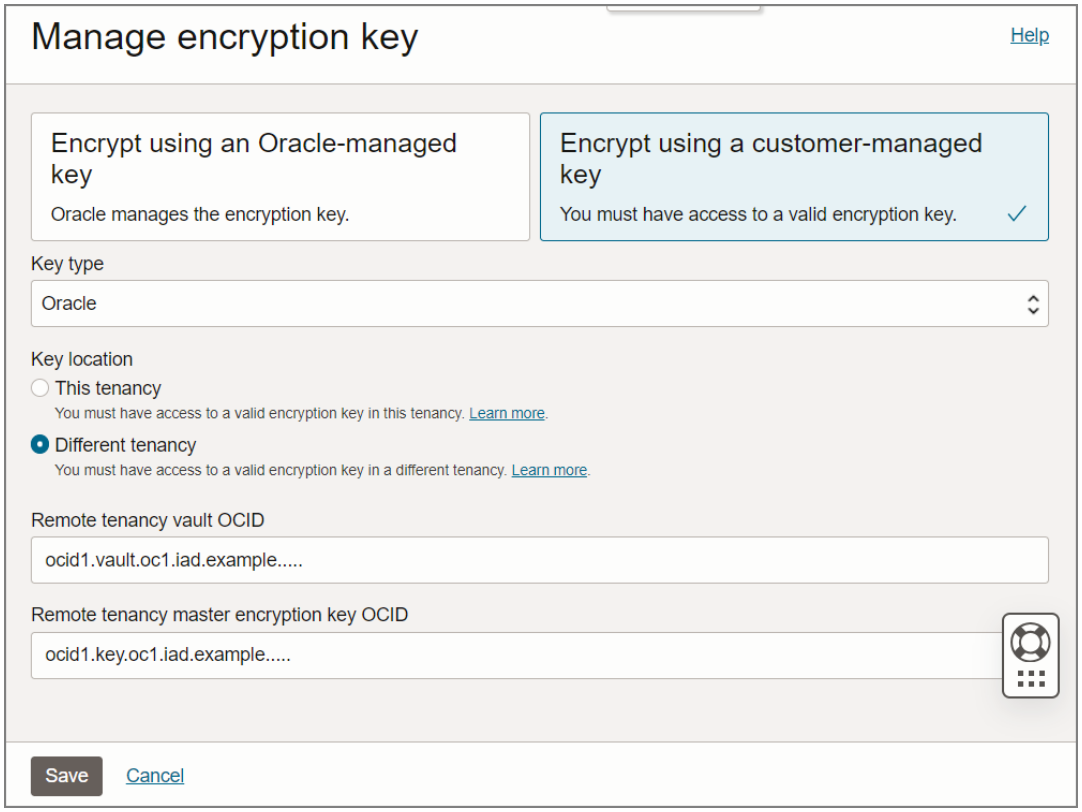
When you use customer-managed master encryption keys with a Vault in a remote tenancy, the Vault and the Autonomous Database instance must be in the same region. To change the tenancy, on the sign-on page click **Change tenancy**. After you change the tenancy, make sure to select the same region for both the Vault and the Autonomous Database instance.

1. Perform the required customer-managed encryption key prerequisite steps as necessary. See Prerequisites to Use Customer-Managed Encryption Keys on Autonomous Database in OCI Vault for more information.
2. On the **Details** page, from the **More actions** drop-down list, select **Manage encryption key**.

3. On the **Manage encryption key** page, select the **Encrypt using a customer-managed key** option.

   If you are already using customer-managed keys and you want to rotate the TDE keys, follow these steps and use a different key OCID with the same vault OCID, or use a new vault OCID and a new key OCID. This lets you use a key that is different from the current master encryption key.

4. For **Key type**, select **Oracle**.
5. For **Key location**, click **Different tenancy**.
6. Enter a remote tenancy vault OCID.
7. Enter a remote tenancy master encryption key OCID.



8. Click **Save**.

The **Lifecycle state** changes to **Updating**. When the request completes, the **Lifecycle state** shows **Available**.

After the request completes, on the Oracle Cloud Infrastructure Console, the key information shows on the Autonomous Database Information page under the heading **Encryption**. This area shows the **Encryption Key** field with a link to the Master Encryption Key and the **Encryption Key OCID** field with the Master Encryption Key OCID.

Rotate Customer-Managed Encryption Keys on Autonomous Database in OCI Vault

Describes how to rotate customer-managed encryption keys on Autonomous Database in OCI Vault.

When you rotate the customer-managed master encryption key, Autonomous Database generates a new TDE master key and uses the new TDE master key to re-encrypt the tablespace encryption keys that encrypt and decrypt your data. This operation is fast and does not require database downtime. It does not change the tablespace keys and does not re-encrypt customer data.

> **Note:**
>
> Using the Oracle Cloud Infrastructure Console you can rotate an Oracle Cloud Infrastructure Vault master encryption key with the **Rotate Key** command. This is a separate action and does not result in a new master encryption key for your Autonomous Database. To rotate the master encryption key of your Autonomous Database, create a new master encryption key in Oracle Cloud Infrastructure Vault and follow the steps described below.

To rotate customer-managed encryption keys:

1. Create a new master encryption key in your Oracle Cloud Infrastructure Vault. If you already have multiple master encryption keys, then select a master encryption key that is different than the key you are using as your master encryption key for your Autonomous Database instance.

   See Prerequisites to Use Customer-Managed Encryption Keys on Autonomous Database in OCI Vault for more information.

2. Rotate the master encryption key from the Oracle Cloud Infrastructure Console:

   See Use Customer-Managed Encryption Keys with Vault Located in Local Tenancy for more information.

Manage Master Encryption Keys in Azure Key Vault

Autonomous Database supports customer-managed Transparent Data Encryption (TDE) keys that reside in Azure Key Vault.

- Prerequisites to Use Customer-Managed Encryption Keys in Azure Key Vault
  Describes prerequisite steps to use customer-managed master encryption keys on Autonomous Database that reside in Azure Key Vault.
- Use Customer-Managed Encryption Keys on Autonomous Database with Azure Key Vault
  Shows the steps to encrypt your Autonomous Database using customer-managed master encryption keys that reside in Azure Key Vault.

Prerequisites to Use Customer-Managed Encryption Keys in Azure Key Vault

Describes prerequisite steps to use customer-managed master encryption keys on Autonomous Database that reside in Azure Key Vault.

Limitations:

- Azure Key Vault is only supported in commercial regions.
- Cross-tenancy access, where the Autonomous Database instance and the Azure Key Vault are in different tenancies, is not supported.
- Azure Key Vault is not supported in cross-region standbys.
- Azure Key Vault is not supported in refreshable clones.

- Supported key shapes and sizes are the following:
  - RSA 2048, 3072, and 4096
  - EC-P256, EC-P256K, EC-P384, EC-P521

Follow these steps:

1. Create an Autonomous Database instance that uses the default Encryption key setting of **Encrypt using an Oracle-managed key**. See Provision an Autonomous Database Instance for more information.

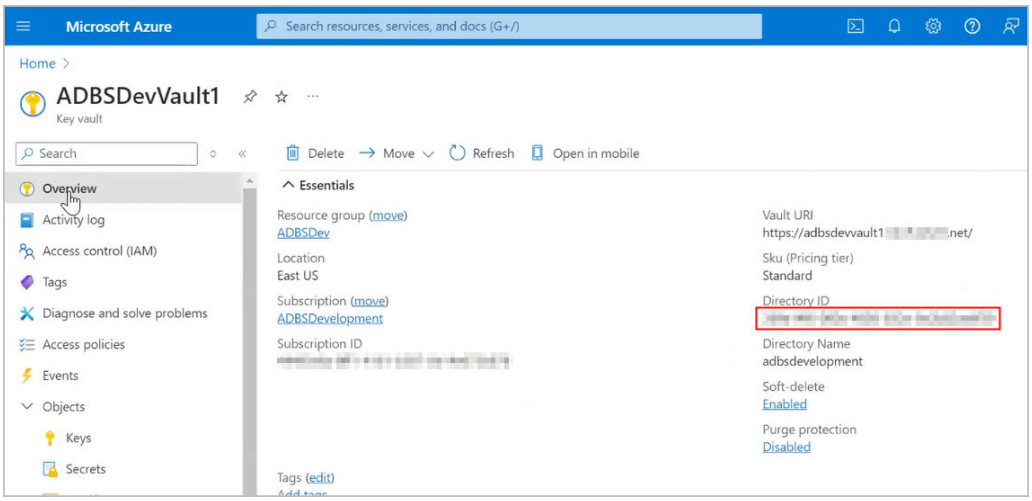
> **Note:**
>
> Encryption key settings for customer-managed keys in Azure Key Vault are not available during the creation process. The options are available post provisioning, when editing the instance.

2. Create an Azure Key Vault with a Transparent Data Encryption (TDE) master key.

   See About Azure Key Vault for more information.

3. Enable Azure service principal authentication with the Azure Directory `tenant_id` to allow your Autonomous Database instance access to the Azure Key Vault.

   a. Obtain your Microsoft Azure Active Directory tenant ID.



   See How to find your Azure Active Directory tenant ID for more information.

   b. Connect to your instance as ADMIN.

   c. Enable Azure service principal with `DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH`. The value of `azure_tenantid` is the Azure Directory ID you obtained in the previous step.

   For example:

```
BEGIN
 DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(
        provider => 'AZURE',
        params   => JSON_OBJECT('azure_tenantid' value
'azure_directoryID'));
```

```
END;
/
```

This enables Azure service principal authentication and creates an Azure application for Autonomous Database on the Azure portal. See Enable Azure Service Principal for more information.

4. Provide the Azure application consent to access Azure resources from Autonomous Database.

   a. On Autonomous Database query `CLOUD_INTEGRATIONS`.

      For example:

```
SELECT * FROM CLOUD_INTEGRATIONS;

PARAM_NAME        PARAM_VALUE
---------------
------------------------------------------------------------------------
------------------------------------------------------------------
azure_tenantid    29...eb
azure_consent_url https://login.microsoftonline.com/f8...5a/oauth2/v2.0/
authorize?client_id=d4f5...d5&response_type=code&scope=User.read
azure_app_name    ADBS_APP_OCID1.AUTONOMOUSDATABASE...
```

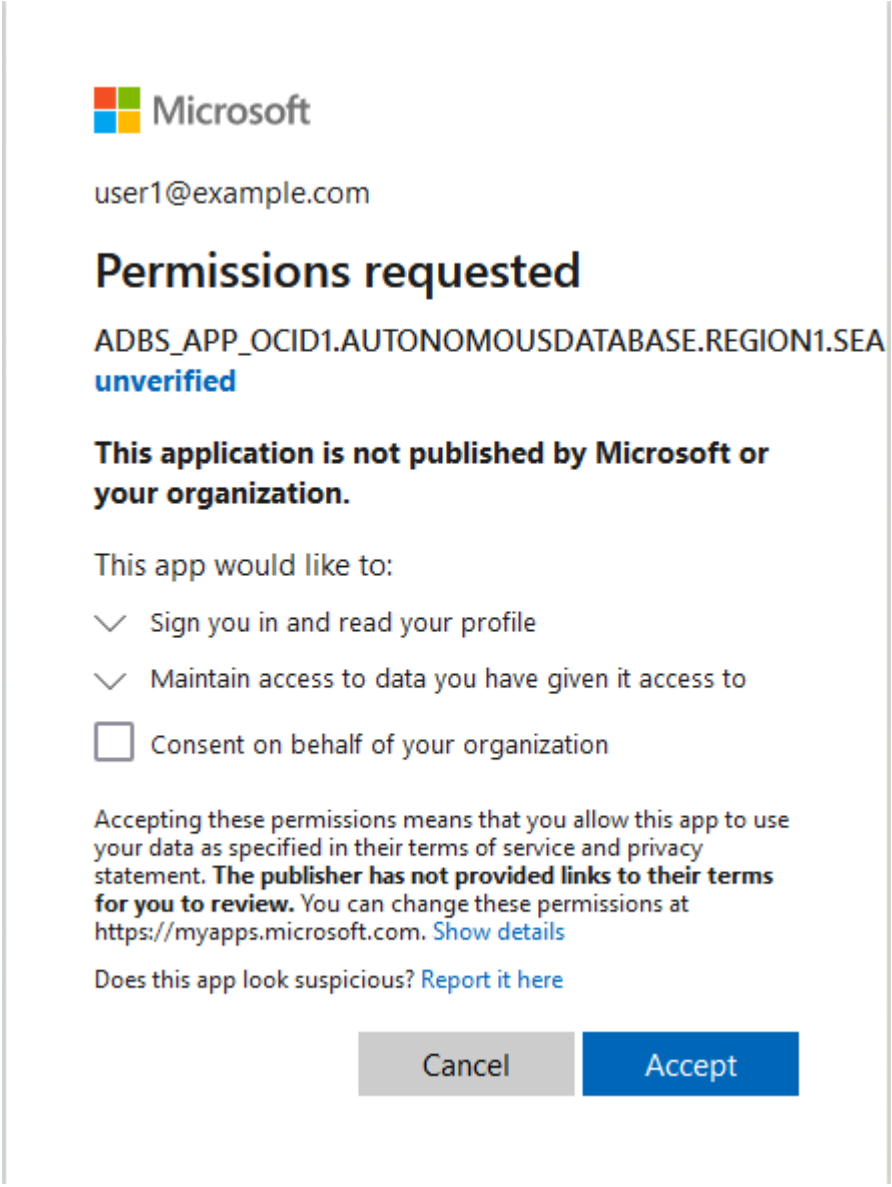
      The view `CLOUD_INTEGRATIONS` is available to the `ADMIN` user or to a user with `DWROLE` role.

   b. In a browser, open the Azure consent URL specified by the `azure_consent_url` parameter.

      For example, copy the `azure_consent_url` from the query results and enter the URL into your browser:

```
https://login.microsoftonline.com/f8...5a/oauth2/v2.0/authorize?
client_id=d4f5...d5&response_type=code&scope=User.read
```

      The **Permissions requested** page opens and shows a consent request, similar to the following:



5. Obtain the Azure application name.
   a. On Autonomous Database query `CLOUD_INTEGRATIONS`.

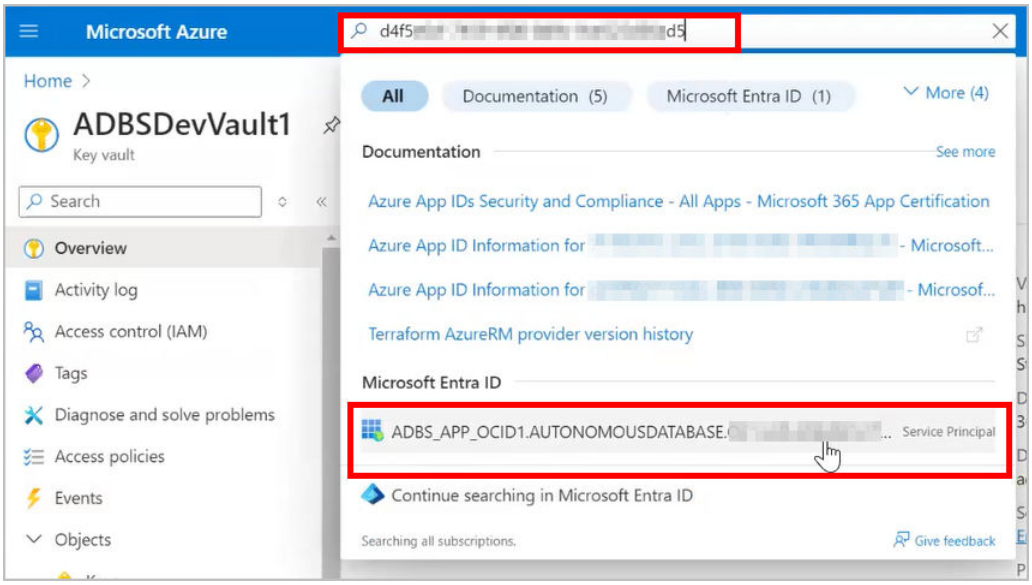
      For example:

      ```
      SELECT * FROM CLOUD_INTEGRATIONS;
      ```

   b. Copy the `client_id` value included in the `consent_url`.

      ```
      PARAM_NAME        PARAM_VALUE
      --------------
      ------------------------------------------------------------------------
      ------------------------------------------------------------------------
      azure_tenantid    29...eb
      azure_consent_url https://login.microsoftonline.com/f8...5a/oauth2/v2.0/
      authorize?client_id=d4f5...d5&response_type=code&scope=User.read
      azure_app_name    ADBS_APP_OCID1.AUTONOMOUSDATABASE...
      ```

c. Search on the `client_id` in the Azure portal. The application ID is displayed under Microsoft Entra ID.



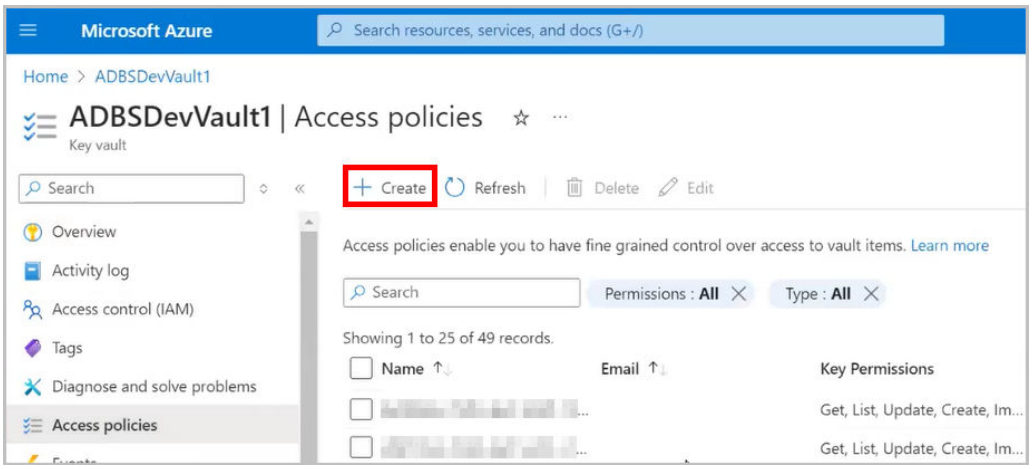
d. Copy the application ID. This value is used in subsequent steps to allow this application access to keys in the Azure Key Vault.

6. Assign necessary roles for the Azure application to access the Azure Key Vault.

a. In the Azure portal, navigate to the Access policies for the Azure Key Vault.

The list of applications with access to this vault are displayed.

b. On the Access policies page, click **+ Create** to create an access policy for the application to access this key vault.



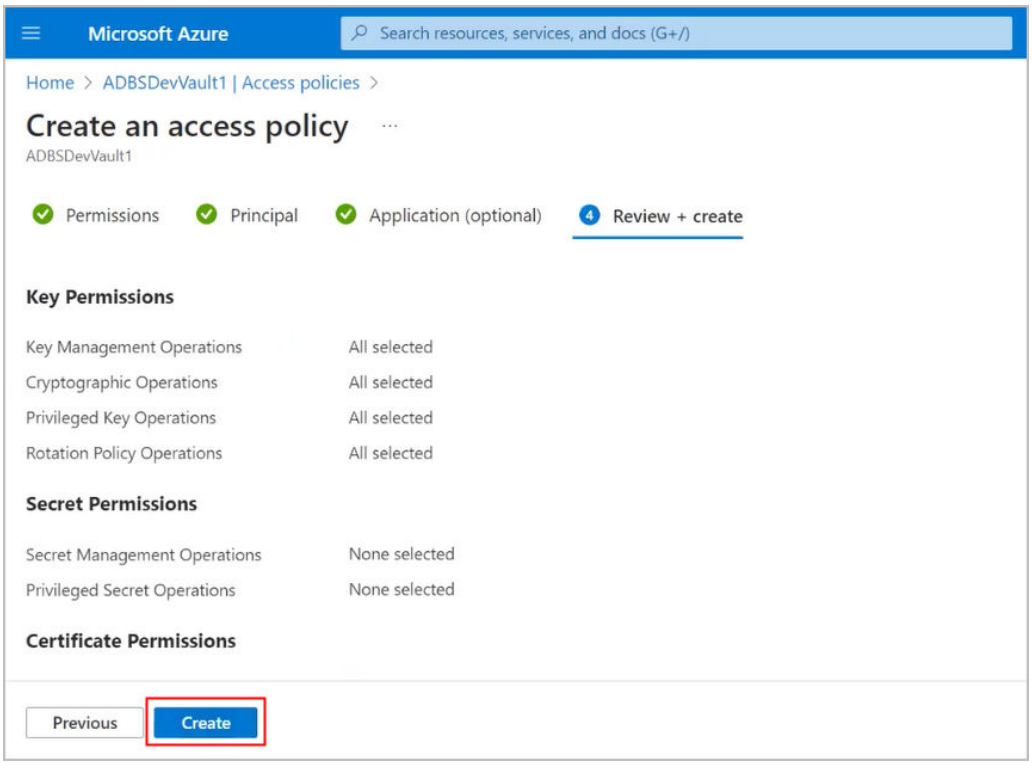
c. For Permissions on the Create an access policy page, select all Key permissions, including: Key Management Operations, Cryptography Operations, Privileged Key Operations, and Rotation Policy Operations and click **Next**.

d. For Principal, search the application name.

e. Select the displayed application name and click **Next**.

f. For Application (optional), select **Next**.

g. For Review + create, review the policy details and click **Create**.



h. On the Azure Key Vault details page, click Refresh and search for the application name. It is included in the displayed list of applications with permission to access keys in this Azure Key Vault.

See Assign a Key Vault access Policy for more information.

Use Customer-Managed Encryption Keys on Autonomous Database with Azure Key Vault

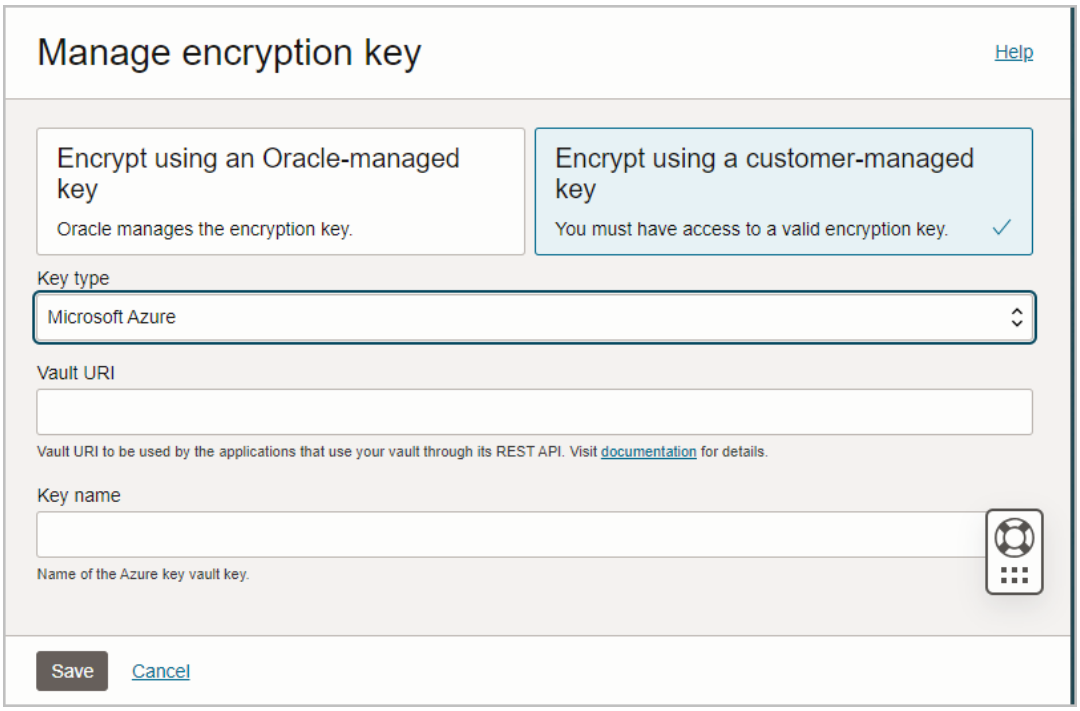
Shows the steps to encrypt your Autonomous Database using customer-managed master encryption keys that reside in Azure Key Vault.

Follow these steps:

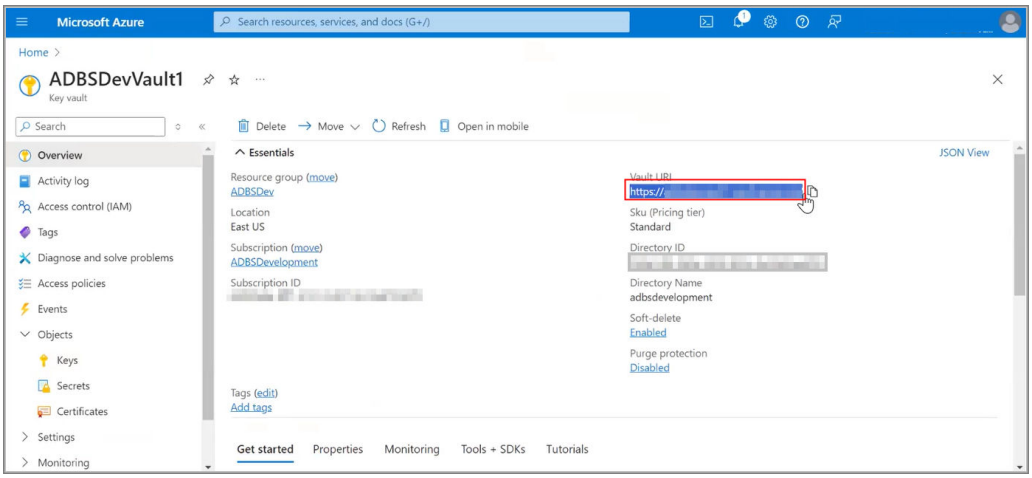
1. Perform the required customer-managed master key prerequisite steps. See Prerequisites to Use Customer-Managed Encryption Keys in Azure Key Vault.
2. On the **Details** page, from the **More actions** drop-down list, select **Manage encryption key**.

> **Note:**
>
> If you are already using customer-managed Transparent Data Encryption (TDE) keys stored in Azure Key Vault and you want to rotate the keys, follow these steps and select a different key (select a key that is different from the currently selected master TDE key).

3. On the **Manage encryption key** page, select **Encrypt using a customer-managed key**.
4. From the **Key type** drop-down, select **Microsoft Azure**.

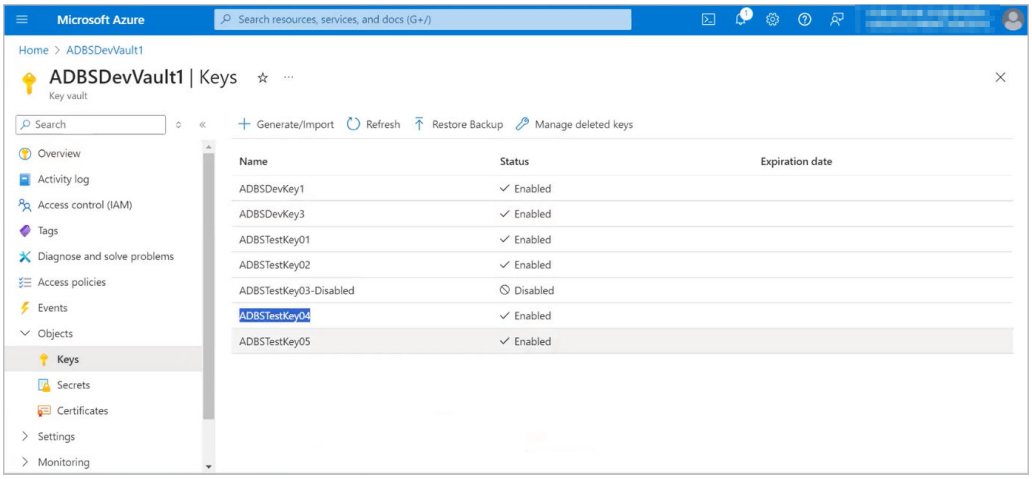


5. In the **Vault URI** field, enter the **Azure Vault URI**.
   a. In the Azure portal, navigate to the Azure Key Vault.
   b. Select the Azure Key Vault **Overview** page and copy the displayed **Vault URI**.



   c. Enter the copied Azure Vault URI into the **Vault URI** field in the Autonomous Database Manage encryption key page.
6. In the **Key name** field, enter the name of the **Azure Key Name**.
   a. In the Azure portal, navigate to the Azure Key Vault and select **Keys**. A list of keys for this vault is displayed.

b. From the list of displayed keys, copy the **Key Name** you want to use.



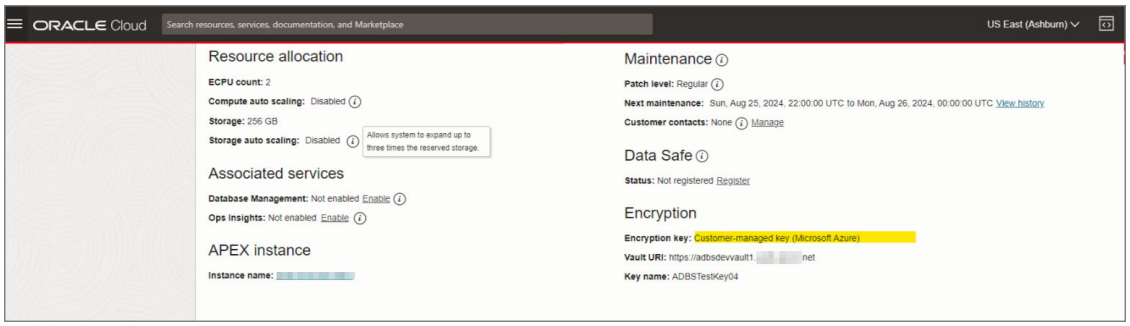
c. Enter the copied Azure key name into the **Key name** field in the Autonomous Database Manage encryption key page.

7. Click **Save**.

The **Lifecycle state** changes to **Updating**. When the request completes, the **Lifecycle state** shows **Available**.

After the request completes, on the Oracle Cloud Infrastructure Console, the key information shows on the Autonomous Database Information page under the heading **Encryption**. This area shows the **Encryption Key** is **Customer-managed key (Microsoft Azure)**, and displays the **Vault URI** and **Key name**.

For example:



See Notes for Using Customer-Managed Keys with Autonomous Database for more information.

Manage Master Encryption Keys in AWS Key Management Service

Autonomous Database supports customer-managed Transparent Data Encryption (TDE) keys that reside in AWS Key Management Service (KMS).

- Prerequisites to Use Customer-Managed Encryption Keys in AWS Key Management Service
  Describes prerequisite steps to use customer-managed master encryption keys that reside in Amazon Web Services (AWS) Key Management Service (KMS) on Autonomous Database.
- Use Customer-Managed Encryption Keys on Autonomous Database with AWS Key Management Service
  Shows the steps to encrypt your Autonomous Database using customer-managed master encryption keys that reside in AWS Key Management Service (KMS).

Prerequisites to Use Customer-Managed Encryption Keys in AWS Key Management Service

Describes prerequisite steps to use customer-managed master encryption keys that reside in Amazon Web Services (AWS) Key Management Service (KMS) on Autonomous Database.

Limitations:

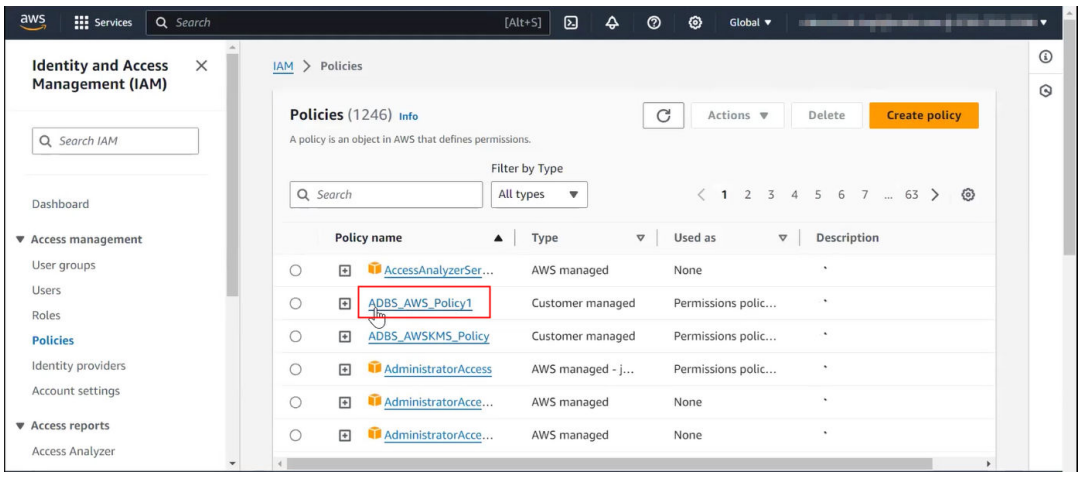
- AWS KMS is only supported in commercial regions.
- Cross-tenancy access, where the Autonomous Database instance and AWS KMS are in different tenancies, is not supported.
- AWS KMS is not supported in cross-region standbys.
- AWS KMS is not supported in refreshable clones.

Follow these steps:

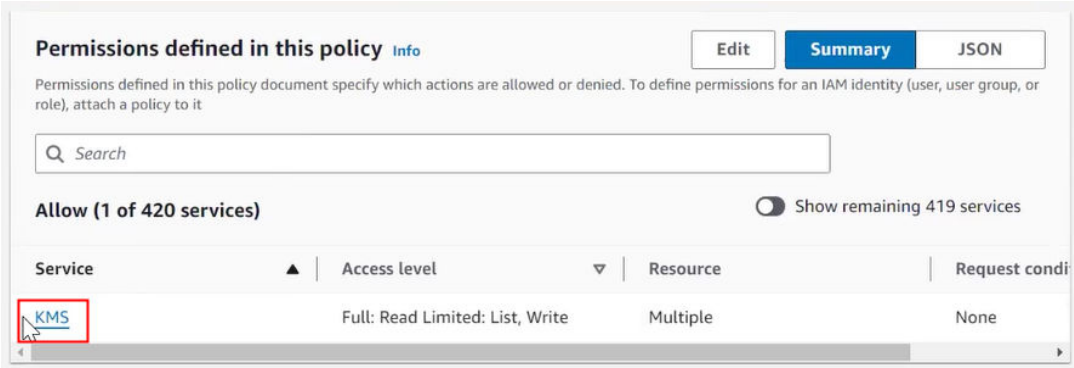
1. Create an AWS policy that grants read access to AWS KMS.

   See Creating an IAM policy to access AWS KMS for instructions, and Perform AWS Management Prerequisites to Use Amazon Resource Names (ARNs) for more information.

   For example, the `ADBS_AWS_Policy1` policy has been created:



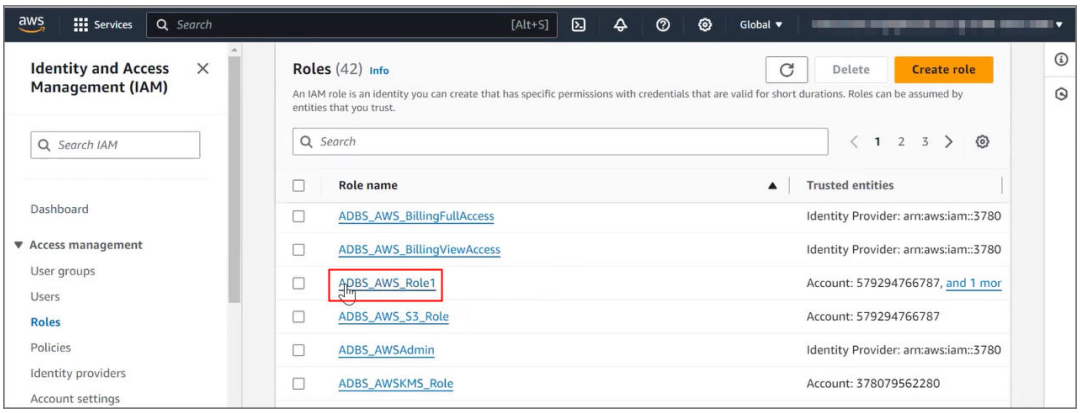
   The `ADBS_AWS_Policy1` policy includes permission to access **KMS**.



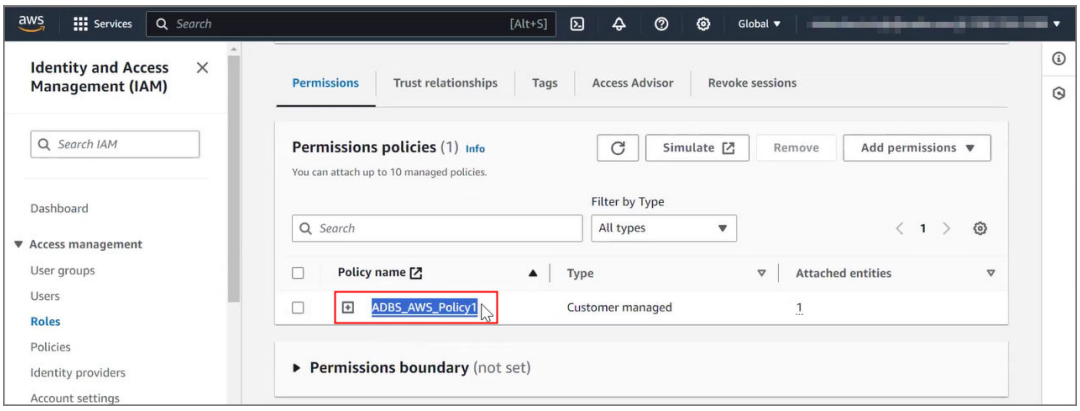
2. Create an AWS role and attach the policy to the role.

   See Creating an IAM role to access AWS services for instructions.

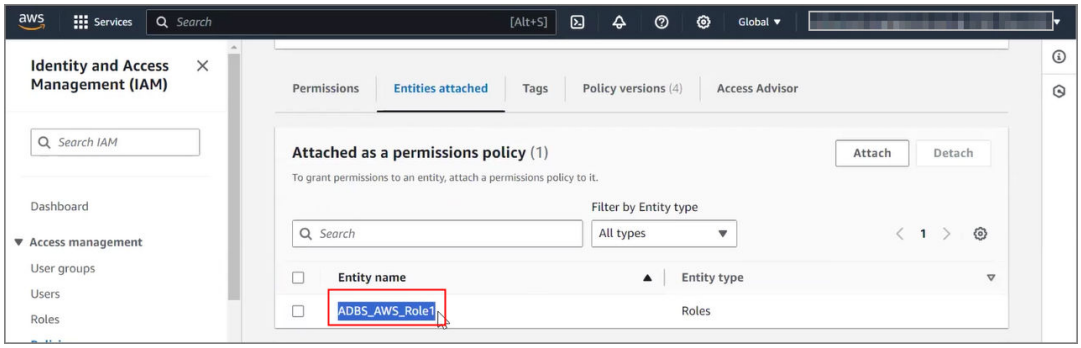
   For example, an `ADBS_AWS_Role1` role has been created:



In this example, the `ADBS_AWS_Policy1` policy is attached to the `ADBS_AWS_Role1` role:



On the policy details page, for this example, the role is listed under **Attached as permissions policy**:



3. Specify a Trust Relationship for the role.

   Edit the AWS Role's Trust Relationship to include Oracle's User ARN, and an External ID (tenancy OCID) for additional security.

   a. On Autonomous Database query `CLOUD_INTEGRATIONS`.

      For example:

      ```
      SELECT * FROM CLOUD.INTEGRATIONS;


      SELECT * FROM CLOUD_INTEGRATIONS;

      PARAM_NAME          PARAM_VALUE
      ---------------
      ------------------------------------------------------------------------
      ----------------------------------------------------------------
      aws_arn             arn:aws:iam:…:user/oraclearn
      ```

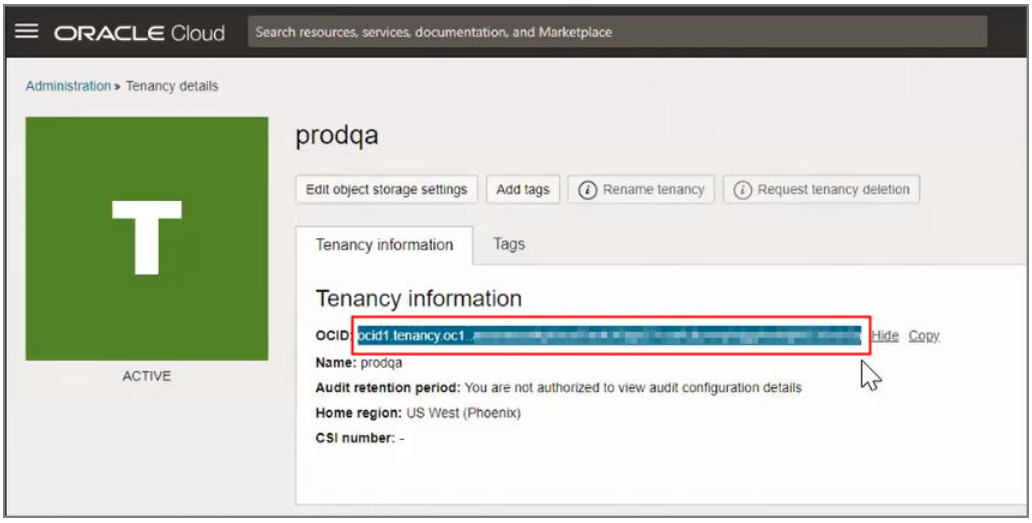
      The view `CLOUD_INTEGRATIONS` is available to the `ADMIN` user or to a user with `DWROLE` role.

   b. Copy the `PARAM_VALUE` for `aws_user_arn` and save the value for a subsequent step.

   c. Get the tenancy OCID, needed for the External ID.

      In the OCI console, click on your **Profile**, and select **Tenancy** to go to the tenancy details page. Copy the tenancy OCID and save it for a subsequent step.

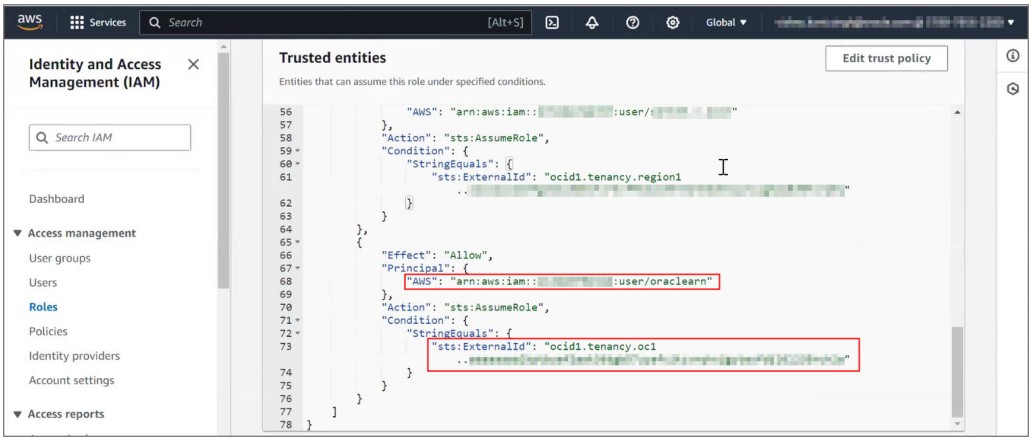
      For example:



d. On the AWS portal, navigate to the **Trusted entities** for the role, and scroll to the **"Principal"** statement.

e. For `"Principal"` specify **"AWS"** as the saved Oracle User ARN and for `"Condition"` specify **"sts:ExternalId"** as the saved OCID.

For example:



Use Customer-Managed Encryption Keys on Autonomous Database with AWS Key Management Service

Shows the steps to encrypt your Autonomous Database using customer-managed master encryption keys that reside in AWS Key Management Service (KMS).

Follow these steps:

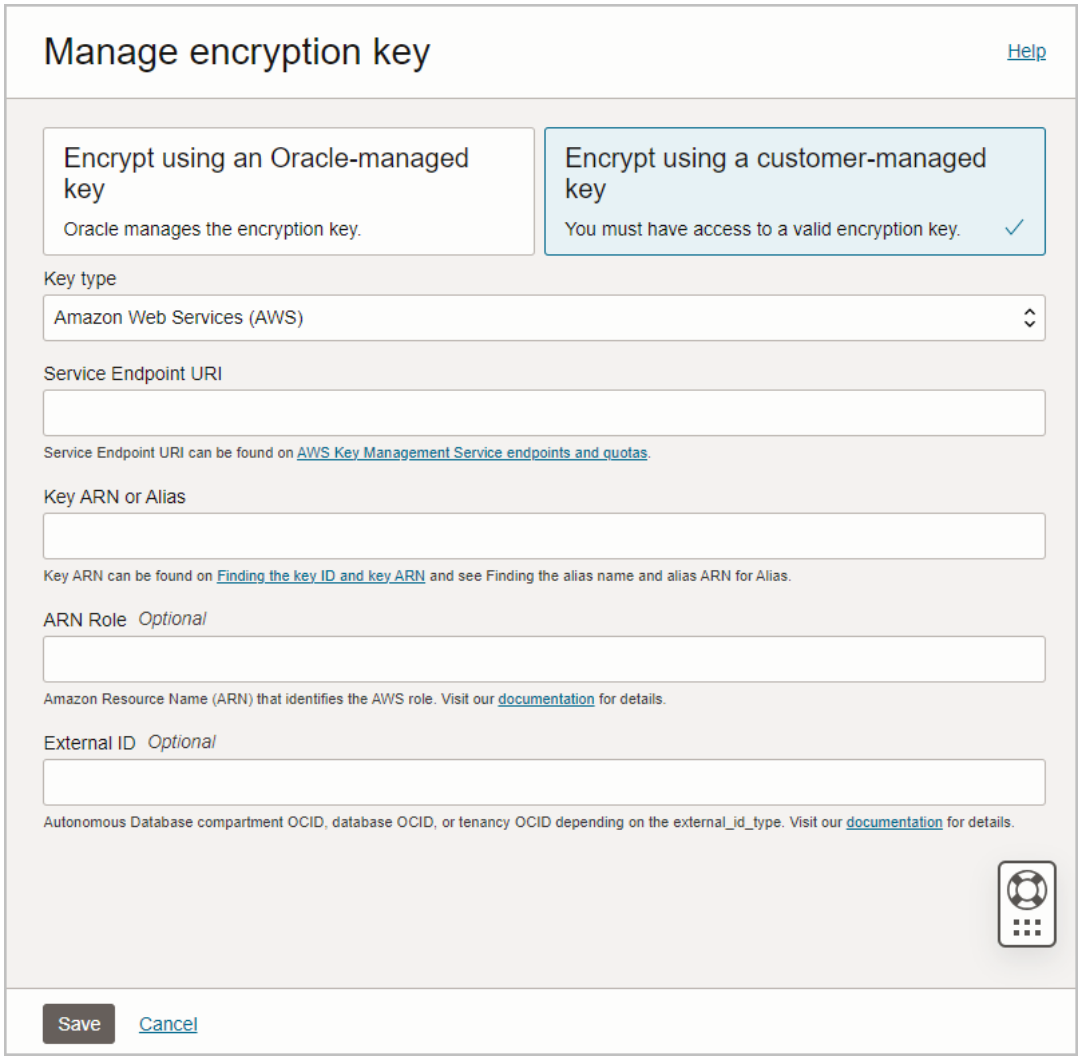
1. Perform the required customer-managed encryption key prerequisite steps as necessary. See Prerequisites to Use Customer-Managed Encryption Keys in AWS Key Management Service for details.
2. Create an Autonomous Database instance that uses the default Encryption key setting of **Encrypt using an Oracle-managed key**. See Provision an Autonomous Database Instance for more information.

> **Note:**
>
> Encryption key settings for customer-managed keys in AWS Key Vault are not available during the Autonomous Database instance creation process. The options are available post provisioning, when editing the instance.

3. On the **Details** page for the Autonomous Database instance, click **More actions**, and select **Manage encryption key**.

> **Note:**
>
> If you are already using customer-managed keys in AWS KMS and you want to rotate the TDE keys, follow these steps and select a different key (select a key that is different from the currently selected master encryption key).

4. On the **Manage encryption key** page, select **Encrypt using a customer-managed key**.
5. From the **Key type** drop-down, select **Amazon Web Services (AWS)**.



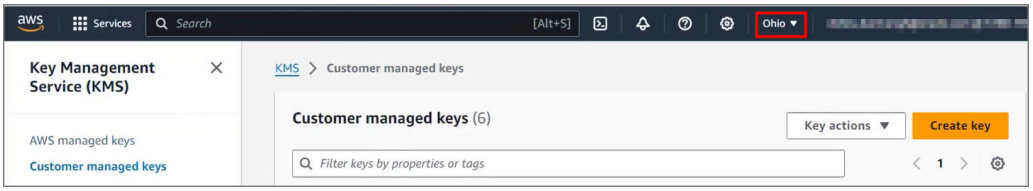
6. Enter the Service Endpoint URI.

   The Service Endpoint URI is the AWS region where the AWS KMS is located.

   a. Go to the AWS portal, navigate to the KMS where your key is located.

   b. Find the region name listed in the top bar of the portal.

      For example, this KMS is in the region named **Ohio**:



   c. Look up the endpoint corresponding to the region. Go to AWS Key Management Service endpoints and quotas and find the endpoint for the AWS region name where your AWS KMS is located.

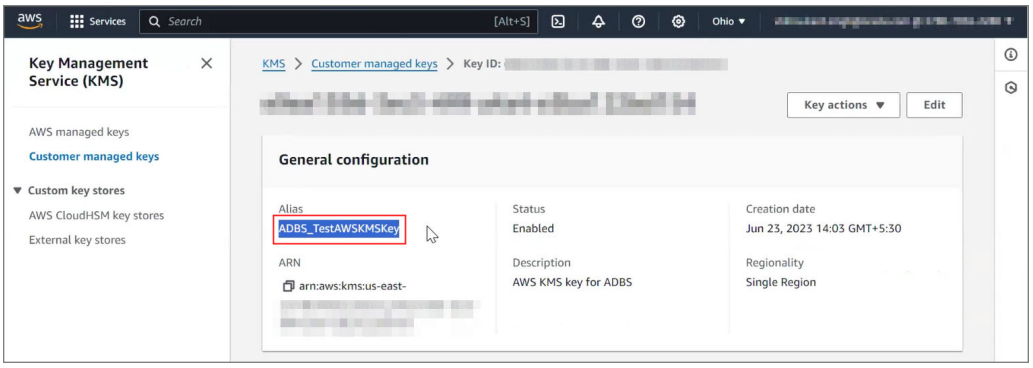
      For example, if the AWS region name is `Ohio` the endpoint is `kms.us-east-2.amazonaws.com`.

   d. Enter the endpoint for the **Service Endpoint URI**.

7. Enter the Key ARN or Alias.

   a. Navigate to the key details page on the AWS portal. Copy the key's Alias or ARN.

      For example, the alias for `ADBS_TestAWSKMSKey` is selected:



   b. Enter the key's Alias or ARN into the **Key ARN or Alias** field.

      If entering the Alias, prefix the entry with `alias/`. For example, if the alias is `ADBS_TestAWSKMSKey` enter:

      ```
      alias/ADBS_TestAWSKMSKey
      ```

      If entering the ARN, no prefix is needed. For example, if the ARN is `arn.aws.kms.us-east-2:37807956...bd154` enter:

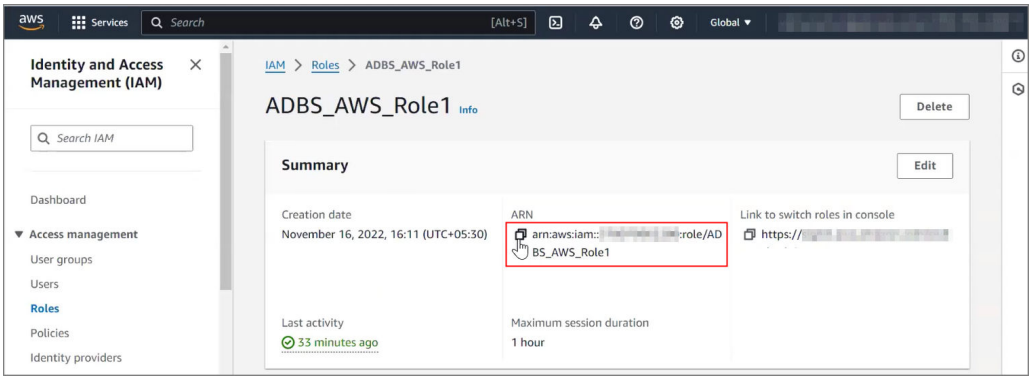
      ```
      arn.aws.kms.us-east-2:37807956...bd154
      ```

8. Enter ARN Role (Optional).

a. Navigate to the role details page on the AWS portal.

b. Copy the role's ARN.

For example, the ARN for `ADBS_AWS_Role1` is copied:



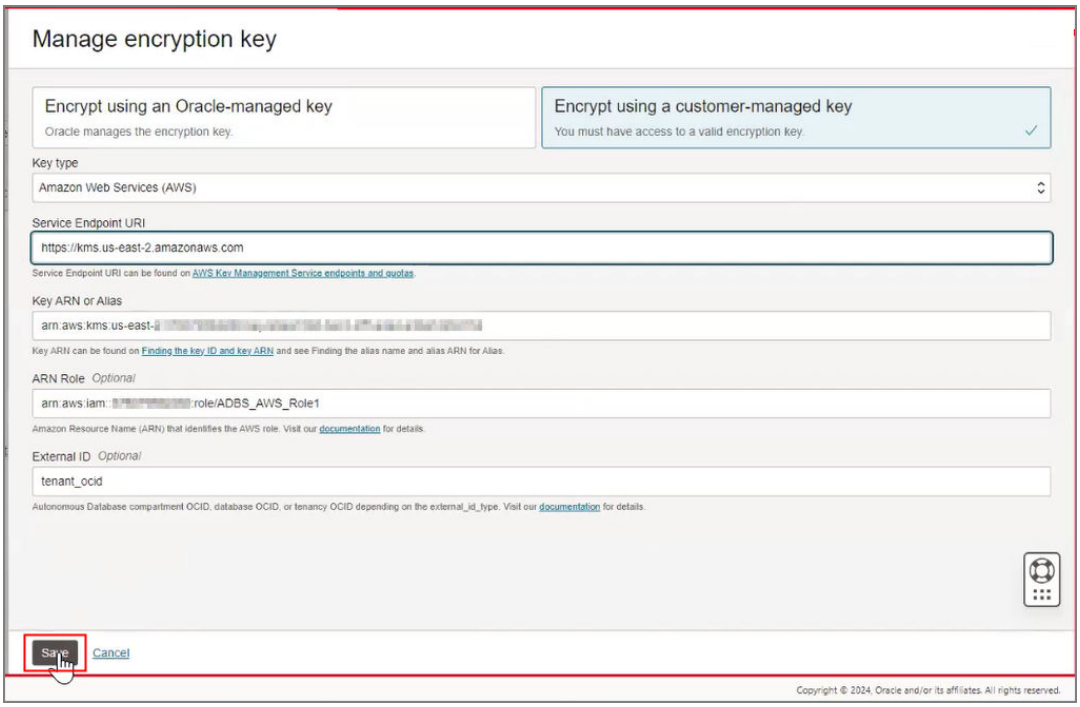
c. Enter the copied ARN into the **ARN Role field**.

9. Enter External ID (Optional).

For the External ID, enter **`tenant_ocid`**.

10. Click **Save**.

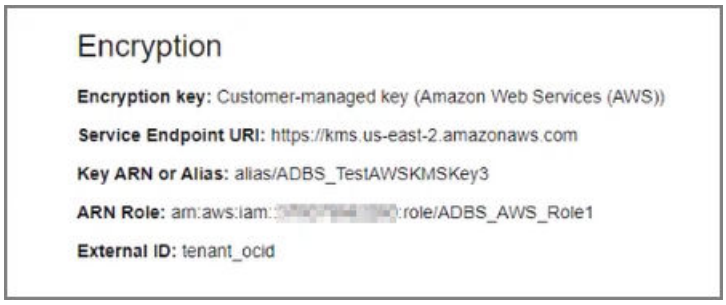
For example:



The **Lifecycle state** changes to **Updating**. When the request completes, the **Lifecycle state** shows **Available**.

After the request completes, on the Oracle Cloud Infrastructure Console, the key information shows on the Autonomous Database instance details page under the heading **Encryption**.

For example:



See Notes for Using Customer-Managed Keys with Autonomous Database for more information.

Manage Master Encryption Keys in Oracle Key Vault

Autonomous Database supports customer-managed Transparent Data Encryption (TDE) keys that reside in Oracle Key Vault (OKV).

- Prerequisites to Use Customer-Managed Encryption Keys in Oracle Key Vault
  Describes the prerequisite steps to use customer-managed master encryption keys that reside in Oracle Key Vault (OKV) on Autonomous Database.
- Use Customer-Managed Encryption Keys on Autonomous Database with Oracle Key Vault
  Shows the steps to encrypt your Autonomous Database using customer-managed master encryption keys that reside in Oracle Key Vault (OKV).

Prerequisites to Use Customer-Managed Encryption Keys in Oracle Key Vault

Describes the prerequisite steps to use customer-managed master encryption keys that reside in Oracle Key Vault (OKV) on Autonomous Database.

Requirements:

- The Autonomous Database instance must use private endpoints.
- Autonomous Database supports OKV instances that are in private networks and are accessible from the same network where the Autonomous Database resides.

  See VCN and Subnet Management and Oracle Key Vault Installation and Configuration, for more information.

Limitations:

- Cross-tenancy access, where the Autonomous Database instance and OKV are in different tenancies, is not supported.
- OKV is not supported in cross-region standbys.
- OKV is not supported in refreshable clones.

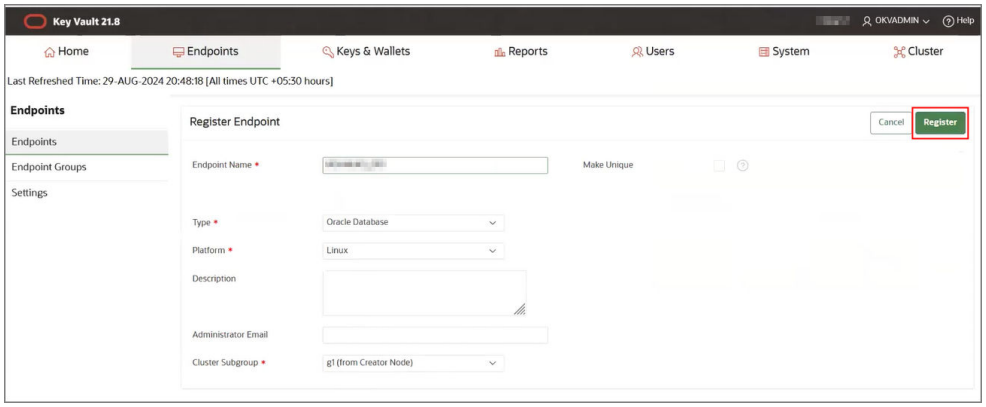
Follow these steps:

1. Create an OKV endpoint, wallet and TDE master key.
   a. Sign into the OKV instance.
      - Copy the **Public IP address**, or **Private IP address** on the OKV instance details page and paste into a browser.

- On the OKV instance login page, enter the username and password.

b. Create and register an OKV endpoint.

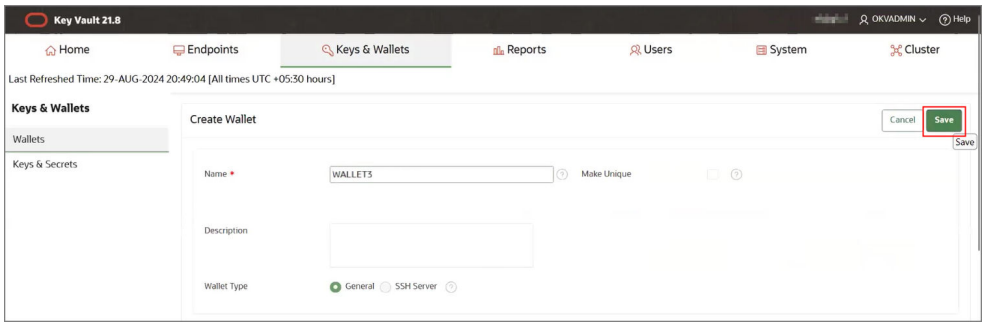
- Click **Endpoints**, on the OKV home page.
- Click **Add**, on the Endpoints details page and enter a name for the endpoint.
- Allow all other settings to default and click **Register**.
  For example:



See Adding, Deleting, or Reenroling Endpoints for more information.

c. Create a wallet.
The OKV wallet holds the TDE master encryption key.

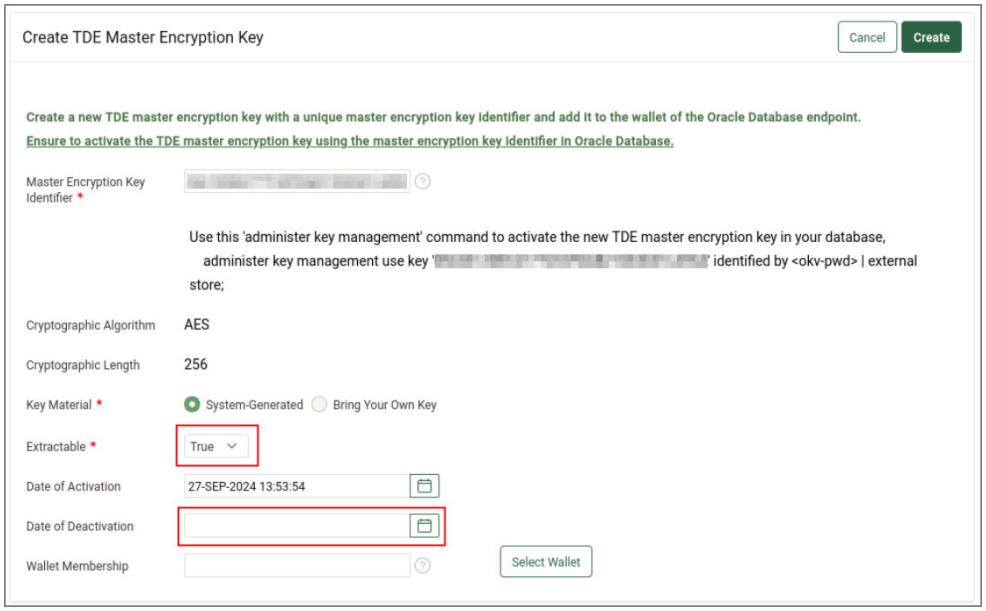
- From the OKV home page, navigate to the **Keys & Wallets** page.
- For **Wallets** on the **Keys & Wallets** page, click **Create**.
- Enter the name of the wallet in the **Name** field and click **Save**.
  For example:



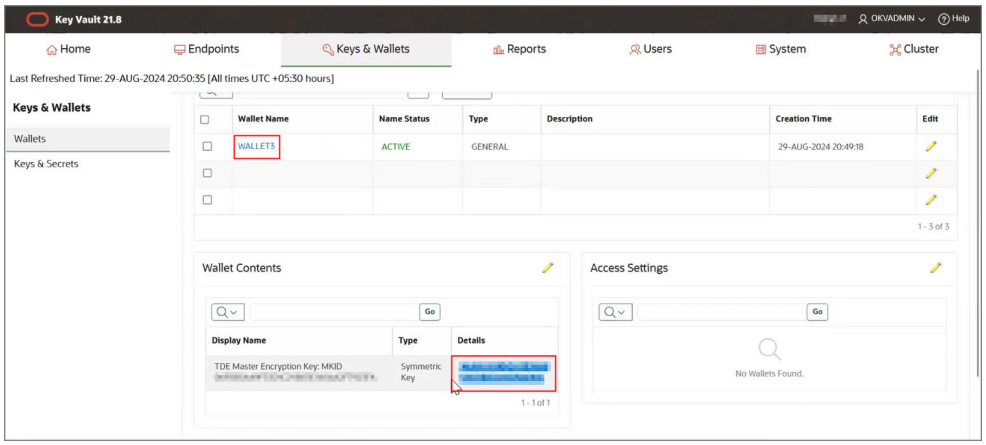
See Creating a Virtual Wallet for more information.

d. Create a TDE Master Encryption Key in the wallet.

- Select **Keys & Secrets**, and click **Create**.
- For **Application Keys**, select **TDE Master Encryption Key**.
- For **Extractable**, select **True**.
- Make sure the **Date of Deactivation** is empty.
  For example:



- For **Wallet Membership**, click **Select Wallet**.
- From the displayed list of wallets, select the wallet created in the previous step and click **Close**.
- Click **Create**. The TDE Master Encryption Key is created and displayed under **Wallet Contents**.
  For example:



e. Modify the endpoint to make the previously created wallet the default wallet.
   - Navigate to the endpoint's details page.
   - On the Default Wallet pane, select **Choose Wallet**.
   - On the Choose Wallet page, select the previously created wallet and click **Select**.
   - Click **Save**.

f. Enable Restful services.

> **Note:**
>
> Restful services must be enabled on the OKV instance to successfully run the curl command, in a subsequent step, to download the wallet.

- From the OKV home page, select the **System** tab.
- On the left navigation pane, click **Setting** .
- Under System Configuration, select **Restful services**.
- Click **All**, then **Save**.

2. Provision an Autonomous Database instance, with the following required settings:
   a. For **Choose network access**, select **Private endpoint access only**.
   b. For **Virtual cloud network**, select the VCN where this database instance is running.
   c. For **Subnet**, select the private subnet where this database instance is running.
   d. For **Network security groups (NSGs)**, select the security group.
   e. For **Encryption key** settings, default to **Encryption using an Oracle-managed key**. These settings are changed to customer managed keys in OKV, after these prerequisite steps are completed. The customer managed keys are disabled during instance provisioning. See Use Customer-Managed Encryption Keys on Autonomous Database with Oracle Key Vault for details.
3. Connect to the Autonomous Database instance and create a directory for the OKV wallet.
   a. Connect to the Private Endpoint Autonomous Database instance as the ADMIN user.

      For example, connect to the OKVDEMO1 database instance:

      ```
      SQL> connect ADMIN/<password>@OKVDEMO1_low
      ```

   b. Create a directory object in the Autonomous Database instance.

      For example:

      ```
      SQL> create directory okv_dir as 'okvdir';
      ```

   c. Check to make sure the directory is created.

      For example, the following statements create the `OKV_DIR` directory object and the statement results display the directory name (`OKV_DIR`) and directory path (`/u03/dbfs/<path data>/data/okvdir`).

      ```
      SQL> connect ADMIN/<admin password>#@OKVDEMO1_low
      Connected
      SQL>
      SQL> create directory okv_dir as 'okvdir';
      Directory created.
      SQL> select * from dba_directories where directory_name = 'OKV_DIR';
      OWNER
      ------------------------------------------------------------------------
      DIRECTORY_NAME
      ------------------------------------------------------------------------
      DIRECTORY_PATH
      ------------------------------------------------------------------------
      ```

```
ORIGIN_CON_ID
-------------
SYS
OKV_DIR
/u03/dbfs/<path data>/data/okvdir
SQL>
```

4. Download the OKV wallet to the directory object created in the Autonomous Database instance.

   a. Download the wallet for the endpoint from the OKV instance.

      Run the following command from the client:

      ```
      curl -k -X POST
       --location https://<OKV server>:5695/okv/cloud/utility/endpoint/
      download/sso
       --data "token=<Enrollment Token>"
       --output cwallet.sso
      ```

      Where:

      - `OKV server`: is the Internal Fully Qualified Domain Name (FQDN) or Private IP found on the OKV instance details page.
      - `Enrollment Token`: is the enrollment token for the endpoint found on the Endpoints page.

      After the wallet is downloaded, the endpoint on the OKV instance shows a status of enrolled. Alternatively, you can also download the wallet by selecting the endpoint name and clicking "Download (wallet)".

   b. Upload the wallet to Object Storage.

      Upload the wallet file from your local machine to your Object Storage bucket using Upload Object. See Uploading an Object to a Bucket, for more information.

   c. In Object Storage, generate a Pre-Authenticated Request (PAR) URL for the uploaded wallet file. See Creating a Pre-Authenticated Request in Object Storage for more information.

   d. From the VM, connect to the database instance as ADMIN user.

   e. In the database instance, run the `DBMS_CLOUD.GET_OBJECT` procedure to download the wallet from Object Storage to the database instance wallet directory.

      For example:

      ```
      BEGIN
          DBMS_CLOUD.GET_OBJECT(
              object_uri => '<PAR URL>',
              directory_name => '<wallet_dir>');
      END;
      /
      ```

      - Where:
        - `object_uri` is the PAR URL generated for the wallet file in Object Storage.
        - `directory_name` is the name of the wallet directory created in the database instance.

See GET_OBJECT Procedure and Function for more information.

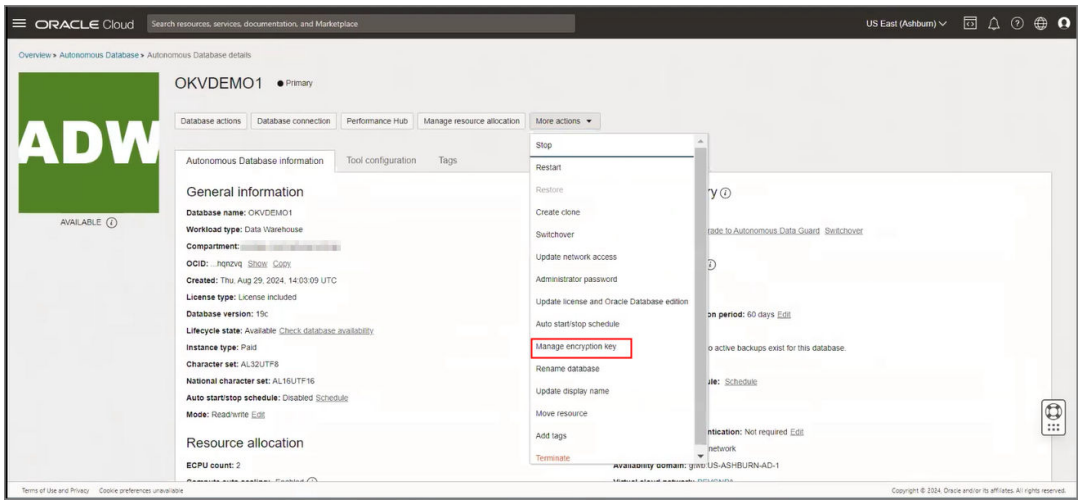
f. Delete the wallet from Object Storage.

Use Customer-Managed Encryption Keys on Autonomous Database with Oracle Key Vault

Shows the steps to encrypt your Autonomous Database using customer-managed master encryption keys that reside in Oracle Key Vault (OKV).

Follow these steps:

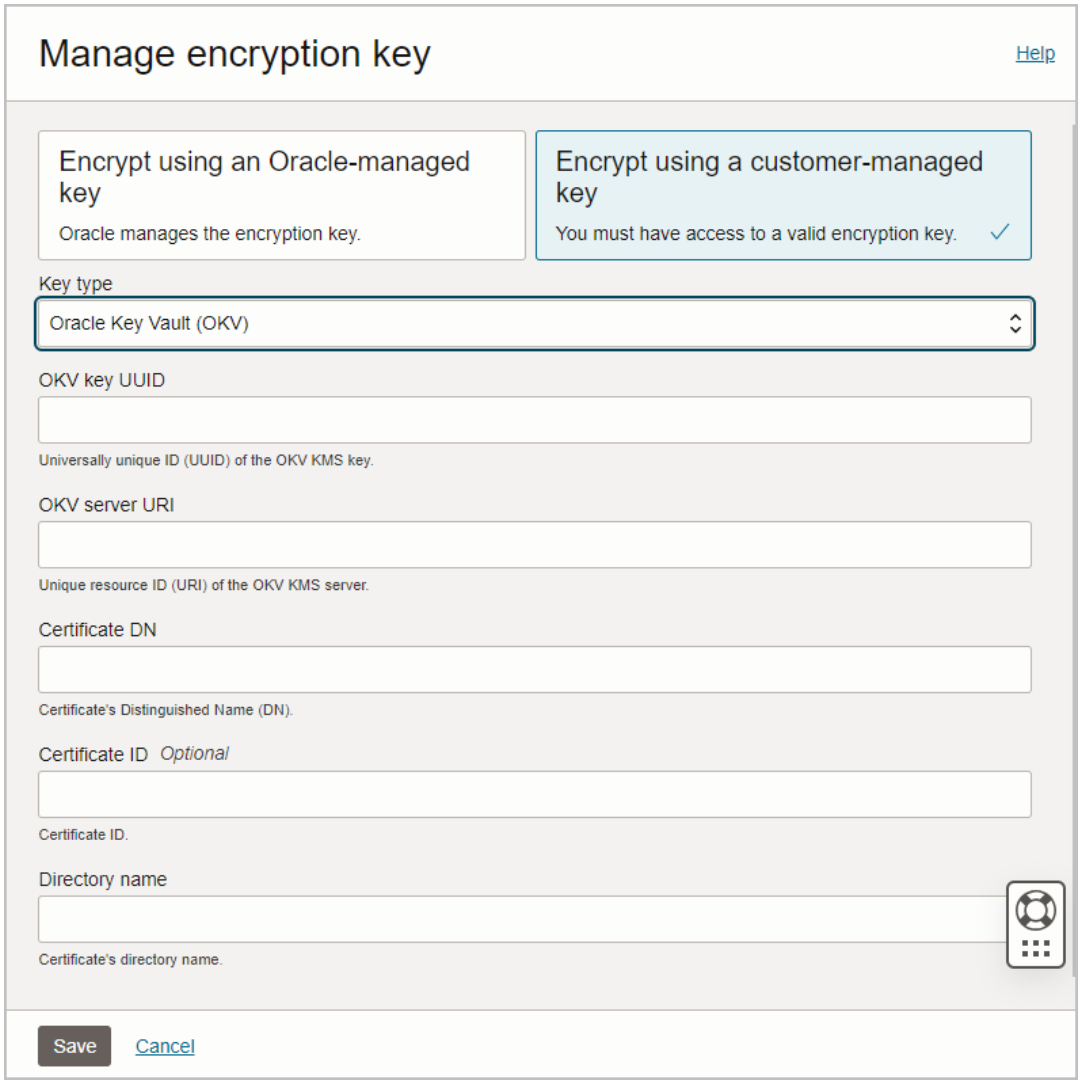
1. Perform the required customer-managed encryption key prerequisite steps as necessary. See Prerequisites to Use Customer-Managed Encryption Keys in Oracle Key Vault for details.
2. On the **Details** page for the Autonomous Database instance, click **More actions**, and select **Manage encryption key**.



> **Note:**
>
> If you are already using customer-managed keys in OKV and you want to rotate the TDE keys, follow these steps and select a different key (select a key that is different from the currently selected master encryption key). However, you cannot use an OKV key that has been used previously on the same Autonomous Database instance.

3. On the **Manage encryption key** page, select **Encrypt using a customer-managed key**.
4. On the **Key type** drop-down, select **Oracle Key Vault (OKV)**.

   The Manage encryption key dialog displays the Oracle Key Vault (OKV) options.



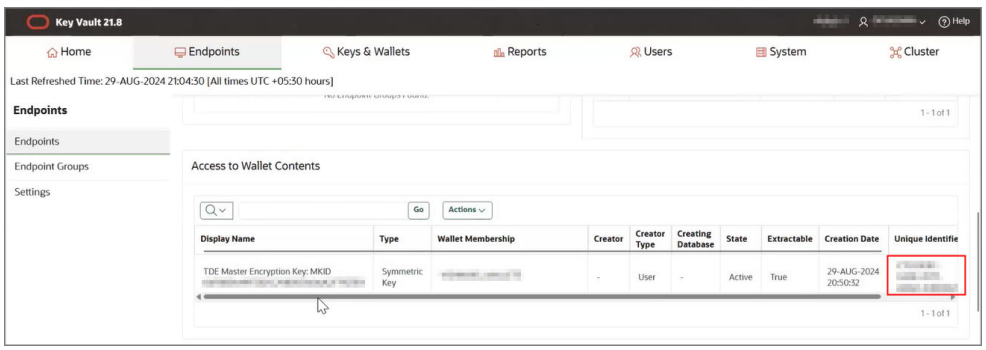
5. Enter the following information:

- **OKV UUID** - Enter the TDE master key's Unique Identifier for the key located in the OKV instance.
  To find this value:

  a. Sign into the OKV instance, and select **Endpoints**.

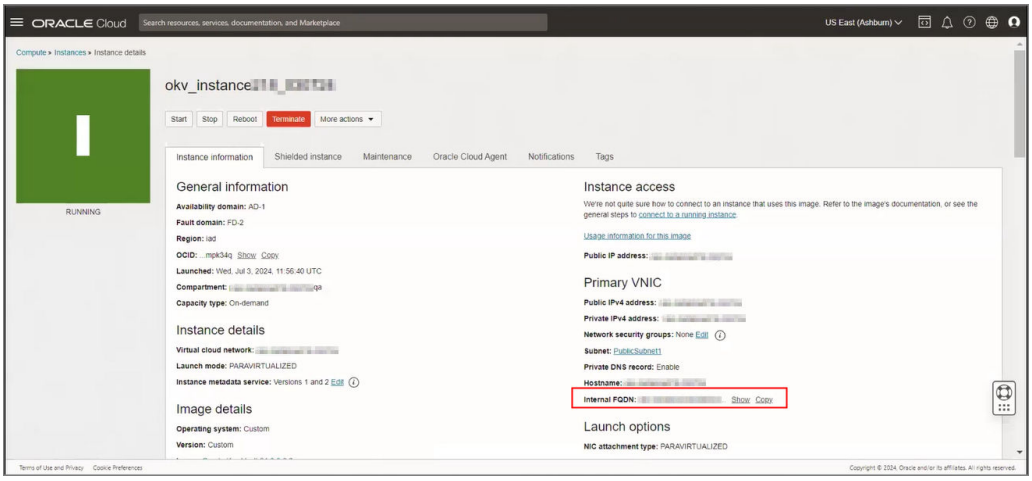
  b. Select the endpoint that contains the wallet and key.

  c. Scroll to **Access to Wallet Contents**, and copy the **Unique Identifier**.



d. Paste the Unique Identifier into the **OKV UUID** field.

- **OKV Server URI** - Enter the **Internal Fully Qualified Domain Name** or the Private IP found on the OKV instance details page.
  For example, if using an OCI VM to host OKV, this value can be found on the OKV instance details page under **Primary VNIC**:



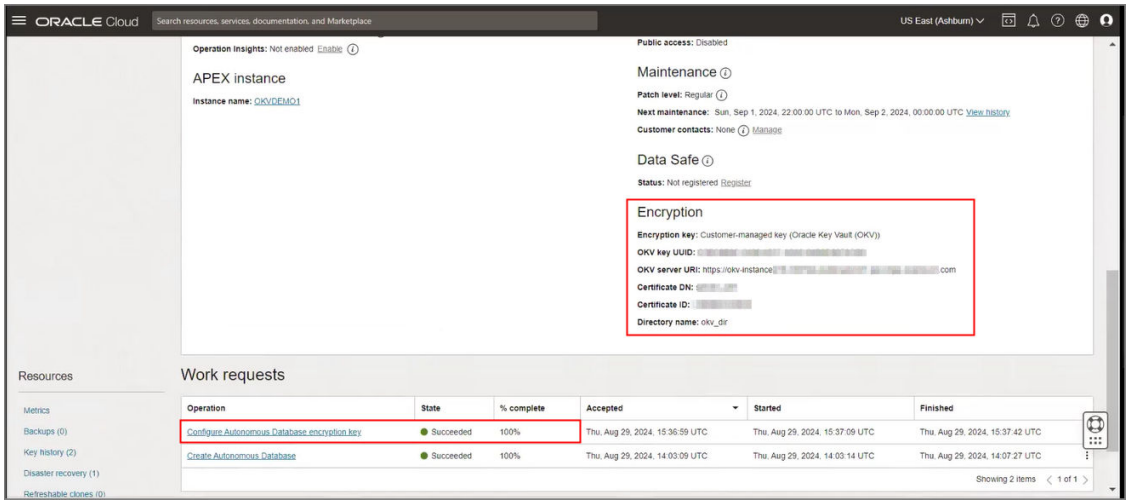
- **Certificate DN** - Enter your certificate Distinguished Name (DN).
- **Certificate ID** (*Optional*) - Enter your certificate ID or leave blank.

> **Note:**
>
> This field is optional if using OKV versions 21.9 and above. If using OKV versions below 21.9, the certificate ID is required.

- **Directory name** - Enter the directory name where the wallet is saved on the Autonomous Database instance.

6. Click **Save**.

When the save completes successfully, encryption settings for the Autonomous Database instance are updated to show **Customer-managed key (Oracle Key Vault (OKV))** and the work request state shows succeeded.



See Notes for Using Customer-Managed Keys with Autonomous Database for more information.

Switch to Oracle-Managed Encryption Keys on Autonomous Database

Shows the steps to switch to Oracle-managed master encryption keys on Autonomous Database if you are using customer-managed encryption keys.

If your Autonomous Database is using customer-managed keys and you want to switch to Oracle-managed keys:

1. On the **Details** page, from the **More actions** drop-down list, select **Manage Encryption Key**.
2. On the **Manage Encryption Key** page, select **Encrypt using Oracle-managed keys**.
3. Click **Save Changes**.

The **Lifecycle state** changes to **Updating**. When the request completes, the **Lifecycle state** shows **Available**.

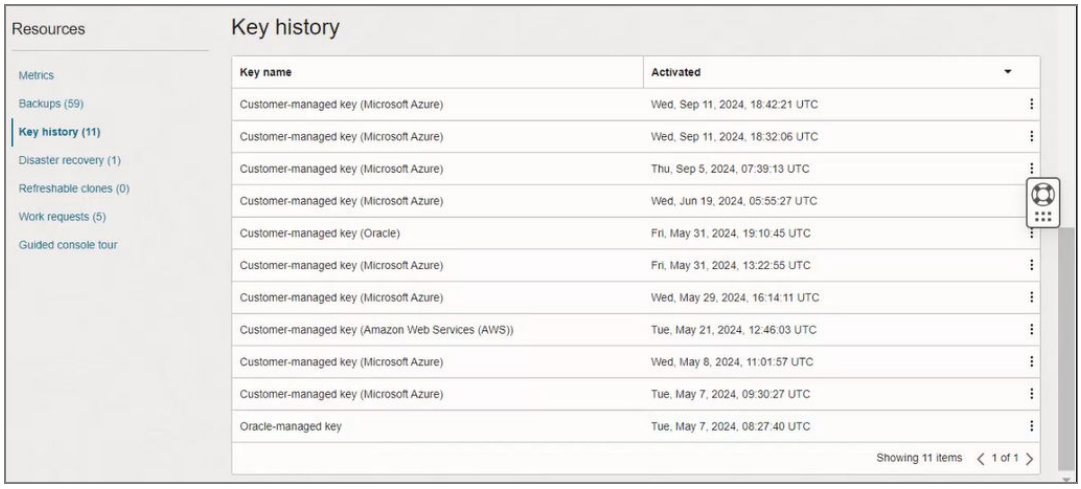
View History for Customer-Managed Encryption Keys on Autonomous Database

You can view the key history from the Oracle Cloud Infrastructure Console or by selecting from the `V$ENCRYPTION_KEYS` view.

To view the key history from the Oracle Cloud Infrastructure Console:

1. On the Autonomous Database Details page, under **Resources**, click **Key History**.
2. This shows the Key History.

   For example:



To view the key history by selecting from the `V$ENCRYPTION_KEYS` view:

1. Connect to the Autonomous Database as the ADMIN user.
2. As the ADMIN user, issue the following command:

```
SELECT activation_time, tag FROM V$ENCRYPTION_KEYS;
```

For example:

```
SELECT activation_time, tag FROM V$ENCRYPTION_KEYS;
ACTIVATION_TIME                TAG
-----------------------        -----------------------
2020-02-19T17:21:57.821Z       {"credential_name":"OCI$RESOURCE_PRINCIPAL",
                                "oci_management_url":"https://aqm-
management.kms.ca-toronto-1.oraclecloud.com/20180608/keys/",
                                "master_key_id":"ocid1.key.oc1.ca-
toronto-1","username":"\"ADMIN\"",
                                "vault_id":"ocid1.vault.oc1.ca-toronto-1"}
```

See Notes for Using Customer-Managed Keys with Autonomous Database for more information.

Notes for Using Customer-Managed Keys with Autonomous Database

Provides additional information and notes for using customer-managed keys with Autonomous Database

- Caution for Customer-Managed Encryption Keys when the Key is Unreachable
  After you switch to customer-managed keys, some database operations might be affected when the key is unreachable.
- Caution for Using Customer-Managed Encryption Keys When the Master Key is Disabled or Deleted
  After you switch to customer-managed keys, some database operations might be affected if the master key is disabled or deleted.
- Caution for Using Customer-Managed Encryption Keys History and Backups
  After you switch to customer-managed keys, some database operations might be affected if the master key is rotated, disabled, or deleted and you do not have a valid key to restore your data from a previously saved backup or from a clone.

- Notes for Customer-Managed Keys in OCI Vault with Autonomous Data Guard
  Autonomous Data Guard is supported when using customer-managed keys with the key provider Oracle Cloud Infrastructure Vault.
- Notes for Customer-Managed Encryption Keys with Cloning
  Autonomous Database does not support using customer-managed encryption keys with a refreshable clone.
- Notes for Customer-Managed Keys with Vault in Remote Tenancy
  To use customer-managed master encryption keys with a Vault in a remote tenancy, note the following:

Caution for Customer-Managed Encryption Keys when the Key is Unreachable

After you switch to customer-managed keys, some database operations might be affected when the key is unreachable.

If the database is unable to access your key for some reason, such as a network outage, then Autonomous Database handles the outage as follows:

- There is a 2-hour grace period where the database remains up and running.
- If the key is not reachable at the end of the 2-hour grace period, the database Lifecycle state is set to **Inaccessible**. In this state existing connections are dropped and new connections are not allowed.
- If Autonomous Data Guard is enabled when using Oracle Cloud Infrastructure Vault, during or after the 2-hour grace period you can manually try to perform a failover operation. Autonomous Data Guard automatic failover is not triggered when you are using customer-managed encryption keys and the Oracle Cloud Infrastructure Vault is unreachable.

  > **Note:**
  >
  > Only Oracle Cloud Infrastructure Vault is supported with Autonomous Data Guard.

- If Autonomous Database is stopped, then you cannot start the database when the keys are unreachable.

  For this case, the work request shown when you click **Work Requests** on the Oracle Cloud Infrastructure console under Resources shows:

  ```
  The Vault service is not accessible.
  Your Autonomous Database could not be started. Please contact Oracle
  Support.
  ```

Caution for Using Customer-Managed Encryption Keys When the Master Key is Disabled or Deleted

After you switch to customer-managed keys, some database operations might be affected if the master key is disabled or deleted.

- For disable/delete key operations where the Master Encryption Key **State** is any of the following:

Key Vault	Key State
Oracle Cloud Infrastructure (OCI) Vault	– DISABLING – DISABLED – DELETING – DELETED – SCHEDULING_DELETION – PENDING_DELETION
AWS Key Management System (KMS)	– Disabled – PendingDeletion
Azure Key Vault	– DISABLED – DELETED
Oracle Key Vault (OKV)	– COMPROMISED – DEACTIVATED – DESTROYED

The database becomes **Inaccessible** after the key Lifecycle state goes into one of these states. When the key is in any of these states, Autonomous Database drops existing connections and does not allow new connections.

The database becoming inaccessible can take up to 15 min after the key operations (Autonomous Database checks the key provider at 15 minute intervals).

- If you disable or delete the Oracle Cloud Infrastructure Vault key used by your Autonomous Database while Autonomous Data Guard is enabled, Autonomous Data Guard will not perform an automatic failover.

> **Note:**
>
> Only Oracle Cloud Infrastructure Vault is supported with Autonomous Data Guard.

- If you enable a disabled key, the database automatically goes into the Available state.
- If you delete the master key you can check the key history in the Oracle Cloud Infrastructure Console to find the keys that were used for the database. The history shows you the key name, along with an activation timestamp. If one of the older keys from the history list is still available, then you can restore the database to the time when the database was using that key, or you can clone from a backup to create a new database with that timestamp.

See View History for Customer-Managed Encryption Keys on Autonomous Database for more information.

Caution for Using Customer-Managed Encryption Keys History and Backups

After you switch to customer-managed keys, some database operations might be affected if the master key is rotated, disabled, or deleted and you do not have a valid key to restore your data from a previously saved backup or from a clone.

- It is recommended that you create a new key that hasn't been used for rotation in the last 60 days and use that key for rotation. This makes sure that you can go back to a backup if the current key is deleted or disabled.
- When you perform multiple encryption key rotations within 60 days, it is recommended that you either create a new key for each master encryption key rotation operation or specify a

key that has not been used in the last 60 days. This helps to save at least one key you can use to recover your data from a backup, in the case where a customer-managed master encryption key is disabled or deleted.

See View History for Customer-Managed Encryption Keys on Autonomous Database for more information.

Notes for Customer-Managed Keys in OCI Vault with Autonomous Data Guard

Autonomous Data Guard is supported when using customer-managed keys with the key provider Oracle Cloud Infrastructure Vault.

> **Note:**
>
> When you are using Oracle-managed keys you can only switch to customer-managed keys from the primary database. Likewise, when you are using customer-managed keys you can only rotate keys or switch to Oracle-managed keys from the primary database. The **Manage encryption key** option under **More actions** is disabled on a standby database.

Note the following when you are using Autonomous Data Guard with a standby database with customer-managed keys:

- In order for a remote standby to use the same master encryption key as primary, the master encryption key must be replicated to the remote region.

  Customer-Managed Encryption Keys are only supported with a single cross-region Autonomous Data Guard standby. Multiple cross-region standbys are not supported because Oracle Cloud Infrastructure Vault only supports replication to one remote region.
- The primary database and the standby database use the same key. In the event of a switchover or failover to the standby, you can continue to use the same key.
- When you rotate the customer-managed key on the primary database, this is reflected on the standby database.
- When you switch from customer-managed keys to Oracle-managed keys the change is reflected on the standby database.
- When the Oracle Cloud Infrastructure Vault is unreachable, there is a 2-hour grace period before the database goes into the `INACCESSIBLE` state. You can perform a failover during or after this 2-hour grace period.
- If you disable or delete the Oracle Cloud Infrastructure master encryption key that your Autonomous Database is using with customer-managed keys while Autonomous Data Guard is enabled, Autonomous Data Guard does not perform an automatic failover.
- Customer-Managed Encryption Keys are not supported with Cross Tenancy Autonomous Data Guard.

See the following for more information:

- To Replicate a Vault and Its Keys
- Autonomous Data Guard with Customer Managed Keys
- Overview of Vault

Notes for Customer-Managed Encryption Keys with Cloning

Autonomous Database does not support using customer-managed encryption keys with a refreshable clone.

- You cannot create a refreshable clone from a source database that uses customer-managed encryption keys. Additionally, you cannot switch to a customer-managed encryption key on a source database that has one or more refreshable clones.

Notes for Customer-Managed Keys with Vault in Remote Tenancy

To use customer-managed master encryption keys with a Vault in a remote tenancy, note the following:

When you use customer-managed master encryption keys with a Vault in a remote tenancy, the Vault and the Autonomous Database instance must be in the same region. To change the tenancy, on the sign-on page click **Change tenancy**. After you change the tenancy, make sure to select the same region for both the Vault and the Autonomous Database instance.

Configure Network Access with Access Control Rules (ACLs) and Private Endpoints

Provides details on how to configure network access with access control rules or using a private endpoint and describes the secure client connection options. Also describes how to enable support for TLS connections (require mutual TLS only or allow both mutual TLS and TLS authentication).

- Configuring Network Access with Access Control Rules (ACLs)
  Specifying an access control list blocks all IP addresses that are not in the ACL list from accessing the database. After you specify an access control list, the Autonomous Database only accepts connections from addresses on the access control list and the database rejects all other client connections.
- Configure Network Access with Private Endpoints
  You can specify that Autonomous Database uses a private endpoint inside your Virtual Cloud Network (VCN) in your tenancy. You can configure a private endpoint during provisioning or cloning your Autonomous Database, or you can switch to using a private endpoint in an existing database that uses a public endpoint. This allows you to keep all traffic to and from your database off of the public internet.
- Use Zero Trust Packet Routing on Autonomous Database
  Oracle Cloud Infrastructure Zero Trust Packet Routing (ZPR) protects sensitive data from unauthorized access through intent-based security policies that you write for resources, such as an Autonomous Database on a private endpoint, that you assign security attributes to.
- Update Network Options to Allow TLS or Require Only Mutual TLS (mTLS) Authentication on Autonomous Database
  Describes how to update the secure client connection authentication options, Mutual TLS (mTLS) and TLS.

Configuring Network Access with Access Control Rules (ACLs)

Specifying an access control list blocks all IP addresses that are not in the ACL list from accessing the database. After you specify an access control list, the Autonomous Database

only accepts connections from addresses on the access control list and the database rejects all other client connections.

- Configure Access Control Lists When You Provision or Clone an Instance
  When you provision or clone Autonomous Database with the **Secure access from allowed IPs and VCNs only** option, you can restrict network access by defining Access Control Lists (ACLs).
- Configure Access Control Lists for an Existing Autonomous Database Instance
  You can control and restrict access to your Autonomous Database by specifying network access control lists (ACLs). On an existing Autonomous Database instance with a public endpoint you can add, change, or remove ACLs.
- Change from Private to Public Endpoints with Autonomous Database
  If your Autonomous Database instance is configured to use a private endpoint you can change the configuration to use a public endpoint.
- Access Control List Restrictions and Notes
  Describes restrictions and notes for access control rules on Autonomous Database.

Configure Access Control Lists When You Provision or Clone an Instance

When you provision or clone Autonomous Database with the **Secure access from allowed IPs and VCNs only** option, you can restrict network access by defining Access Control Lists (ACLs).

See Provision an Autonomous Database Instance for information on provisioning your Autonomous Database.

Configure ACLs as follows:

1. In the Choose network access area, select **Secure access from allowed IPs and VCNs only**.

   With **Secure access from allowed IPs and VCNs only** selected, the console shows the fields and options to specify ACLs:

2. In the Choose network access area, specify access control rules by selecting an **IP notation type** and entering **Values** appropriate for the type you select:

   - **IP address**:
     In **Values** field enter values for the **IP address**. An IP address specified in a network ACL entry is the public IP address of the client that is visible on the public internet that you want to grant access. For example, for an Oracle Cloud Infrastructure VM, this is the IP address shown in the **Public IP** field on the Oracle Cloud Infrastructure console for that VM.

     > **Note:**
     >
     > Optionally select **Add my IP address** to add your current IP address to the ACL entry.

   - **CIDR block**:
     In **Values** field enter values for the **CIDR block**. The CIDR block specified is the public CIDR block of the clients that are visible on the public internet that you want to grant access.
   - **Virtual cloud network**:

Use this option when the network route from the client to the database is going through an Oracle Cloud Infrastructure Service Gateway. See Access to Oracle Services: Service Gateway for more information.

Use this option to specify the VCN for use with an Oracle Cloud Infrastructure Service Gateway:

- In **Virtual cloud network** field select the VCN that you want to grant access from. If you do not have the privileges to see the VCNs in your tenancy this list is empty. In this case use the selection **Virtual cloud network (OCID)** to specify the OCID of the VCN.
- Optionally, in the **IP addresses or CIDRs** field enter private IP addresses or private CIDR blocks as a comma separated list to allow specific clients in the VCN.

- **Virtual cloud network (OCID)**:
  Use this option when the network route from the client to the database is going through an Oracle Cloud Infrastructure Service Gateway. See Access to Oracle Services: Service Gateway for more information.
  - In the **Values** field enter the OCID of the VCN you want to grant access from.
  - Optionally, in the **IP addresses or CIDRs** field enter private IP addresses or private CIDR blocks as a comma separated list to allow specific clients in the VCN.

If you want to specify multiple IP addresses or CIDR ranges within the same VCN, do not create multiple ACL entries. Use one ACL entry with the values for the multiple IP addresses or CIDR ranges separated by commas.

3. Click **Add access control rule** to add a new value to the access control list.
4. Click **x** to remove an entry.

   You can also clear the value in the **IP addresses** or **CIDR blocks** field to remove an entry.
5. Require mutual TLS (mTLS) authentication.

   After you enter an IP notation type and a value, you have the option to select this option. The options are:

   - When **Require mutual TLS (mTLS) authentication** is selected, only mTLS connections are allowed (TLS authentication is not allowed).
   - When **Require mutual TLS (mTLS) authentication** is deselected, TLS and mTLS connections are allowed. This is the default configuration.

   See Update Network Options to Allow TLS or Require Only Mutual TLS (mTLS) Authentication on Autonomous Database for more information.
6. Complete the remaining provisioning or cloning steps, as specified in Provision an Autonomous Database Instance, Clone an Autonomous Database Instance, or Clone an Autonomous Database from a Backup.

After provisioning completes, you can update public endpoint ACLs or you can change the Autonomous Database configuration to use a private endpoint.

See Configure Access Control Lists for an Existing Autonomous Database Instance for information on updating ACLs.

See Change from Public to Private Endpoints with Autonomous Database for information on changing to a private endpoint.

Configure Access Control Lists for an Existing Autonomous Database Instance

You can control and restrict access to your Autonomous Database by specifying network access control lists (ACLs). On an existing Autonomous Database instance with a public endpoint you can add, change, or remove ACLs.

Configure ACLs, or add, remove, or update existing ACLs for an Autonomous Database instance as follows:

1. On the **Details** page, in the Network area, next to the **Access control list field**, click **Edit**.

   This shows the Update network access pane.

   As an alternative you can click **More actions** and select select **Update network access**, and in the pane, under **Access type**, select **Secure access from allowed IPs and VCNs only**.

2. Specify the access control rules by selecting an **IP notation type** and values:

   Select one of:

   - **IP address**:
     In **Values** field enter values for the **IP address**. An IP address specified in a network ACL entry is the public IP address of the client that is visible on the public internet that you want to grant access. For example, for an Oracle Cloud Infrastructure VM, this is the IP address shown in the **Public IP** field on the Oracle Cloud Infrastructure console for that VM.

     > **Note:**
     >
     > Optionally select **Add my IP address** to add your current IP address to the ACL entry.

   - **CIDR block**:
     In **Values** field enter values for the **CIDR block**. The CIDR block specified is the public CIDR block of the clients that are visible on the public internet that you want to grant access.

   - **Virtual cloud network**:
     Use this option when the network route from the client to the database is going through an Oracle Cloud Infrastructure Service Gateway. See Access to Oracle Services: Service Gateway for more information.

     Use this option to specify the VCN for use with an Oracle Cloud Infrastructure Service Gateway:

     – In **Virtual cloud network** field select the VCN that you want to grant access from. If you do not have the privileges to see the VCNs in your tenancy this list is empty. In this case use the selection **Virtual cloud network (OCID)** to specify the OCID of the VCN.

     – Optionally, in the **IP addresses or CIDRs** field enter private IP addresses or private CIDR blocks as a comma separated list to allow specific clients in the VCN.

   - **Virtual cloud network (OCID)**:
     Use this option when the network route from the client to the database is going through an Oracle Cloud Infrastructure Service Gateway. See Access to Oracle Services: Service Gateway for more information.

     – In the **Values** field enter the OCID of the VCN you want to grant access from.

– Optionally, in the **IP addresses or CIDRs** field enter private IP addresses or private CIDR blocks as a comma separated list to allow specific clients in the VCN.

If you want to specify multiple IP addresses or CIDR ranges within the same VCN, do not create multiple ACL entries. Use one ACL entry with the values for the multiple IP addresses or CIDR ranges separated by commas.

3. Click **Add access control** to add a new value to the access control list.
4. Click **x** to remove an entry.

   You can also clear the value in the **IP addresses** or **CIDR blocks** field to remove an entry.
5. Click **Update**.

If the Lifecycle state is **Available** when you click **Update** the Lifecycle state changes to **Updating** until the ACL is set. The database is still up and accessible, there is no downtime. When the update is complete the Lifecycle state returns to **Available** and the network ACLs from the access control list are in effect.

Change from Private to Public Endpoints with Autonomous Database

If your Autonomous Database instance is configured to use a private endpoint you can change the configuration to use a public endpoint.

There are several prerequisites to change an instance from a private to a public endpoint, as follows:

- The Autonomous Database instance must be in the available state (Lifecycle state: **Available**).
- Before changing the network configuration from a private endpoint to a public endpoint, you must change the configuration to not allow TLS connections. This closes any existing TLS connections. See Update your Autonomous Database Instance to Require mTLS and Disallow TLS Authentication for more information.

To specify a public endpoint for your Autonomous Database do the following:

1. On the **Details** page, from the **More actions** drop-down list, select **Update network access**.
2. In the **Update network access** dialog, select one of **Secure access from everywhere** or **Secure access from allowed IPs and VCNs only**.

   For example, if you select **Secure access from allowed IPs and VCNs only** the dialog shows fields to configure access control rules:
3. In the dialog, under Configure access control rules specify rules by selecting an **IP notation type** and values:

   - **IP address**:
     In **Values** field enter values for the **IP address**. An IP address specified in a network ACL entry is the public IP address of the client that is visible on the public internet that you want to grant access. For example, for an Oracle Cloud Infrastructure VM, this is the IP address shown in the **Public IP** field on the Oracle Cloud Infrastructure console for that VM.

     > **Note:**
     >
     > Optionally select **Add my IP address** to add your current IP address to the ACL entry.

- **CIDR block**:
  In **Values** field enter values for the **CIDR block**. The CIDR block specified is the public CIDR block of the clients that are visible on the public internet that you want to grant access.
- **Virtual cloud network**:
  Use this option when the network route from the client to the database is going through an Oracle Cloud Infrastructure Service Gateway. See Access to Oracle Services: Service Gateway for more information.
  - In **Virtual cloud network** field select the VCN that you want to grant access from. If you do not have the privileges to see the VCNs in your tenancy this list is empty. In this case use the selection **Virtual cloud network (OCID)** to specify the OCID of the VCN.
  - Optionally, in the **IP addresses or CIDRs** field enter private IP addresses or private CIDR blocks as a comma separated list to allow specific clients in the VCN.
- **Virtual cloud network (OCID)**:
  Use this option when the network route from the client to the database is going through an Oracle Cloud Infrastructure Service Gateway. See Access to Oracle Services: Service Gateway for more information.
  - In the **Values** field enter the OCID of the VCN you want to grant access from.
  - Optionally, in the **IP addresses or CIDRs** field enter private IP addresses or private CIDR blocks as a comma separated list to allow specific clients in the VCN.

If you want to specify multiple IP addresses or CIDR ranges within the same VCN, do not create multiple ACL entries. Use one ACL entry with the values for the multiple IP addresses or CIDR ranges separated by commas.

4. Click **Add access control rule** to add a new value to the access control list.
5. Click **x** to remove an entry.

   You can also clear the value in the **IP addresses** or **CIDR blocks** field to remove an entry.
6. Click **Update**.
7. In the **Confirm** dialog, type the Autonomous Database name to confirm the change.
8. In the **Confirm** dialog, click **Update**.

The Lifecycle state changes to **Updating** until the operation completes.

Notes for changing from private endpoint to public endpoint network access:

- After updating the network access type all database users must obtain a new wallet and use the new wallet to access the database. See Download Client Credentials (Wallets) for more information.
- After the update completes, you can change or define new access control rules ACLs for the public endpoint. See Configure Access Control Lists for an Existing Autonomous Database Instance for more information.
- The URL for Database Actions and for the Database Tools are different when a database uses a private endpoint compared to using a public endpoint. Click Database Actions on the Oracle Cloud Infrastructure Console to find the updated Database Actions URL and in Database Actions click the appropriate cards to find the updated Database Tools URLs, after changing from a private endpoint to a public endpoint.

Access Control List Restrictions and Notes

Describes restrictions and notes for access control rules on Autonomous Database.

- See IP Address Ranges for information about the public IP address ranges in Oracle Cloud Infrastructure. You must allow traffic to these CIDR blocks to ensure access to an Autonomous Database instance on a public endpoint.
- If you want to only allow connections coming through a service gateway you need to use the IP address of the service gateway in your ACL definition. To do this you need to add an ACL definition with the CIDR source type with the value `240.0.0.0/4`. Note that this is not recommended, instead of this you can specify individual VCNs in your ACL definition for the VCNs you want to allow access from.

  See Access to Oracle Services: Service Gateway for more information.
- When you restore a database the existing ACLs are not overwritten by the restore.
- The network ACLs apply to the database connections and Oracle Machine Learning notebooks. If an ACL is defined, if you try to login to Oracle Machine Learning Notebooks from a client whose IP is not specified on the ACL this shows the "login rejected based on access control list set by the administrator" error.
- The following Autonomous Database tools are subject to ACLs. You can use Virtual Cloud Network, Virtual Cloud Network (OCID), IP address, or CIDR block ACLs to control access to these tools:
  - Database Actions
  - Oracle APEX
  - Oracle Graph Studio
  - Oracle Machine Learning Notebooks
  - Oracle REST Data Services
- If you have a private subnet in your VCN that is configured to access the public internet through a NAT Gateway, you need to enter the public IP address of the NAT Gateway in your ACL definition. Clients in the private subnet do not have public IP addresses. See NAT Gateway for more information.
- If you are using ACLs and TLS connections are allowed, you must change your network configuration to not allow TLS connections before removing all ACLs. See Update your Autonomous Database Instance to Require mTLS and Disallow TLS Authentication for more information.
- To view the network information for your instance, see View Network Information on the OCI Console.

Configure Network Access with Private Endpoints

You can specify that Autonomous Database uses a private endpoint inside your Virtual Cloud Network (VCN) in your tenancy. You can configure a private endpoint during provisioning or cloning your Autonomous Database, or you can switch to using a private endpoint in an existing database that uses a public endpoint. This allows you to keep all traffic to and from your database off of the public internet.

Specifying the virtual cloud network configuration allows traffic only from the virtual cloud network you specify and blocks access to the database from all public IPs or VCNs. This allows you to define security rules with Security Lists or at the Network Security Group (NSG) level to specify ingress/egress for your Autonomous Database instance. Using a private endpoint and defining Security Lists or NSGs allows you to control traffic to and from your Autonomous Database instance.

> **Note:**
>
> If you configure your Autonomous Database instance to use a private endpoint and you also want to allow connections from specific public IP addresses or from specific VCNs if those VCNs are configured to privately connect to Autonomous Database using a Service Gateway, select the **Allow public access** option. This adds a public endpoint for a database that is configured with a private endpoint. See Use a Private Endpoint with Public Access Allowed for more information.
> The **Allow public access option** is available only when the database uses the ECPU compute model.

Topics

- Configure Private Endpoints
  You can specify that Autonomous Database uses a private endpoint and configure a Virtual Cloud Network (VCN) in your tenancy to use with the private endpoint.
- Enhanced Security for Outbound Connections with Private Endpoints
  When you use a private endpoint with your Autonomous Database instance you can provide enhanced security by setting the `ROUTE_OUTBOUND_CONNECTIONS` database property to the value `PRIVATE_ENDPOINT`.
- Private Endpoints Notes
  Describes restrictions and notes for private endpoints on Autonomous Database.
- Private Endpoints Configuration Examples on Autonomous Database
  Shows several Private Endpoint (VCN) configuration samples for Autonomous Database.

Configure Private Endpoints

You can specify that Autonomous Database uses a private endpoint and configure a Virtual Cloud Network (VCN) in your tenancy to use with the private endpoint.

- Prerequisite Steps to Configure Private Endpoints
  Describes the prerequisite steps you need to perform before you configure a private endpoint for an Autonomous Database instance.
- IAM Policies Required to Manage Private Endpoints
  In addition to the policies required to provision and manage an Autonomous Database, some network policies are needed to use private endpoints.
- Configure Private Endpoints When You Provision or Clone an Instance
  You can configure a private endpoint when you provision or clone an Autonomous Database instance.
- Change from Public to Private Endpoints with Autonomous Database
  If your Autonomous Database instance is configured to use a public endpoint you can change the configuration to a private endpoint.
- Update the Configuration for a Private Endpoint
  You can change some options in the configuration of a private endpoint on an existing Autonomous Database instance.
- Configure Private Endpoint Advanced Options
  The private endpoint access advanced options allow you to enter a user specified private IP address and host name, select one or more network security groups, or specify details to allow public access to a private endpoint database.

- Use a Private Endpoint with Public Access Allowed
  Select the **Allow public access** option when you want to configure an Autonomous Database to use a private endpoint and you also want to allow connections from specific public IP addresses or from specific VCNs (if the VCNs are configured to privately connect to Autonomous Database using a Service Gateway).

Prerequisite Steps to Configure Private Endpoints

Describes the prerequisite steps you need to perform before you configure a private endpoint for an Autonomous Database instance.

Perform the following prerequisite steps before configuring a private endpoint:

- Set required policies for the resources you are working with. See IAM Policies Required to Manage Private Endpoints for more information.
- Create a VCN within the region that will contain your Autonomous Database. See VCNs and Subnets for more information.
- Configure a subnet within your VCN configured with default DHCP options. See DNS in Your Virtual Cloud Network for more information.
- (Optional) Perform the following optional step before configuring a private endpoint:

  Specify a Network Security Group (NSG) within your VCN. The NSG specifies rules for connections to your Autonomous Database. See Network Security Groups for more information.

IAM Policies Required to Manage Private Endpoints

In addition to the policies required to provision and manage an Autonomous Database, some network policies are needed to use private endpoints.

The following table lists the IAM policies required for a cloud user to add a private endpoint. The listed policies are the minimum requirements to add a private endpoint. You can also use a policy rule that is broader. For example, if you set the policy rule:

```
Allow group MyGroupName to manage virtual-network-family in tenancy
```

This rule also works because it is a superset that contains all the required policies.

Operation	Required IAM Policies
Configure a private endpoint	`use vcns` for the compartment which the VCN is in `use subnets` for the compartment which the VCN is in `use network-security-groups` for the compartment which the network security group is in `manage private-ips` for the compartment which the VCN is in `manage vnics` for the compartment which the VCN is in `manage vnics` for the compartment which the database is provisioned or is to be provisioned in

Autonomous Database relies on the IAM (Identity and Access Management) service to authenticate and authorize cloud users to perform operations that use any of the Oracle Cloud Infrastructure interfaces (the Console, REST API, CLI, SDK, or others).

The IAM service uses **groups**, **compartments** and **policies** to control which cloud users can access which resources. In particular, a policy defines what kind of access a group of users

has to a particular kind of resource in a particular compartment. For more information, see Getting Started with Policies.

Configure Private Endpoints When You Provision or Clone an Instance

You can configure a private endpoint when you provision or clone an Autonomous Database instance.

These steps assume you are provisioning or cloning an instance and you have completed the prerequisite steps, and you are at the **Choose network access** step of the provisioning or cloning steps:

1. Select **Private endpoint access only**.

   This expands the Virtual cloud network private access configuration area.

   If you select **Private endpoint access only**, this only allows connections from the specified private network (VCN), from peered VCNs, and from on-prem networks connected to your VCN. You can configure an Autonomous Database instance on a private endpoint to allow connections from on-prem networks. See Example: Connecting from Your Data Center to Autonomous Database for an example.

   If you want to allow connections from public IP addresses or from allowed IPs and VCNs, you have the following options:

   - Select **Secure access from everywhere**.
   - Select **Secure access from allowed IPs and VCNs only**.
   - If you select **Private endpoint access only**, expand **Show advanced options** and select **Allow public access**. See Configure Private Endpoint Advanced Options for more information.

2. Select a **Virtual cloud network** in your compartment or if the VCN is in a different compartment click **Change Compartment** and select the compartment that contains the VCN and then select a virtual cloud network.

   See VCNs and Subnets for more information.

3. Select the **Subnet** in your compartment to attach the Autonomous Database to or if the Subnet is in a different compartment click **Change Compartment** and select the compartment that contains the Subnet and then select a subnet.

   See VCNs and Subnets for more information.

4. (Optional) Click **Show advanced options** to show additional private endpoint options.

   See Configure Private Endpoint Advanced Options for details on the advanced options.

5. Require mutual TLS (mTLS) authentication.

   The **Require mutual TLS (mTLS) authentication** options are:

   - When **Require mutual TLS (mTLS) authentication** is deselected, TLS and mTLS connections are allowed. This is the default configuration.
   - When **Require mutual TLS (mTLS) authentication** is selected, only mTLS connections are allowed (TLS authentication is not allowed).

   See Update Network Options to Allow TLS or Require Only Mutual TLS (mTLS) Authentication on Autonomous Database for more information.

6. Complete the remaining provisioning or cloning steps, as specified in Provision an Autonomous Database Instance, Clone an Autonomous Database Instance, or Clone an Autonomous Database from a Backup.

See Private Endpoints Notes for more information.

Change from Public to Private Endpoints with Autonomous Database

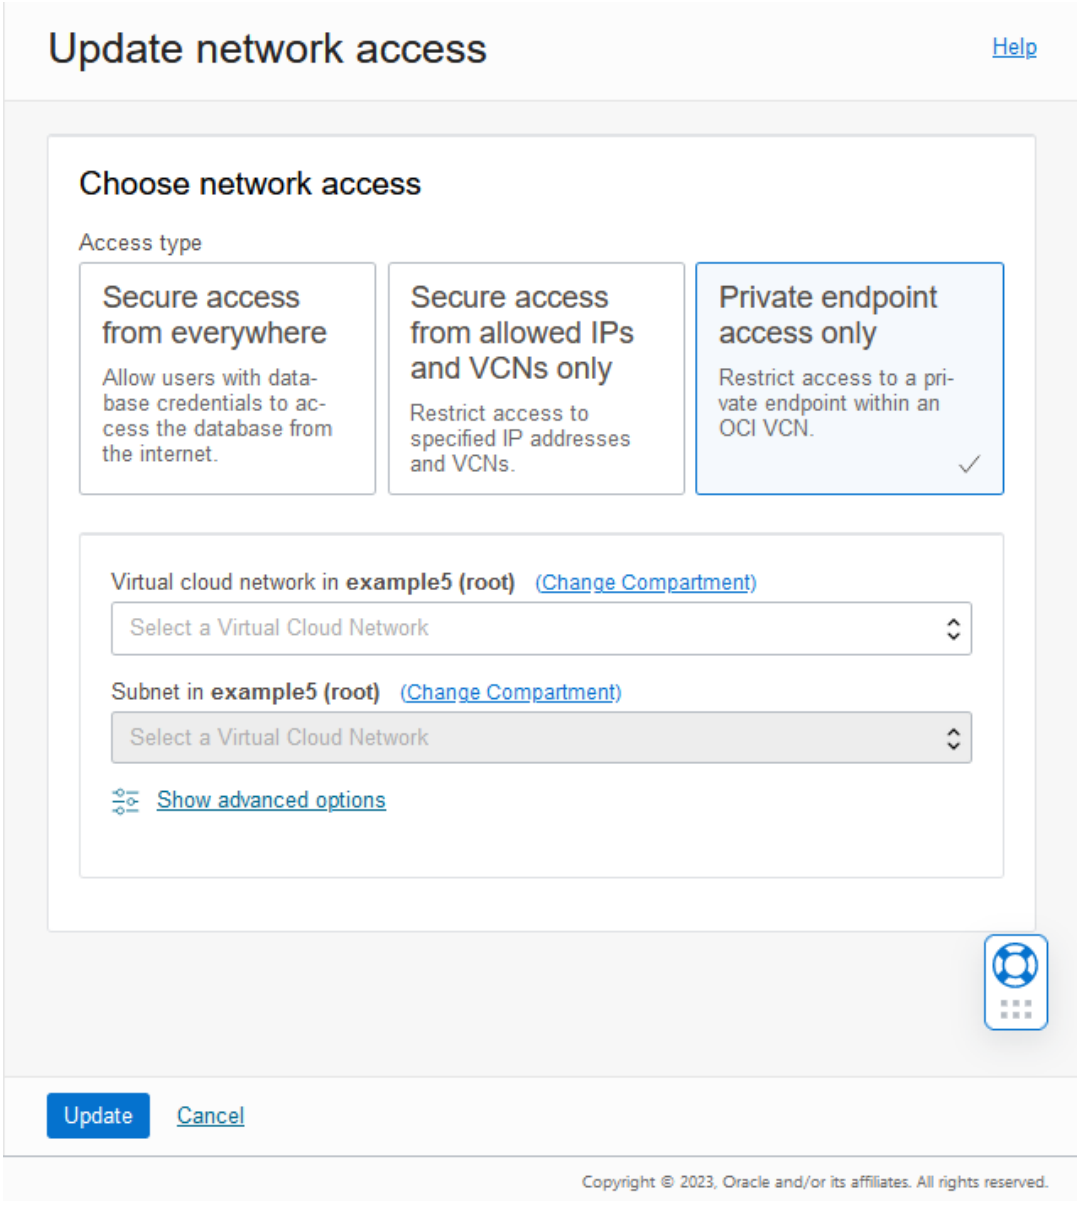
If your Autonomous Database instance is configured to use a public endpoint you can change the configuration to a private endpoint.

1. On the **Details** page, from the **More actions** drop-down list, select **Update network access**.

   To change an instance from a public to a private endpoint, the Autonomous Database instance must be in the **Available** state (Lifecycle State: **Available**).

2. In the **Update network access** dialog, select **Private endpoint access only**.

   This expands the Virtual cloud network private access configuration area.



If you select **Private endpoint access only**, this only allows connections from the specified private network (VCN), from peered VCNs, and from on-prem networks connected to your VCN. You can configure an Autonomous Database instance on a private

endpoint to allow connections from on-prem networks. See Example: Connecting from Your Data Center to Autonomous Database for an example.

If you want to allow connections from public IP addresses or from allowed IPs and VCNs, you have the following options:

- Select **Secure access from everywhere**.
- Select **Secure access from allowed IPs and VCNs only**.
- If you select **Private endpoint access only**, expand **Show advanced options** and select **Allow public access**. See Configure Private Endpoint Advanced Options for more information.

3. Select a **Virtual cloud network** in your compartment or if the VCN is in a different compartment click **Change Compartment** and select the compartment that contains the VCN and then select a virtual cloud network.

   See VCNs and Subnets for more information.

4. Select the **Subnet** in your compartment to attach the Autonomous Database to or if the Subnet is in a different compartment click **Change Compartment** and select the compartment that contains the Subnet and then select a subnet.

   See VCNs and Subnets for more information.

5. (Optional) Click **Show advanced options** to see additional options.

   See Configure Private Endpoint Advanced Options for details on the advanced options.

6. Click **Update**.
7. In the **Confirm** dialog, type the Autonomous Database name to confirm the change.
8. In the **Confirm** dialog, click **Update**.

The Lifecycle State changes to **Updating** until the operation completes.

Notes for changing from public to private network access:

- After updating the network access type all database users must obtain a new wallet and use the new wallet to access the database. See Download Client Credentials (Wallets) for more information.
- If you had ACLs defined for the public endpoint, the ACLs do not apply for the private endpoint.
- After you update the network access to use a private endpoint, the URL for the Database Tools is different compared to using a public endpoint. You can find the updated URLs on the console, after changing from a public endpoint to a private endpoint.

Update the Configuration for a Private Endpoint

You can change some options in the configuration of a private endpoint on an existing Autonomous Database instance.

1. On the **Details** page, from the **More actions** drop-down list, select **Update network access**.

   This shows the Update network access panel.

2. Select **Private endpoint access only**.

   If you want to allow connections from public IP addresses or from allowed IPs and VCNs, you have the following options:

   - Select **Secure access from everywhere**.

- Select **Secure access from allowed IPs and VCNs only**.
- When you select **Private endpoint access only** show advanced options and select **Allow public access**. This defines a private endpoint database that includes both a private endpoint and a public endpoint.

a. Optionally add **Network security groups (NSGs)**.

Optionally, to allow connections to the Autonomous Database instance define security rules in an NSG; this creates a virtual firewall for your Autonomous Database.

- Select a Network Security Group in your compartment to attach the Autonomous Database to, or if the Network Security Group is in a different compartment, click **Change Compartment** and select a different compartment and then select a Network Security Group in that compartment.
- Click **+ Another Network Security Group** to add another Network Security Group.
- Click **x** to remove a Network Security Group entry.

For the NSG you select for the private endpoint define a security rule as follows:

- For mutual TLS (mTLS) authentication, add a stateful ingress rule with the source set to the address range you want to allow to connect to your database, the IP Protocol set to TCP, and the Destination Port Range set to 1522.
- For TLS authentication, add a stateful ingress rule with the source set to the address range you want to allow to connect to your database, the IP Protocol set to TCP, and the Destination Port Range set to 1521.
- To use Oracle APEX, Database Actions, and Oracle REST Data Services, add port 443 to the NSG rule.

> **Note:**
>
> Incoming and outgoing connections are limited by the combination of ingress and egress rules defined in NSGs and the Security Lists defined with the VCN. When there are no NSGs, ingress and egress rules defined in the Security Lists for the VCN still apply. See Security Lists for more information on working with Security Lists.

See Secure Connections to Autonomous Database with mTLS or with TLS for more information.

See Private Endpoints Configuration Examples on Autonomous Database for examples.

See Network Security Groups for more information.

b. Optionally, select **Allow public access** or if this is already selected, you can configure access control rules to the public endpoint that is configured with the private endpoint database.

The **Allow public access option** is available only when the database uses the ECPU compute model.

When you select **Allow public access**, this shows the Configure access control options to enter the allowed IP addresses, CIDR blocks, or Virtual cloud networks that can connect to the database.

Select one of:

- **IP address**:
  In **Values** field enter values for the **IP address**. An IP address specified in a network ACL entry is the public IP address of the client that is visible on the public internet that you want to grant access. For example, for an Oracle Cloud Infrastructure VM, this is the IP address shown in the **Public IP** field on the Oracle Cloud Infrastructure console for that VM.

  Optionally select **Add my IP address** to add your current IP address to the ACL entry.
- **CIDR block**:
  In **Values** field enter values for the **CIDR block**. The CIDR block specified is the public CIDR block of the clients that are visible on the public internet that you want to grant access.
- **Virtual cloud network**:
  Use this option when the network route from the client to the database is going through an Oracle Cloud Infrastructure Service Gateway. See Access to Oracle Services: Service Gateway for more information.

  Use this option to specify the VCN for use with an Oracle Cloud Infrastructure Service Gateway:

  – In **Virtual cloud network** field select the VCN that you want to grant access from. If you do not have the privileges to see the VCNs in your tenancy this list is empty. In this case use the selection **Virtual cloud network (OCID)** to specify the OCID of the VCN.

  – Optionally, in the **IP addresses or CIDRs** field enter private IP addresses or private CIDR blocks as a comma separated list to allow specific clients in the VCN.
- **Virtual cloud network (OCID)**:
  Use this option when the network route from the client to the database is going through an Oracle Cloud Infrastructure Service Gateway. See Access to Oracle Services: Service Gateway for more information.

  – In the **Values** field enter the OCID of the VCN you want to grant access from.

  – Optionally, in the **IP addresses or CIDRs** field enter private IP addresses or private CIDR blocks as a comma separated list to allow specific clients in the VCN.

If you want to specify multiple IP addresses or CIDR ranges within the same VCN, do not create multiple ACL entries. Use one ACL entry with the values for the multiple IP addresses or CIDR ranges separated by commas.

3. Click **Update**.

If the Lifecycle state is **Available** when you click **Update**, the Lifecycle state changes to **Updating** until the changes are applied. The database is still up and accessible, there is no downtime. When the update is complete the Lifecycle state returns to **Available**.

See Private Endpoints Notes for more information.

Configure Private Endpoint Advanced Options

The private endpoint access advanced options allow you to enter a user specified private IP address and host name, select one or more network security groups, or specify details to allow public access to a private endpoint database.

These steps assume you are provisioning or cloning an Autonomous Database instance or changing from public access to private access for an existing Autonomous Database instance and you are at the **Choose network access** step.

1. Select **Private endpoint access only**.

   This shows the Virtual cloud network private access configuration area.

2. (Optional) Click **Show advanced options** to show additional private endpoint options.

   This displays the advanced options.

   a. Optionally enter a **Private IP address**.

      Use this field to enter a custom private IP address. The private IP address you enter must be within the selected subnet's CIDR range.

      If you do not provide a custom private IP address the IP address is automatically assigned.

   b. Optionally enter a **Hostname prefix**.

      This specifies a hostname prefix for the Autonomous Database and associates a DNS name with the database instance, in the following form:

      ```
      hostname_prefix.adb.region.oraclecloud.com
      ```

      If you do not specify a hostname prefix, a system generated hostname prefix is supplied.

   c. Optionally add **Network security groups (NSGs)**.

      Optionally, to allow connections to the Autonomous Database instance define security rules in an NSG; this creates a virtual firewall for your Autonomous Database.

      - Select a Network Security Group in your compartment to attach the Autonomous Database to, or if the Network Security Group is in a different compartment, click **Change Compartment** and select a different compartment and then select a Network Security Group in that compartment.
      - Click **+ Another Network Security Group** to add another Network Security Group.
      - Click **x** to remove a Network Security Group entry.

      For the NSG you select for the private endpoint define a security rule as follows:

      - For mutual TLS (mTLS) authentication, add a stateful ingress rule with the source set to the address range you want to allow to connect to your database, the IP Protocol set to TCP, and the Destination Port Range set to 1522.
      - For TLS authentication, add a stateful ingress rule with the source set to the address range you want to allow to connect to your database, the IP Protocol set to TCP, and the Destination Port Range set to 1521.
      - To use Oracle APEX, Database Actions, and Oracle REST Data Services, add port 443 to the NSG rule.

> **Note:**
>
> Incoming and outgoing connections are limited by the combination of ingress and egress rules defined in NSGs and the Security Lists defined with the VCN. When there are no NSGs, ingress and egress rules defined in the Security Lists for the VCN still apply. See Security Lists for more information on working with Security Lists.

See Secure Connections to Autonomous Database with mTLS or with TLS for more information.

See Private Endpoints Configuration Examples on Autonomous Database for examples.

See Network Security Groups for more information.

d. Optionally, select **Allow public access** and configure access control rules to add a public endpoint for the private endpoint database.

   The **Allow public access option** is available only when the database uses the ECPU compute model.

   When you select **Allow public access**, this shows the Configure access control options to enter the allowed IP addresses, CIDR blocks, or Virtual cloud networks that can connect to the database.

   Select one of:

   - **IP address**:
     In **Values** field enter values for the **IP address**. An IP address specified in a network ACL entry is the public IP address of the client that is visible on the public internet that you want to grant access. For example, for an Oracle Cloud Infrastructure VM, this is the IP address shown in the **Public IP** field on the Oracle Cloud Infrastructure console for that VM.

     Optionally select **Add my IP address** to add your current IP address to the ACL entry.

     Optionally select **Add my IP address** to add your current IP address to the ACL entry.

   - **CIDR block**:
     In **Values** field enter values for the **CIDR block**. The CIDR block specified is the public CIDR block of the clients that are visible on the public internet that you want to grant access.

   - **Virtual cloud network**:
     Use this option when the network route from the client to the database is going through an Oracle Cloud Infrastructure Service Gateway. See Access to Oracle Services: Service Gateway for more information.

     Use this option to specify the VCN for use with an Oracle Cloud Infrastructure Service Gateway:

     – In **Virtual cloud network** field select the VCN that you want to grant access from. If you do not have the privileges to see the VCNs in your tenancy this list is empty. In this case use the selection **Virtual cloud network (OCID)** to specify the OCID of the VCN.

  - Optionally, in the **IP addresses or CIDRs** field enter private IP addresses or private CIDR blocks as a comma separated list to allow specific clients in the VCN.

- **Virtual cloud network (OCID)**:
  Use this option when the network route from the client to the database is going through an Oracle Cloud Infrastructure Service Gateway. See Access to Oracle Services: Service Gateway for more information.

  - In the **Values** field enter the OCID of the VCN you want to grant access from.
  - Optionally, in the **IP addresses or CIDRs** field enter private IP addresses or private CIDR blocks as a comma separated list to allow specific clients in the VCN.

  If you want to specify multiple IP addresses or CIDR ranges within the same VCN, do not create multiple ACL entries. Use one ACL entry with the values for the multiple IP addresses or CIDR ranges separated by commas.

3. Complete the remaining private endpoint configuration steps.

Use a Private Endpoint with Public Access Allowed

Select the **Allow public access** option when you want to configure an Autonomous Database to use a private endpoint and you also want to allow connections from specific public IP addresses or from specific VCNs (if the VCNs are configured to privately connect to Autonomous Database using a Service Gateway).

This option adds a public endpoint for a database that is configured on a private endpoint. You configure a private endpoint for your Autonomous Database instance when you provision or clone the instance, or when you update the network configuration for an existing Autonomous Database. See the following for details on the steps to configure an Autonomous Database instance with a private endpoint:

- Configure Private Endpoints When You Provision or Clone an Instance
- Change from Public to Private Endpoints with Autonomous Database

When public access is enabled with **Allow public access** on a private endpoint database, the instance has both a private endpoint and a public endpoint:

- The private hostname, endpoint URL, and private IP address allow you to connect to the database from the VCN where the database resides.
- The public hostname allows you to connect to the database from specific public IP addresses or from specific VCNs if those VCNs are configured to privately connect to Autonomous Database using a Service Gateway.

Autonomous Database Connection String Additions for a Private Endpoint Database with Allow Public Access Enabled

When **Allow public access** is enabled for a private endpoint database, there are additional connection strings that allow you to connect to the database from the public endpoint:

- The connection strings in `tnsnames.ora` in the Autonomous Database wallet zip include the public connection strings to use with connections coming from the public internet. The connection strings for the public endpoint use the following naming convention:

  dbname`_public_`*consumerGroup*

For example:

```
adbfinance_public_low
```

See Download Client Credentials (Wallets) for more information.

- You can view the connection strings for both the public endpoint and the private endpoint from the Oracle Cloud Infrastructure Console (or using the API).

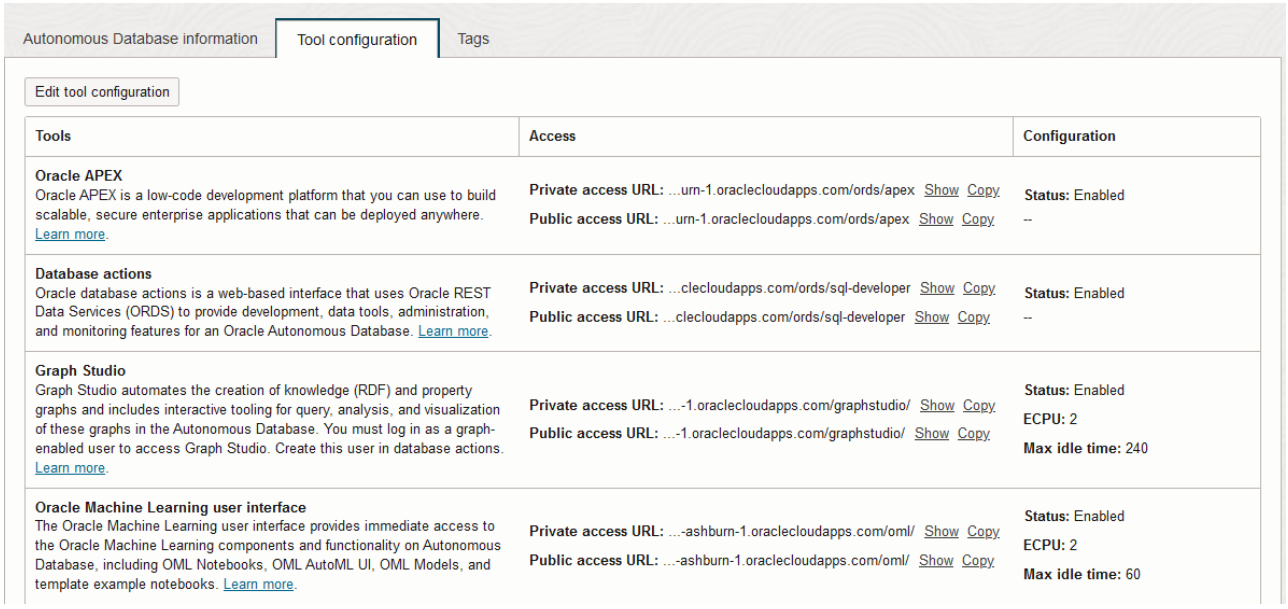
  See View TNS Names and Connection Strings for an Autonomous Database Instance for more information.

Autonomous Database Tools Additions for a Private Endpoint Database with Allow Public Access Enabled

When **Allow public access** is enabled for a private endpoint database, the database tools allow you to connect from specific public IP addresses or from specific VCNs if those VCNs are configured to privately connect to Autonomous Database using a Service Gateway:

- Each tool has a private access URL and a public access URL displayed in the Tool configuration table. Use the public access URL to access the tool from specific public IP addresses or from specific VCNs if those VCNs are configured to privately connect to Autonomous Database using a Service Gateway.

  For example:



See View Autonomous Database Built-in Tools Status for more information.

- The `README` file in the wallet zip file provides both an Access Link for the private endpoint for each database tool and a Public Access Link.

  See Wallet README File for more information.

Enhanced Security for Outbound Connections with Private Endpoints

When you use a private endpoint with your Autonomous Database instance you can provide enhanced security by setting the `ROUTE_OUTBOUND_CONNECTIONS` database property to the value `PRIVATE_ENDPOINT`.

Setting the `ROUTE_OUTBOUND_CONNECTIONS` database property to the value `PRIVATE_ENDPOINT` enforces that all outgoing connections to a target host are subject to and limited by the private endpoint's egress rules. You define egress rules in the Virtual Cloud Network (VCN) security list or in the Network Security Group (NSG) associated with the Autonomous Database instance private endpoint.

Before you set the `ROUTE_OUTBOUND_CONNECTIONS` database property, configure your Autonomous Database instance to use a private endpoint. See Configure Private Endpoints for more information.

Set the `ROUTE_OUTBOUND_CONNECTIONS` database property to `PRIVATE_ENDPOINT` to specify that all outgoing connections are subject to the Autonomous Database instance private endpoint VCN's egress rules. With the value `PRIVATE_ENDPOINT` the database restricts outgoing connections to locations specified by the private endpoint's egress rules and also changes DNS resolution such that hostnames are resolved using your VCN's DNS resolver (not using a public DNS resolver).

> **Note:**
>
> With `ROUTE_OUTBOUND_CONNECTIONS` not set to `PRIVATE_ENDPOINT`, all outgoing connections to the public internet pass through the Network Address Translation (NAT) Gateway of the service VCN. In this case, if the target host is on a public endpoint the outgoing connections are not subject to the Autonomous Database instance private endpoint VCN or NSG egress rules.

When you configure a private endpoint for your Autonomous Database instance and set `ROUTE_OUTBOUND_CONNECTIONS` to `PRIVATE_ENDPOINT`, this setting changes the handling of outbound connections and DNS resolution for the following:

- Database links
- APEX_LDAP, APEX_MAIL, and APEX_WEB_SERVICE
- UTL_HTTP, UTL_SMTP, and UTL_TCP
- DBMS_LDAP
- CMU with Microsoft Active Directory

  See Use Microsoft Active Directory with Autonomous Database for more information.

When you configure a private endpoint for your Autonomous Database instance and set `ROUTE_OUTBOUND_CONNECTIONS` to `PRIVATE_ENDPOINT`, this setting does not change the handling of outbound connections and DNS resolution for the following:

- Oracle REST Data Services (ORDS)
- Database Actions

To set `ROUTE_OUTBOUND_CONNECTIONS`:

1. Connect to your database.

2. Set the database property `ROUTE_OUTBOUND_CONNECTIONS`.

   For example:

   ```
   ALTER DATABASE PROPERTY SET ROUTE_OUTBOUND_CONNECTIONS =
   'PRIVATE_ENDPOINT';
   ```

Notes for setting `ROUTE_OUTBOUND_CONNECTIONS`:

- Use the following command to restore the default parameter value:

  ```
  ALTER DATABASE PROPERTY SET ROUTE_OUTBOUND_CONNECTIONS = '';
  ```

- Use the following command to query the current parameter value:

  ```
  SELECT * FROM DATABASE_PROPERTIES
          WHERE PROPERTY_NAME = 'ROUTE_OUTBOUND_CONNECTIONS';
  ```

  If the property is not set the query does not return results.

- This property only applies for database links that you create after you set the property to the value `PRIVATE_ENDPOINT`. Thus, database links that you created prior to setting the property continue to use the NAT Gateway of the service VCN and are not subject to the Autonomous Database instance private endpoint's egress rules.
- Only set `ROUTE_OUTBOUND_CONNECTIONS` to the value `PRIVATE_ENDPOINT` when you are using Autonomous Database with a private endpoint.
- When your database is on a private endpoint and you want your outbound connections to be resolved by your VCN, you need to set the `ROUTE_OUTBOUND_CONNECTIONS` parameter to `PRIVATE_ENDPOINT`.

See NAT Gateway for more information on Network Address Translation (NAT) gateway.

Private Endpoints Notes

Describes restrictions and notes for private endpoints on Autonomous Database.

- After you update the network access to use a private endpoint, or after the provisioning or cloning completes where you configure a private endpoint, you can view the network configuration on the Autonomous Database Details page under the **Network** section.

  The **Network** section shows the following information for a private endpoint:

  - **Access type**: Specifies the access type for the Autonomous Database configuration. Private endpoint configurations show the access type: **Virtual cloud network**.
  - **Availability domain**: Specifies the availability domain of your Autonomous Database instance.
  - **Virtual cloud network**: This includes a link for the VCN associated with the private endpoint.
  - **Subnet**: This includes a link for the subnet associated with the private endpoint.
  - **Private endpoint IP**: Shows the private endpoint IP for the private endpoint configuration.
  - **Private endpoint URL**: Shows the private endpoint URL for the private endpoint configuration.

  - **Network security groups**: This field includes links to the NSG(s) configured with the private endpoint.
  - **Public access**: This field indicates whether public access is enabled for the private endpoint. Click the `Edit` link to view or change the allowed ACLs or VCNs.
  - **Public endpoint URL**: This shows when **Allow public access** is enabled on the private endpoint. This is the public endpoint URL that you can use to connect from allowed IPs or VCNs on the public internet.

  For more details on the network information on the Oracle Cloud Infrastructure Console, see View Network Information on the OCI Console.

- After provisioning or cloning completes, you can change the Autonomous Database configuration to use a public endpoint.

  See Change from Private to Public Endpoints with Autonomous Database for information on changing to a public endpoint.

- You can specify up to five NSGs to control access to your Autonomous Database.
- You can change the private endpoint Network Security Group (NSG) for the Autonomous Database.

  To change the NSG for a private endpoint, do the following:

  1. On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.
  2. On the **Autonomous Database Details** page, under **Network** in the **Network Security Groups** field, click **Edit**.

- You can connect your Oracle Analytics Cloud instance to your Autonomous Database that has a private endpoint using the Data Gateway like you do for an on-premises database. See Configure and Register Data Gateway for Data Visualization for more information.
- The following Autonomous Database tools are supported in databases configured with a private endpoint:
  - Database Actions
  - Oracle APEX
  - Oracle Graph Studio
  - Oracle Machine Learning Notebooks
  - Oracle REST Data Services
  - Oracle Database API for MongoDB

  Additional configuration is required to access these Autonomous Database tools from on-premises environments. See Example: Connecting from Your Data Center to Autonomous Database to learn more.

  Accessing Oracle APEX, Database Actions, Oracle Graph Studio, or Oracle REST Data Services using a private endpoint from on-premises environments without completing the additional private endpoint configuration shows the error:

  ```
  404 Not Found
  ```

- After you update the network access to use a private endpoint, the URL for the Database Tools is different compared to using a public endpoint. You can find the updated URLs on the console, after changing from a public endpoint to a private endpoint.
- In addition to the default Oracle REST Data Services (ORDS) preconfigured with Autonomous Database, you can configure an alternative ORDS deployment that provides

more configuration options and that can be used with private endpoints. See About Customer Managed Oracle REST Data Services on Autonomous Database to learn about an alternative ORDS deployment that can be used with private endpoints.

- Modifying a private IP address is not allowed after you provision or clone an instance, whether the IP address is automatically assigned when you enter a value in the **Private IP address** field.

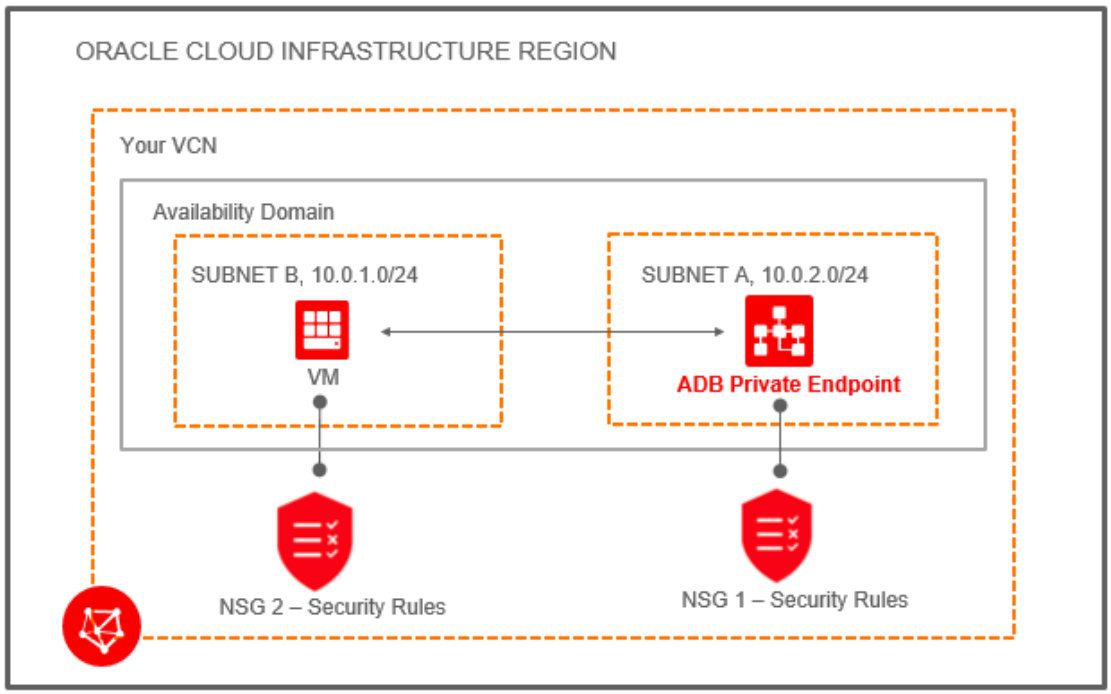
Private Endpoints Configuration Examples on Autonomous Database

Shows several Private Endpoint (VCN) configuration samples for Autonomous Database.

- Example: Connecting from Inside Oracle Cloud Infrastructure VCN
  Demonstrates an application running inside Oracle Cloud Infrastructure on a virtual machine (VM) in the same VCN which is configured with your Autonomous Database.
- Example: Connecting from Your Data Center to Autonomous Database
  Demonstrates how to connect privately to an Autonomous Database from your on-premise data center. In this scenario, traffic never goes over the public internet.

Example: Connecting from Inside Oracle Cloud Infrastructure VCN

Demonstrates an application running inside Oracle Cloud Infrastructure on a virtual machine (VM) in the same VCN which is configured with your Autonomous Database.



There is an Autonomous Database instance which has a private endpoint in the VCN named "Your VCN". The VCN includes two subnets: "SUBNET B" (CIDR 10.0.1.0/24) and "SUBNET A" (CIDR 10.0.2.0/24).

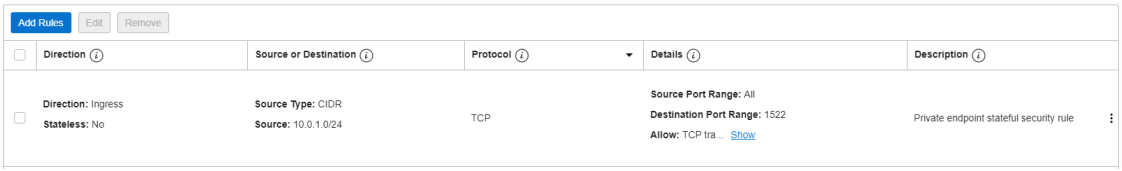
The Network Security Group (NSG) associated with the Autonomous Database instance is shown as "NSG 1 - Security Rules". This Network Security Group defines security rules that allow incoming and outgoing traffic to and from the Autonomous Database instance. Define a rule for the Autonomous Database instance as follows:

- For Mutual TLS authentication, add a stateful ingress rule to allow connections from the source to the Autonomous Database instance; the source is set to the address range you

want to allow to connect to your database, IP Protocol is set to TCP, and the Destination Port Range is set to 1522.

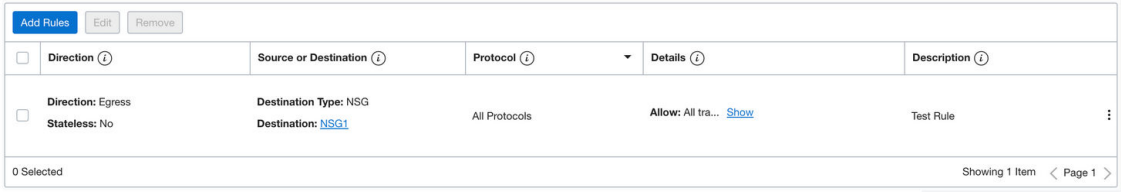
- For TLS authentication, add a stateful ingress rule to allow connections from the source to the Autonomous Database instance; the source is set to the address range you want to allow to connect to your database, IP Protocol is set to TCP, and the Destination Port Range is set to either 1521 or 1522.
- To use Oracle APEX, Database Actions, and Oracle REST Data Services, add port 443 to the NSG rule.

The following figure shows a sample stateful security rule to control traffic for the Autonomous Database instance:



The application connecting to the Autonomous Database is running on a VM in SUBNET B. You also add a security rule to allow traffic to and from the VM (as shown, with label "NSG 2 Security Rules"). You can use a stateful security rule for the VM, so simply add a rule for egress to NSG 2 Security Rules (this allows access to the destination subnet A).

The following figure shows sample security rules that control traffic for the VM:

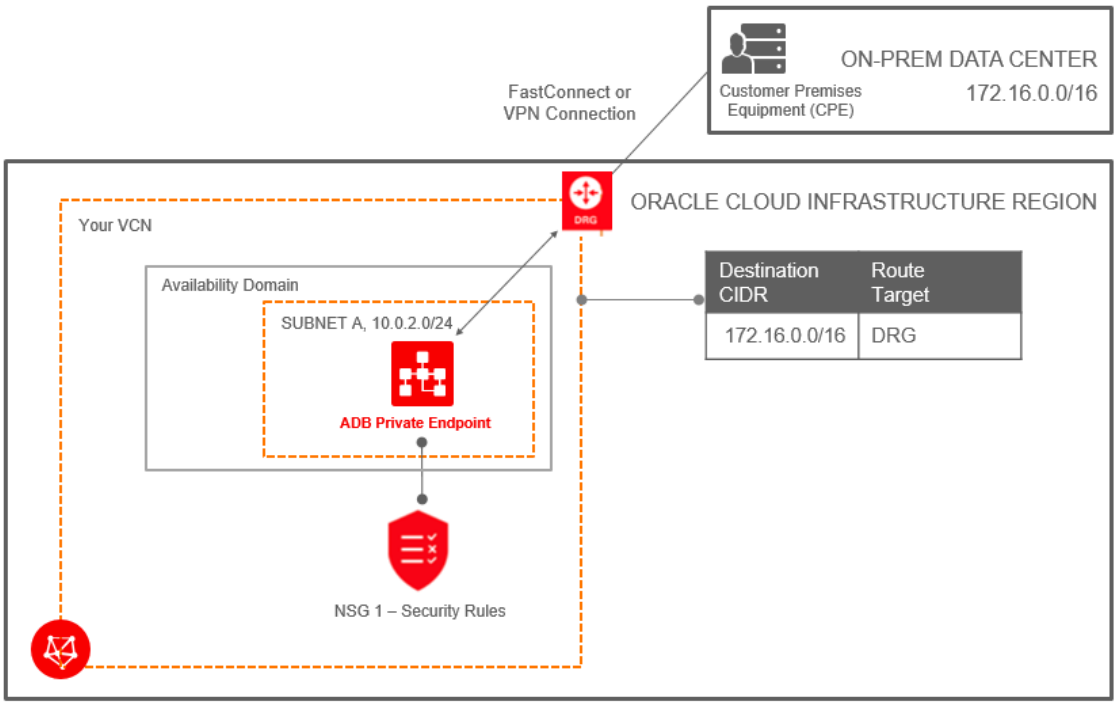


After you configure the security rules, your application can connect to the Autonomous Database instance using the client credentials wallet. See Download Client Credentials (Wallets) for more information.

See Network Security Groups for information on configuring Network Security Groups.

Example: Connecting from Your Data Center to Autonomous Database

Demonstrates how to connect privately to an Autonomous Database from your on-premise data center. In this scenario, traffic never goes over the public internet.



To connect from your data center, you connect the on-premise network to the VCN with FastConnect and then set up a Dynamic Routing Gateway (DRG). To resolve the Autonomous Database private endpoint, a Fully Qualified Domain Name (FQDN), requires that you add an entry in your on-premise client's hosts file. For example, `/etc/hosts` file for Linux machines. For example:

```
/etc/hosts entry -> 10.0.2.7 example.adb.ca-toronto-1.oraclecloud.com
```

To use Oracle APEX, Database Actions, and Oracle REST Data Services, add another entry with the same IP. For example:

```
/etc/hosts entry -> 10.0.2.7 example.adb.ca-toronto-1.oraclecloudapps.com
```

You find the private endpoint IP and the FQDN as follows:

- The Private IP is shown on the Oracle Cloud Infrastructure console Autonomous Database details page for the instance.
- The FQDN is shown in the `tnsnames.ora` file in the Autonomous Database client credential wallet.

Alternatively you can use Oracle Cloud Infrastructure private DNS to provide DNS name resolution. See Private DNS for more information.

In this example there is a Dynamic Routing Gateway (DRG) between the on-premise data center and "Your VCN". The VCN contains the Autonomous Database. This also shows a route table for the VCN associated with the Autonomous Database, for outgoing traffic to CIDR 172.16.0.0/16 through the DRG.

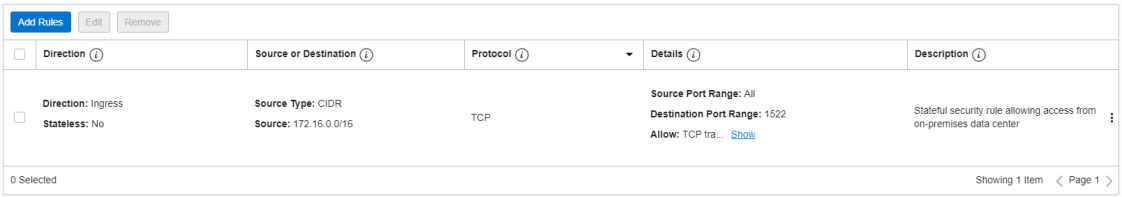
In addition to setting up the DRG, define a Network Security Group (NSG) rule to allow traffic to and from the Autonomous Database, by adding a rule for the data center CIDR range (172.16.0.0/16). In this example, define a security rule in "NSG 1" as follows:

- For Mutual TLS authentication, create a stateful rule to allow ingress traffic from the data center. This is a stateful ingress rule with the source set to the address range you want to

allow to connect to your database, protocol set to TCP, source port range set to CIDR range (172.16.0.0/16), and destination port set to 1522.

- For TLS authentication, create a stateful rule to allow ingress traffic from the data center. This is a stateful ingress rule with the source set to the address range you want to allow to connect to your database, protocol set to TCP, source port range set to CIDR range (172.16.0.0/16), and destination port set to either 1521 or 1522.
- To use Oracle APEX, Database Actions, and Oracle REST Data Services, add port 443 to the NSG rule.

The following figure shows the security rule that controls traffic for the Autonomous Database instance:



After you configure the security rule, your on-premise database application can connect to the Autonomous Database instance using the client credentials wallet. See Download Client Credentials (Wallets) for more information.

Use Zero Trust Packet Routing on Autonomous Database

Oracle Cloud Infrastructure Zero Trust Packet Routing (ZPR) protects sensitive data from unauthorized access through intent-based security policies that you write for resources, such as an Autonomous Database on a private endpoint, that you assign security attributes to.

Security attributes are labels that ZPR uses to identify and organize resources. ZPR enforces policy at the network level each time access is requested regardless of potential network architecture changes or misconfigurations. You can write ZPR based policies in your tenancy to ensure an Autonomous Database instance on a private endpoint is only accessed by authorized users/resources. ZPR is built on top of existing network security group (NSG) and security control list (SCL) rules. For a packet to reach a target, it must pass all NSG and SCL rules, and ZPR policies. If any NSG, SCL, or ZPR rule or policy doesn't allow traffic, the request is dropped.

> **Note:**
>
> On Autonomous Database Zero Trust Packet Routing (ZPR) applies only for incoming connections when the instance is configured with a private endpoint. See Configure Private Endpoints for more information.

You can secure networks with Zero Trust Packet Routing (ZPR) in three steps:

1. Create and manage security attribute namespaces and security attributes
2. Write policies using security attributes to control access to resources
3. Apply security attributes to specified resources

Avoid entering confidential information when assigning descriptions, tags, security attributes, or friendly names to cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

You have the following options to apply Zero Trust Packet Routing (ZPR) security attributes to a private endpoint on Autonomous Database:

- Apply security attributes when you provision an Autonomous Database instance and you choose **Private endpoint access only** network access during provisioning. See Provision an Autonomous Database Instance for more information.
- Apply security attributes when you clone an Autonomous Database instance and you choose **Private endpoint access only** network access during cloning. See Clone an Autonomous Database Instance for more information.
- Apply security attributes to a private endpoint on an existing Autonomous Database instance. See Configure Zero Trust Packet Routing (ZPR) for a Private Endpoint for more information.

See Getting Started with Zero Trust Packet Routing for more information.

- Configure Zero Trust Packet Routing (ZPR) for a Private Endpoint
  Shows the steps to configure Zero Trust Packet Routing (ZPR) for a private endpoint.

Configure Zero Trust Packet Routing (ZPR) for a Private Endpoint

Shows the steps to configure Zero Trust Packet Routing (ZPR) for a private endpoint.

Set or change the values for ZPR security attributes on an existing Autonomous Database instance configured with a private endpoint as follows:

1. From the Oracle Cloud Infrastructure Console, on the **Details** page, select the **Security** tab.
2. Click **Add security attributes**.
3. In the Add security attributes dialog, select a **Namespace** and enter a **Key** and a **Value**.
4. Optionally click **Add security attribute** to add additional security attributes.
5. Click **Add security attributes** to confirm.

Update Network Options to Allow TLS or Require Only Mutual TLS (mTLS) Authentication on Autonomous Database

Describes how to update the secure client connection authentication options, Mutual TLS (mTLS) and TLS.

- Network Access Prerequisites for TLS Connections
  Describes the network access configuration prerequisites for TLS connections.
- Update your Autonomous Database Instance to Allow both TLS and mTLS Authentication
- Update your Autonomous Database Instance to Require mTLS and Disallow TLS Authentication
  If your Autonomous Database instance is configured to allow TLS connections, you can update the instance to require mTLS connections and disallow TLS connections.

Network Access Prerequisites for TLS Connections

Describes the network access configuration prerequisites for TLS connections.

To allow an Autonomous Database instance to use TLS connections, either ACLs must be defined or a private endpoint must be configured:

- When an Autonomous Database instance is configured to operate over the public internet, one or more Access Control Lists (ACLs) must be defined before you use TLS authentication to connect to the database. To validate that ACLs are defined, in the **Network** area on the Autonomous Database Details page view the **Access Control List** field. This field shows **Enabled** when ACLs are defined and shows **Disabled** when ACLs are not defined.

  See Configuring Network Access with Access Control Rules (ACLs) for more information.
- When an Autonomous Database instance is configured with a private endpoint you can use TLS authentication to connect to the database. To validate that a private endpoint is defined, in the **Network** area on the Autonomous Database Details page view the **Access type** field. This field shows **Virtual Cloud Network** when a private endpoint is defined.

  See Configure Network Access with Private Endpoints for more information.

> **Note:**
>
> When an Autonomous Database instance is configured with the network access type: **Secure access from everywhere**, you can only use TLS connections to connect to the database if you specify ACLs to restrict access.

Update your Autonomous Database Instance to Allow both TLS and mTLS Authentication

If your Autonomous Database instance is configured to only allow mTLS connections, you can update the instance to allow both mTLS and TLS connections.

When you update your configuration to allow both mTLS and TLS, you can use both authentication types at the same time and connections are no longer restricted to require mTLS authentication.

You can allow TLS connections when network access is configured as follows:

- With network access configured with ACLs defined.
- With network access configured with a private endpoint defined.

> **Note:**
>
> When you configure your Autonomous Database instance network access with ACLs or a private endpoint, the ACLs or the private endpoint apply for both mTLS and TLS connections.

Perform the network access configuration prerequisites. See Network Access Prerequisites for TLS Connections for more information.

Perform the following steps as necessary:

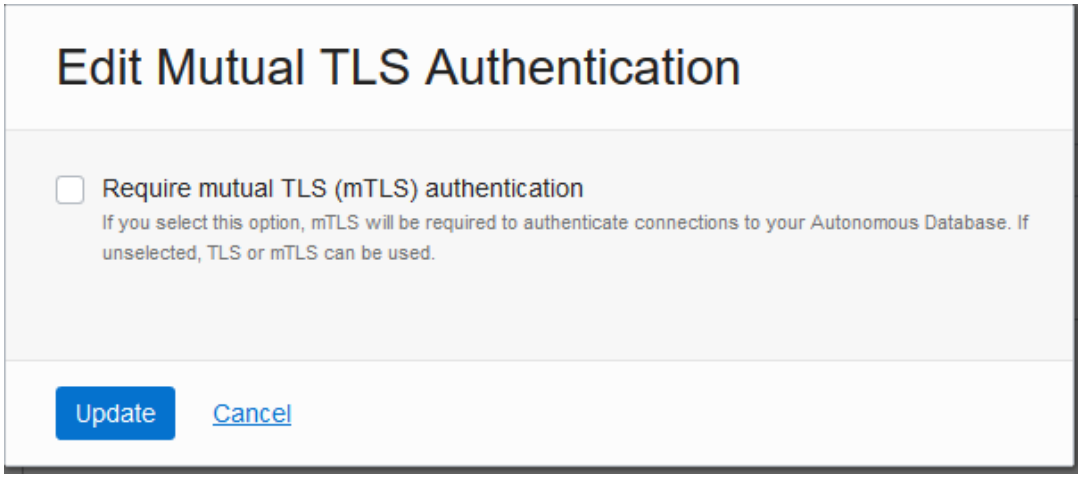
- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To change the Autonomous Database instance to allow TLS authentication, do the following:

1. On the **Autonomous Database Details** page, under **Network**, click **Edit** in the **Mutual TLS (mTLS) Authentication** field.

   This shows the Edit Mutual TLS Authentication page.

2. To change the value to allow TLS authentication, deselect **Require mutual TLS (mTLS) authentication**.



3. Click **Update**.

   The Autonomous Database Lifecycle state changes to **Updating**.

   After some time, the **Lifecycle state** shows **Available** and the **Mutual TLS (mTLS) Authentication** field changes to show **Not Required**.

After you define ACLs or configure a private endpoint and the **Mutual TLS (mTLS) Authentication** field shows **Not Required**, the ACLs or the private endpoint you specify apply to all connection types (mTLS and TLS).

Depending on the type of client, TLS connections have the following support with Autonomous Database:

- If the client is connecting with JDBC Thin using TLS authentication, the client can connect without providing a wallet. See Connect with JDBC Thin Driver for more information.
- If the client is connecting with managed ODP.NET or ODP.NET Core versions 19.13 or 21.4 (or above) using TLS authentication, the client can connect without providing a wallet. See Connect Microsoft .NET, Visual Studio Code, and Visual Studio Without a Wallet for more information.
- If the client is connecting with SQLNet and Oracle Call Interface (OCI), and for certain other connection types with TLS authentication, the clients must provide the CA certificate in a wallet. See Prepare for Oracle Call Interface, ODBC, and JDBC OCI Connections Using TLS Authentication and Connect SQL*Plus Without a Wallet for more information.

Update your Autonomous Database Instance to Require mTLS and Disallow TLS Authentication

If your Autonomous Database instance is configured to allow TLS connections, you can update the instance to require mTLS connections and disallow TLS connections.

> **Note:**
>
> When you update an Autonomous Database instance to require Mutual TLS (mTLS) connections, existing TLS connections are disconnected.

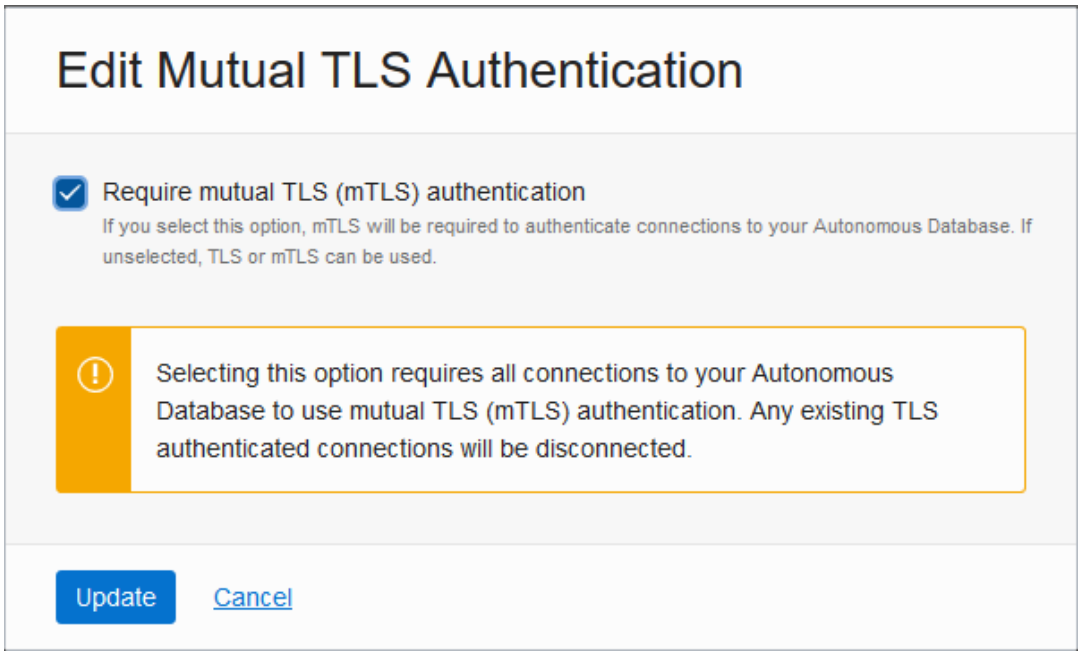
Perform the following steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select your Autonomous Database from the links under the **Display name** column.

To change the Autonomous Database instance to require mTLS authentication and to not allow TLS authentication, do the following:

1. On the **Autonomous Database Details** page, under **Network**, click **Edit** in the **Mutual TLS (mTLS) Authentication** field.

   This shows the Edit Mutual TLS Authentication page.
2. Select **Require mutual TLS (mTLS) authentication**.



3. Click **Update**.

   The Autonomous Database Lifecycle state changes to **Updating**.

After some time, the **Lifecycle state** shows **Available** and the **Mutual TLS (mTLS) Authentication** field changes to show **Required**.

Audit Autonomous Database

Autonomous Database provides auditing that allows you to monitor Oracle database activities.

- About Auditing Autonomous Database
  Autonomous Database provides auditing to track, monitor, and record database actions. Auditing can help you detect security risks and improve regulatory compliance for your database.
- Register Oracle Data Safe on Autonomous Database
  Use Oracle Data Safe to apply auditing policies for database users, for administrative users, to apply predefined auditing policies or to extend the audit data record retention for your Autonomous Database instance.
- Extend Audit Record Retention with Oracle Data Safe on Autonomous Database
  Use Oracle Data Safe to extend the audit data record retention to a specified number of months.
- View and Manage Oracle Data Safe Audit Trails on Autonomous Database
  Data Safe uses audit trails to define where to retrieve the audit data and to collect Autonomous Database audit records. During the registration process, Oracle Data Safe discovers the audit trails and creates an audit trail resource.
- View and Manage Audit Policies with Oracle Data Safe on Autonomous Database
  Use Oracle Data Safe to set audit policies for your Autonomous Database instance.
- Generate Audit Reports with Data Safe on Autonomous Database
  Data Safe includes out-of-box audit data reports, and you can create custom reports to suit your needs.

About Auditing Autonomous Database

Autonomous Database provides auditing to track, monitor, and record database actions. Auditing can help you detect security risks and improve regulatory compliance for your database.

- Audit Features on Autonomous Database
  Autonomous Database includes extensive, sophisticated audit capabilities that allow you capture the audit information you need for your organization. Autonomous Database provides default auditing.
- Audit Data on Autonomous Database
  Autonomous Database protects audit data and writes its audit trail to the `UNIFIED_AUDIT_TRAIL` data dictionary view.
- Default Audit Policies on Autonomous Database
  Autonomous Database provides auditing to track, monitor, and record activities on your database.

Audit Features on Autonomous Database

Autonomous Database includes extensive, sophisticated audit capabilities that allow you capture the audit information you need for your organization. Autonomous Database provides default auditing.

In addition, you can use either of the following to apply auditing policies:

- Use Oracle Data Safe to apply auditing policies for database users, for administrative users, and to apply predefined auditing policies or to apply customized auditing policies. See Activity Auditing Overview for more information.
- Configure Oracle Database Audit Policies. See Configuring Audit Policies for more information.

You can configure auditing to accomplish the following:

- Enable accountability for actions. These include actions taken in a particular schema, table, or row, or affecting specific content.
- Deter users, or others, such as intruders, from inappropriate actions based on their accountability.
- Investigate suspicious activity. For example, if a user is logging into the database using the application's database credentials, then auditing connections to the database lets you determine that the login came from a user's workstation instead of from the application server.
- Notify an auditor of the actions of an unauthorized user. For example, notify an auditor when an unauthorized user attempts to delete data from a table.
- Monitor and gather data about specific database activities. For example, you can gather statistics about which tables are being updated, the number of failed logins, or how many concurrent users connect at peak times.
- Detect problems with an authorization or access control implementation. For example, you can create audit policies that you expect will never generate an audit record because the data is protected in other ways. However, if these policies generate audit records, then you will know the other security controls are not properly implemented.
- Address auditing requirements for compliance. Regulations such as the following have common auditing-related requirements:
  - European Union General Data Protection Regulation (GDPR)
  - Sarbanes-Oxley Act
  - Health Insurance Portability and Accountability Act (HIPAA)
  - International Convergence of Capital Measurement and Capital Standards: a Revised Framework (Basel II)
  - Japan Privacy Law
  - European Union Directive on Privacy and Electronic Communications

Audit Data on Autonomous Database

Autonomous Database protects audit data and writes its audit trail to the `UNIFIED_AUDIT_TRAIL` data dictionary view.

The underlying table storing audit data on Autonomous Database is `AUDSYS.AUD$UNIFIED`. This table is protected and does not allow users to perform DML/DDL operations or to purge the table (any attempt to perform these actions automatically produces an audit record). After an audit record is written, the only activity allowed is for the ADMIN user to perform a `PURGE`. The ADMIN has the AUDIT_ADMIN role that is required to run a `PURGE`. If you assign the AUDIT_ADMIN role to another user, then that user could also perform a `PURGE`.

Depending on the number and type of audit policies you use and the amount of activity, over time the audit trail can grow to use a large amount of storage. Autonomous Database provides the following ways to limit the storage required for audit data:

- Each Autonomous Database instance runs an automated purge job once a day to remove all audit records older than fourteen (14) days.
- Users with the AUDIT_ADMIN role can purge audit records manually using the `DBMS_AUDIT_MGMT.CLEAN_AUDIT_TRAIL` procedure. See DBMS_AUDIT_MGMT for more information.

If you need a longer audit data retention period than 14 days, use Oracle Data Safe to retain audit data. See Extend Audit Record Retention with Oracle Data Safe on Autonomous Database for more information.

Autonomous Database audits and logs every operation carried out in your database by the Oracle Cloud Infrastructure Operations teams. See View Oracle Cloud Infrastructure Operations Actions for more information on how to audit Operations activities.

Default Audit Policies on Autonomous Database

Autonomous Database provides auditing to track, monitor, and record activities on your database.

By default, Autonomous Database applies audit policies to audit the following database activities:

- All activity by Oracle Cloud Operations
- All login failures to the database
- All password changes
- Attempts to create or alter procedures
- Execution of certain procedures, including procedures in the packages: `UTL_HTTP` or `UTL_SMTP` that connect to the network

In addition, you can use either of the following to apply additional auditing policies:

- Use Oracle Data Safe to apply auditing policies for database users, for administrative users, and to apply predefined auditing policies or to apply customized auditing policies. For more information, see:
  - View and Manage Audit Trails
  - View and Manage Audit Policies
- Configure Oracle Database Audit Policies. See Configuring Audit Policies for more information.

Register Oracle Data Safe on Autonomous Database

Use Oracle Data Safe to apply auditing policies for database users, for administrative users, to apply predefined auditing policies or to extend the audit data record retention for your Autonomous Database instance.

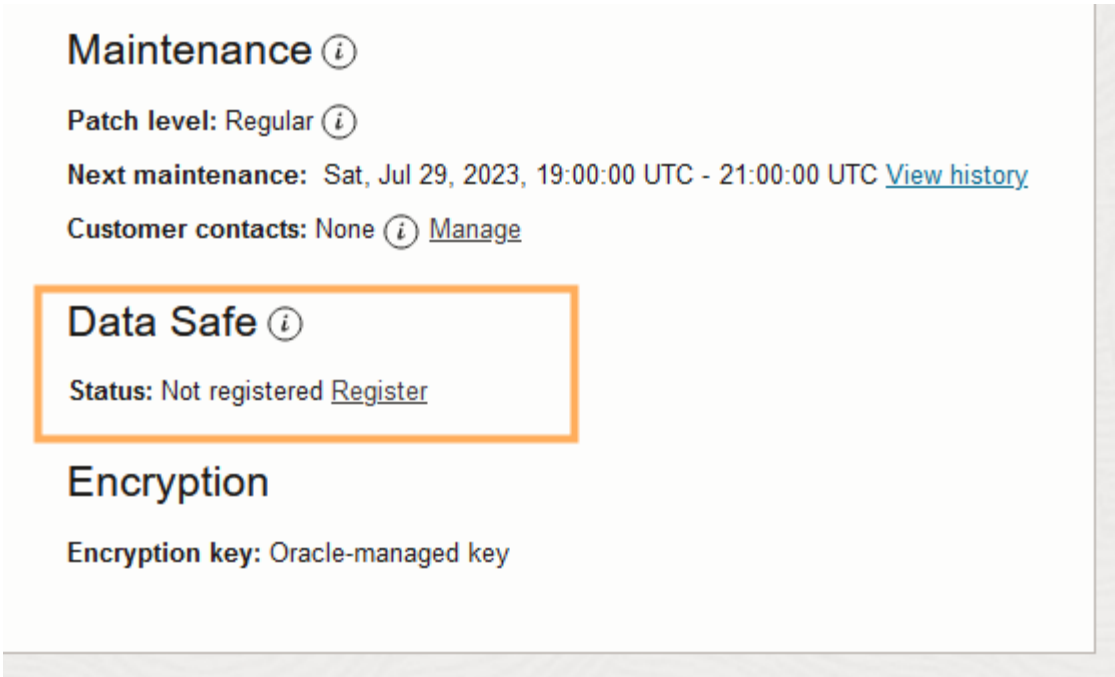
Register your Autonomous Database instance with Oracle Data Safe as follows:

1. Access your Autonomous Database instance from the Oracle Cloud Infrastructure Console.
   a. Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
   b. From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.

c. On the Autonomous Databases page select an Autonomous Database from the links under the **Display Name** column.

2. Register your Autonomous Database instance with Oracle Data Safe.

a. On the Autonomous Database Details page, under **Data Safe**, click **Register**.



b. In the Register database with Data Safe dialog, click **Confirm**.

The Data Safe status shows: **Registering**. This step takes about 15 to 20 minutes.

After Oracle Data Safe is registered, the Data Safe status shows **Registered** and two links: **View** and **Deregister**.

Click **View** to show the Data Safe register database details page.

Click **Deregister** to disable Oracle Data Safe.

Extend Audit Record Retention with Oracle Data Safe on Autonomous Database

Use Oracle Data Safe to extend the audit data record retention to a specified number of months.

First register your Autonomous Database instance with Oracle Data Safe. See Register Oracle Data Safe on Autonomous Database for more information.

After your Autonomous Database instance is registered, you can specify the Data Safe retention period.

See Update Retention Periods for a Target Database for more information.

View and Manage Oracle Data Safe Audit Trails on Autonomous Database

Data Safe uses audit trails to define where to retrieve the audit data and to collect Autonomous Database audit records. During the registration process, Oracle Data Safe discovers the audit trails and creates an audit trail resource.

Oracle Data Safe lists resources on the Audit Trails page. To access the Audit Trails page, under Security Center in Data Safe click **Activity Auditing**, and then, on the Activity Auditing page under **Related Resources** click **Audit Trails**. You can discover new audit trails at any time and remove audit trail resources in Oracle Data Safe as needed.

First register your Autonomous Database instance with Oracle Data Safe. See Register Oracle Data Safe on Autonomous Database for more information.

When the Autonomous Database is stopped or restarted, the following happens:

- The audit trail switches to a retrying state and Data Safe makes multiple attempts to reconnect for four (4) hours. The Audit Trail **Collection State** field shows: **RETRYING**.

  In this case, if the Autonomous Database (target database) starts, the audit trail automatically resumes.
- After four hours, the audit trail switches to a stopped state. The Audit Trail **Collection State** field shows: **STOPPED_NEEDS_ATTN**.

  In this case, when the Autonomous Database (target database) starts and the audit trail collection state is **STOPPED_NEEDS_ATTN**, you can manually resume the audit trail.

You can also manually stop or delete the audit trail. Deleting the audit trail does not remove audit records that have already been collected. Those records remain in Data Safe until the retention period is reached.

See View and Manage Audit Trails for more information.

View and Manage Audit Policies with Oracle Data Safe on Autonomous Database

Use Oracle Data Safe to set audit policies for your Autonomous Database instance.

First register your Autonomous Database instance with Oracle Data Safe. See Register Oracle Data Safe on Autonomous Database for more information.

After your Autonomous Database instance is registered, access Oracle Data Safe to set audit policies.

See View and Manage Audit Policies for more information.

Generate Audit Reports with Data Safe on Autonomous Database

Data Safe includes out-of-box audit data reports, and you can create custom reports to suit your needs.

After you enable and register Oracle Data Safe, and you add a trail to collect audit data from your Autonomous Database instance, then you can use the reports to monitor activity for your database.

See View and Manage Audit Reports for more information.

Rotate Wallets for Autonomous Database

Wallet rotation lets you invalidate existing client certification keys for a database instance or for all Autonomous Database instances that a cloud account owns in a region.

- About Wallet Rotation
  You have the option to perform one of two types of wallet rotation: immediate or with a grace period.
- Rotate Wallets with Immediate Rotation
  Immediate wallet rotation lets you invalidate existing client certification keys for an Autonomous Database instance or for all Autonomous Database instances that a cloud account owns in a region.
- Rotate Wallets with Grace Period

About Wallet Rotation

You have the option to perform one of two types of wallet rotation: immediate or with a grace period.

- **Immediate** wallet rotation initiates immediately, without delay.
- **After a grace period** wallet rotation occurs with a grace period. During the grace period the old client certification keys remain valid for a selected time of 1 hour to 24 hours. After the grace period expires only the new client certification keys are valid.

You may want to rotate wallets for any of the following reasons:

- If your organization's policies require regular client certification key rotation.
- When a client certification key or a set of keys is suspected to be compromised.

Rotate Wallets with Immediate Rotation

Immediate wallet rotation lets you invalidate existing client certification keys for an Autonomous Database instance or for all Autonomous Database instances that a cloud account owns in a region.

There are two options for immediate client certification key rotation:

- Per-database with **Instance wallet** selected:
  - For the database whose certification key is rotated, any existing database specific instance wallets will be void. After you rotate a wallet you have to download a new wallet to connect to the database.
  - Regional wallets containing all database certification keys continue to work.
  - All user sessions are terminated for the database whose wallet is rotated. User session termination begins after wallet rotation completes, however this process does not happen immediately.

> **Note:**
>
> If you want to terminate all connections immediately after the wallet rotation completes, Oracle recommends that you restart the Autonomous Database instance. This provides the highest level of security for your database.

- Regional level with **Regional wallet** selected:
  - For the region whose certification key is rotated, both regional and database specific instance wallets will be void. After you rotate a wallet you have to download new regional or instance wallets to connect to any database in the region.
  - All user sessions are terminated for the databases in the region whose wallet is rotated. User session termination begins after wallet rotation completes, however this process does not happen immediately.

> **Note:**
>
> If you want to terminate all connections immediately after the wallet rotation completes, Oracle recommends that you restart the Autonomous Database instances in the region. This provides the highest level of security for your database.

To immediately rotate the client certification key for a given database or for all Autonomous Database instances that a cloud account owns in a region:

1. Navigate to the Autonomous Database details page.
2. Click **Database connection**.
3. On the **Database connection** page select the **Wallet type**:
   - **Instance wallet**: Wallet rotation for a single database only; this provides a database-specific wallet rotation.
   - **Regional wallet**: Wallet rotation for all Autonomous Databases for a given tenant and region (this option rotates the client certification key for all service instances that a cloud account owns).
4. Click **Rotate wallet**.
5. Enter the name as shown in the dialog to confirm the wallet rotation.
6. In the Rotate Wallet dialog, click **Rotate**.

The Database Connection page shows: **Rotation in Progress**.

After the rotation completes, the **Wallet last rotated** field shows the last rotation date and time.

Oracle recommends you provide a database-specific instance wallet to end users and for application use whenever possible, with Wallet type set to **Instance wallet** when you use **Download wallet**. Regional wallets should only be used for administrative purposes that require potential access to all Autonomous Databases within a region.

You can also use the Autonomous Database API to rotate wallets using `UpdateAutonomousDatabaseRegionalWallet` and `UpdateAutonomousDatabaseWallet`. See Autonomous Database Wallet Reference for more information.

Rotate Wallets with Grace Period

Autonomous Database allows you to rotate wallets for an Autonomous Database instance or for all instances that a cloud account owns in a region. with a grace period of 1 hour to 24 hours.

Setting a grace period allows you to perform wallet rotation without down time. During the grace period you can inform users to download the new wallet and to update their applications to use the new wallet. During the grace period both the old and new client certification keys are

valid. When the grace period expires, Autonomous Database invalidates the old client certification keys and only the new client certification keys are valid.

There are two options for client certification key rotation with a grace period:

- Per-database with **Instance wallet** selected:
  - For the database whose certification key is rotated, database specific instance wallets that were in use before the wallet rotation be void after the grace period expires.
  - After you perform client certification key rotation with a grace period, you can immediately download a wallet and use the new wallet to connect to the database.
  - Regional wallets containing all database certification keys continue to work.
  - After the grace period expires, existing connections using the old wallet continue to work.

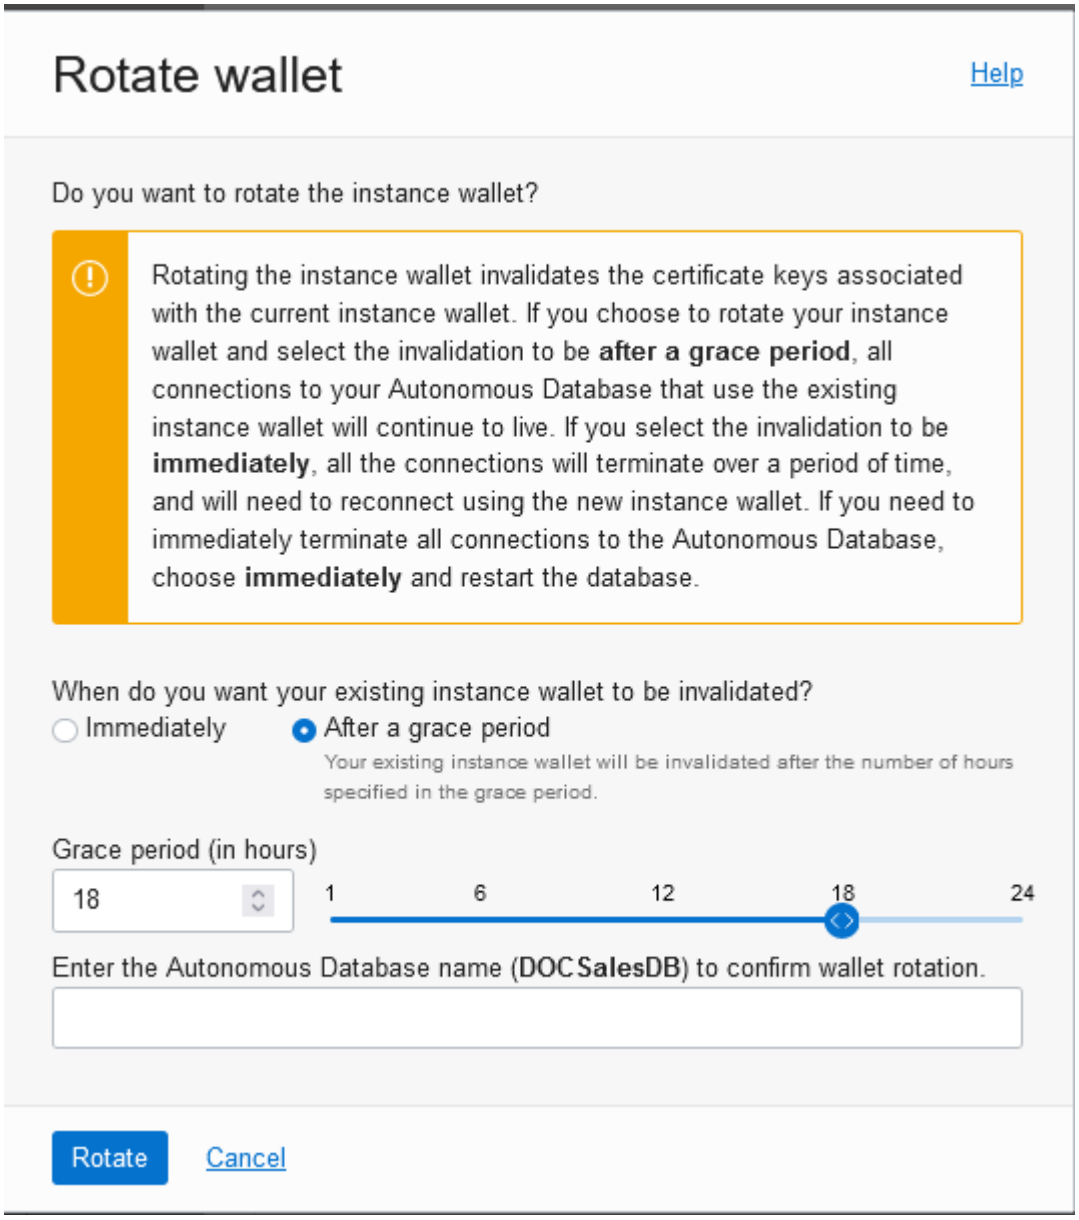
> **Note:**
>
> After the grace period completes, if you want to terminate any connections using the old wallet, Oracle recommends that you restart the Autonomous Database instance.

- Regional level with **Regional wallet** selected:
  - For the region whose certification key is rotated, both regional and database specific instance wallets will be void. After the grace period expires you have to download new regional or instance wallets to connect to any database in the region.
  - After the grace period expires, existing connections using the old wallets continue to work.

> **Note:**
>
> After the grace period completes, if you want to terminate any connections using the old wallets, Oracle recommends that you restart every Autonomous Database instance in the region.

To rotate the client certification key with a grace period for a given database or for all for all Autonomous Database instances that a cloud account owns in a region:

1. Navigate to the Autonomous Database details page.
2. Click **Database connection**.
3. On the **Database connection** page select the **Wallet type**:
   - **Instance wallet**: Wallet rotation for a single database only; this provides a database-specific wallet rotation.
   - **Regional wallet**: Wallet rotation for all Autonomous Databases for a given tenant and region (this option rotates the client certification key for all service instances in the region that a cloud account owns).
4. Click **Rotate wallet**.
5. Select **After a grace period**.
6. In the **Grace period (in hours)** area, either enter a value in the text field or use the slider to select a value.



7. Enter the name as shown in the dialog to confirm the wallet rotation.
8. In the Rotate Wallet dialog, click **Rotate**.

The Database Connection page shows: **Rotation in Progress**.

After the rotation completes, the **Wallet last rotated** field shows the last rotation date and time.

Notes for wallet rotation with a grace period:

- Always Free Autonomous Databases only support immediate wallet rotation (wallet rotation with a grace period is not supported).
- Oracle recommends you provide a database-specific instance wallet to end users and for application use whenever possible, with Wallet type set to **Instance wallet** when you use **Download wallet**. Regional wallets should only be used for administrative purposes that require potential access to all Autonomous Databases within a region.

You can also use the Autonomous Database API to rotate wallets using `UpdateAutonomousDatabaseRegionalWallet` and `UpdateAutonomousDatabaseWallet`. See Autonomous Database Wallet Reference for more information.

Use Oracle Database Vault with Autonomous Database

Oracle Database Vault implements powerful security controls for your database. These unique security controls restrict access to application data by privileged database users, reducing the risk of insider and outside threats and addressing common compliance requirements.

See What Is Oracle Database Vault? for more information.

- Oracle Database Vault Users and Roles on Autonomous Database
  Oracle Database Vault provides powerful security controls to help protect application data from unauthorized access, and to implement separation of duties between administrators and data owners to comply with privacy and regulatory requirements.
- Enable Oracle Database Vault on Autonomous Database
  Shows the steps to enable Oracle Database Vault on Autonomous Database.
- Disable Oracle Database Vault on Autonomous Database
  Shows the steps to disable Oracle Database Vault on Autonomous Database.
- Disable User Management with Oracle Database Vault on Autonomous Database
  Shows how to disallow user management related operations for specified components on Autonomous Database with Oracle Database Vault enabled.
- Enable User Management with Oracle Database Vault on Autonomous Database
  Shows the steps to allow user management for a specified component on Autonomous Database with Oracle Database Vault enabled.

Oracle Database Vault Users and Roles on Autonomous Database

Oracle Database Vault provides powerful security controls to help protect application data from unauthorized access, and to implement separation of duties between administrators and data owners to comply with privacy and regulatory requirements.

By default the ADMIN user has the `DV_OWNER` and `DV_ACCTMGR` roles. If you want to set up separate users for `DV_OWNER` and `DV_ACCTMGR` accounts. See Oracle Database Vault Schemas, Roles, and Accounts for more information.

The user management is by default enabled for the APEX component when Oracle Database Vault is enabled. When user management is enabled, the APEX users who have the necessary roles to `CREATE | ALTER | DROP` users have the needed privileges to perform these operations when Database Vault is enabled. To change this, see Disable User Management with Oracle Database Vault on Autonomous Database.

On Autonomous Database with Oracle Database Vault enabled, grant the following privileges:

- When using Oracle GoldenGate, grant the `GGADMIN` user `DV_GOLDENGATE_ADMIN` and `DV_GOLDENGATE_REDO_ACCESS`.
- The ADMIN user must grant the `BECOME USER` privilege to users who need to use Oracle Data Pump. To perform some Oracle Data Pump operations additional Oracle Database Vault authorization may be needed. For example to run a full database export or to export a realm protected schema requires using `DBMS_MACADM.AUTHORIZE_DATAPUMP_USER`.

  See AUTHORIZE_DATAPUMP_USER Procedure for more information.

- In order for `DBMS_CLOUD` credential related APIs to work when Oracle Database Vault is enabled and the credential owner's schema is protected using a Database Vault realm, you must add authorizations for the `C##CLOUD$SERVICE` user to the Database Vault realm.

  For example:

  ```
  BEGIN
      DBMS_MACADM.ADD_AUTH_TO_REALM(realm_name    => 'PROTECT_ADMIN',
          grantee       => 'C##CLOUD$SERVICE',
          rule_set_name => 'Enabled',
          auth_options  => DBMS_MACUTL.G_REALM_AUTH_PARTICIPANT);
  END;
  /
  ```

  Where *PROTECT_ADMIN* is your Oracle Database Vault realm.

  See ADD_AUTH_TO_REALM Procedure for more information.

Enable Oracle Database Vault on Autonomous Database

Shows the steps to enable Oracle Database Vault on Autonomous Database.

Oracle Database Vault is disabled by default on Autonomous Database. To configure and enable Oracle Database Vault on Autonomous Database, do the following:

1. Configure Oracle Database Vault using the following command:

   ```
   EXEC DBMS_CLOUD_MACADM.CONFIGURE_DATABASE_VAULT('adb_dbv_owner',
   'adb_dbv_acctmgr');
   ```

   Where:

   - *adb_dbv_owner* is the Oracle Database Vault owner.
   - *adb_dbv_acctmgr* is the account manager.

   See CONFIGURE_DATABASE_VAULT Procedure for more information.

2. Enable Oracle Database Vault:

   ```
   EXEC DBMS_CLOUD_MACADM.ENABLE_DATABASE_VAULT;
   ```

   See ENABLE_DATABASE_VAULT Procedure for more information.

3. Restart the Autonomous Database instance.

   See Restart Autonomous Database for more information.

Use the following command to check if Oracle Database Vault is enabled or disabled:

```
SELECT * FROM DBA_DV_STATUS;
```

Output similar to the following appears:

```
NAME                 STATUS
-------------------- -----------
```

```
DV_CONFIGURE_STATUS  TRUE
DV_ENABLE_STATUS     TRUE
```

The `DV_ENABLE_STATUS` value `TRUE` indicates Oracle Database Vault is enabled.

> **Note:**
>
> Autonomous Database maintenance operations such as backups and patching are not affected when Oracle Database Vault is enabled.

See Disable Oracle Database Vault on Autonomous Database for information on disabling Oracle Database Vault.

Disable Oracle Database Vault on Autonomous Database

Shows the steps to disable Oracle Database Vault on Autonomous Database.

To disable Oracle Database Vault on Autonomous Database, do the following:

1. Disable Oracle Database Vault.

   ```
   EXEC DBMS_CLOUD_MACADM.DISABLE_DATABASE_VAULT;
   ```

   See DISABLE_DATABASE_VAULT Procedure for more information.
2. Restart the Autonomous Database instance.

   See Restart Autonomous Database for more information.

Use the following command to check if Oracle Database Vault is enabled or disabled:

```
SELECT * FROM DBA_DV_STATUS;
```

Output similar to the following appears:

```
NAME                 STATUS
-------------------- -----------
DV_CONFIGURE_STATUS  TRUE
DV_ENABLE_STATUS     FALSE
```

The `DV_ENABLE_STATUS` value `FALSE` indicates Oracle Database Vault is disabled.

Disable User Management with Oracle Database Vault on Autonomous Database

Shows how to disallow user management related operations for specified components on Autonomous Database with Oracle Database Vault enabled.

Autonomous Database with Oracle Database Vault enabled has user management, by default, enabled for the Oracle APEX console. If you want to enforce stricter separation of duty and disallow user management from this console, use `DBMS_CLOUD_MACADM.DISABLE_USERMGMT_DATABASE_VAULT`.

1. As a user granted `DV_ACCTMGR` and `DV_ADMIN` roles you can disable user management for specified components.

2. To disable user management for a specified component, for example for the APEX component, use the following command:

```
EXEC DBMS_CLOUD_MACADM.DISABLE_USERMGMT_DATABASE_VAULT('APEX');
```

See DISABLE_USERMGMT_DATABASE_VAULT Procedure for more information.

Enable User Management with Oracle Database Vault on Autonomous Database

Shows the steps to allow user management for a specified component on Autonomous Database with Oracle Database Vault enabled.

Autonomous Database with Oracle Database Vault enabled has user management, by default, enabled for the Oracle APEX console. This allows user management for operations such as `CREATE USER`, `ALTER USER`, and `DROP USER` from the specified component in Autonomous Database.

Use `DBMS_CLOUD_MACADM.ENABLE_USERMGMT_DATABASE_VAULT` to allow specified user accounts to perform user management when Oracle Database Vault is enabled. Use this procedure if user management is disabled and you want to enable it again.

1. A user granted `DV_ACCTMGR` and `DV_ADMIN` roles can enable user management for specified components.
2. To enable user management for a specified component, for example for the APEX component, use the following command:

```
EXEC DBMS_CLOUD_MACADM.ENABLE_USERMGMT_DATABASE_VAULT('APEX');
```

See ENABLE_USERMGMT_DATABASE_VAULT Procedure for more information.

IAM Policies for Autonomous Database

Provides information on IAM policies required for API operations on Autonomous Database.

Oracle Autonomous Database relies on the IAM (Identity and Access Management) service to authenticate and authorize cloud users to perform operations that use any of the Oracle Cloud Infrastructure interfaces (the console, REST API, CLI, or SDK).

The IAM service uses **groups**, **compartments,** and **policies** to control which cloud users can access which resources.

- Policy Details for Autonomous Database
  This topic covers details for writing policies to control access to Autonomous Database resources.
- IAM Permissions and API Operations for Autonomous Database
  This topic covers the available IAM permissions for operations on Autonomous Database.
- Provide Specific Privileges in IAM Policies to Manage Autonomous Database
  Lists IAM policies that you can use with an authorization verb and a condition to grant more granular operations to a group.

Policy Details for Autonomous Database

This topic covers details for writing policies to control access to Autonomous Database resources.

A policy defines what kind of access a group of users has to a specific resource in an individual compartment. For more information, see Getting Started with Policies.

Resource-Types

An aggregate resource-type covers the list of individual resource-types that directly follow. For example, writing one policy to allow a group to have access to the `autonomous-database-family` is equivalent to writing four separate policies for the group that would grant access to the `autonomous-databases`, `autonomous-backups` resource-types. For more information, see Resource-Types.

Resource-Types for Autonomous Database
Aggregate Resource-Type:

`autonomous-database-family`

Individual Resource-Types:

`autonomous-databases`

`autonomous-backups`

Details for Verb + Resource-Type Combinations

The level of access is cumulative as you go from `inspect > read > use > manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

For example, the `read` verb for the `autonomous-databases` resource-type covers the same permissions and API operations as the `inspect` verb, plus the `AUTONOMOUS_DATABASE_CONTENT_READ` permission. The `read` verb partially covers the `CreateAutonomousDatabaseBackup` operation, which also needs manage permissions for `autonomous-backups`.

The following tables show the Permissions and API operations covered by each verb. For information about permissions, see Permissions.

> **Note:**
>
> The resource family covered by **autonomous-database-family** can be used to grant access to database resources associated with all the Autonomous Database workload types.

autonomous-databases Resource Types

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	AUTONOMOUS_DATABASE_INSPECT	GetAutonomousDatabase, GetAutonomousDatabaseBackupConfig, GetAutonomousDatabaseCapability, ListAutonomousDatabases, ListAutonomousDatabaseClones, ListAutonomousDatabasePeers, ListAutonomousDatabaseRefreshableClones, ResourcePoolShapes	*none*
read	*INSPECT* + AUTONOMOUS_DATABASE_CONTENT_READ	GenerateAutonomousDatabasePerformanceData, GenerateAutonomousDatabaseWallet, GetAutonomousDatabaseRegionalWallet, GetAutonomousDatabaseWallet, RetrieveDatabasePerformanceBulkData	CreateAutonomousDatabaseBackup (also needs manage autonomous-backups)

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
use	*READ* + AUTONOMOUS_DATABASE_CONTENT_WRITE AUTONOMOUS_DATABASE_UPDATE	AutonomousDatabaseManualRefresh, CancelAutonomousDatabaseSession, ChangeDisasterRecoveryConfiguration, ConfigureAutonomousDatabaseVaultKey, DeregisterAutonomousDatabaseDataSafe, DisableAutonomousDatabaseOperationsInsights, DisableDatabaseManagement, EnableAutonomousDatabaseOperationsInsights, EnableDatabaseManagement, FailOverAutonomousDatabase, GetAutonomousDatabaseConsoleToken, RegisterAutonomousDatabaseDataSafe, RestartAutonomousDatabase, RotateAutonomousDatabaseEncryptionKey, ShrinkAutonomousDatabase, StartAutonomousDatabase, StopAutonomousDatabase, SwitchOverAutonomousDatabase, UpdateAutonomousDatabaseRegionalWallet, UpdateAutonomousDatabase	RestoreAutonomousDatabase (also needs read autonomous-backups) ChangeAutonomousDatabaseCompartment (also needs read autonomous-backups)
manage	*USE* + AUTONOMOUS_DATABASE_CREATE AUTONOMOUS_DATABASE_DELETE	CreateAutonomousDatabase, DeleteAutonomousDatabase	*none*

List of Operations and Required IAM Policies to Manage an Autonomous Database Instance

Operation	Required IAM Policies
Add peer database	use autonomous-databases
Add security attributes	use autonomous-databases

Operation	Required IAM Policies
Change compute model	`use autonomous-databases`
Change database mode	`use autonomous-databases`
Change Network	`use autonomous-databases`
Change workload type	`use autonomous-databases`
Clone an Autonomous Database	`manage autonomous-databases` See IAM Permissions and API Operations for Autonomous Database for additional cloning permissions on Autonomous Database.
Create an Autonomous Database	`manage autonomous-databases` `read autonomous-databases`
Edit Database Tools Configuration	`use autonomous-databases`
Edit start/stop schedule	`use autonomous-databases`
Enable elastic pool	`use autonomous-databases`
Enable or disable auto scaling for an Autonomous Database	`use autonomous-databases`
Join elastic pool	`use autonomous-databases`
Manage customer contacts	`use autonomous-databases`
Manage encryption key	`use autonomous-databases`
Move an Autonomous Database to another compartment	`use autonomous-databases` in the database's current compartment and in the compartment you are moving it to `read autonomous-backups`
Rename an Autonomous Database	`use autonomous-databases`
Restart an Autonomous Database	`use autonomous-databases`
Restore an Autonomous Database	`use autonomous-databases` `read autonomous-backups`
Scale the ECPU count or storage of an Autonomous Database	`use autonomous-databases`
Set ADMIN user password	`use autonomous-databases`
Stop or start an Autonomous Database	`use autonomous-databases`
Switchover	`use autonomous-databases`
Terminate an Autonomous Database	`manage autonomous-databases`
Update disaster recovery	`use autonomous-databases`
Update display name	`use autonomous-databases`
Update license and Oracle Database Edition	`use autonomous-databases`
Update network access for ACLs	`use autonomous-databases`
Update network access for a private endpoint	`use autonomous-databases`
View a list of an Autonomous Databases	`inspect autonomous-databases`
View details of an Autonomous Database	`inspect autonomous-databases`

autonomous-backups

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	`AUTONOMOUS_DB_BACKUP_INSPECT`	`ListAutonomousDatabaseBackups`, `GetAutonomousDatabaseBackup`	*none*
manage	*`USE`* + `AUTONOMOUS_DB_BACKUP_CREATE` `AUTONOMOUS_DB_BACKUP_DELETE`	`DeleteAutonomousDatabaseBackup`	`CreateAutonomousDatabaseBackup` (also needs `read autonomous-databases`)
read	*`INSPECT`* + `AUTONOMOUS_DB_BACKUP_CONTENT_READ`	*no extra*	`RestoreAutonomousDatabase` (also needs `use autonomous-databases`) `ChangeAutonomousDatabaseCompartment` (also needs `use autonomous-databases`)
use	*READ* + *no extra*	*no extra*	*none*

Supported Variables

All of the general OCI Identity and Access Management variables are supported. See General Variables for All Requests for more information.

Additionally, you can use the `target.id` variable with the OCID of a database after creation of a database and the `target.workloadType` variable with a value as shown in the following table:

target.workloadType Value	Description
`OLTP`	Online Transaction Processing, used for Autonomous Databases with Transaction Processing workload.
`DW`	Data Warehouse, used for Autonomous Databases with Data Warehouse workload.
`AJD`	Autonomous JSON Database used for Autonomous Databases with JSON workload.
`APEX`	APEX Service used for Autonomous Database APEX Service.

Example policy using the `target.id` variable:

```
Allow group ADB-Admins to manage autonomous-databases in tenancy where
target.id = 'OCID'
```

Example policy using the `target.workloadType` variable:

```
Allow group ADB-Admins to manage autonomous-databases in tenancy where
target.workloadType = 'AJD'
```

IAM Permissions and API Operations for Autonomous Database

This topic covers the available IAM permissions for operations on Autonomous Database.

The following are the IAM permissions for Autonomous Database:

- AUTONOMOUS_DATABASE_CONTENT_READ
- AUTONOMOUS_DATABASE_CONTENT_WRITE
- AUTONOMOUS_DATABASE_CREATE

  See Cloning Permissions for additional cloning limitations.
- AUTONOMOUS_DATABASE_DELETE
- AUTONOMOUS_DATABASE_INSPECT
- AUTONOMOUS_DATABASE_UPDATE
- AUTONOMOUS_DB_BACKUP_CONTENT_READ
- AUTONOMOUS_DB_BACKUP_CREATE
- AUTONOMOUS_DB_BACKUP_INSPECT
- NETWORK_SECURITY_GROUP_UPDATE_MEMBERS
- VNIC_ASSOCIATE_NETWORK_SECURITY_GROUP

Example policy for a group to have permissions to create Autonomous Database in a compartment:

```
Allow group group-name to manage autonomous-database in compartment id
compartment-ocid
    where all{request-permission = 'AUTONOMOUS_DATABASE_UPDATE'}
```

Permissions Required to Use Operation	API Operation
AUTONOMOUS_DATABASE_CONTENT_READ	GenerateAutonomousDatabaseWallet GetAutonomousDatabaseRegionalWallet GetAutonomousDatabaseWallet RetrieveDatabasePerformanceBulkData
AUTONOMOUS_DATABASE_CREATE	CreateAutonomousDatabase
AUTONOMOUS_DATABASE_DELETE	DeleteAutonomousDatabase
AUTONOMOUS_DATABASE_INSPECT	GetAutonomousDatabase ListAutonomousDatabaseClones ListAutonomousDatabasePeers ListAutonomousDatabaseRefreshableClones ListAutonomousDatabases ResourcePoolShapes

Permissions Required to Use Operation	API Operation
AUTONOMOUS_DATABASE_UPDATE	AutonomousDatabaseManualRefresh ConfigureAutonomousDatabaseVaultKey DeregisterAutonomousDatabaseDataSafe DisableAutonomousDatabaseOperationsInsights DisableDatabaseManagement EnableAutonomousDatabaseOperationsInsights EnableDatabaseManagement FailOverAutonomousDatabase RegisterAutonomousDatabaseDataSafe RestartAutonomousDatabase RotateAutonomousDatabaseEncryptionKey ShrinkAutonomousDatabase StartAutonomousDatabase StopAutonomousDatabase SwitchOverAutonomousDatabase UpdateAutonomousDatabaseWallet UpdateAutonomousDatabaseRegionalWallet
AUTONOMOUS_DB_BACKUP_INSPECT	GetAutonomousDatabaseBackup ListAutonomousDatabaseBackups
AUTONOMOUS_DB_BACKUP_UPDATE	UpdateAutonomousDatabaseBackup
AUTONOMOUS_DB_BACKUP_CREATE AUTONOMOUS_DATABASE_CONTENT_READ	CreateAutonomousDatabaseBackup
AUTONOMOUS_DB_BACKUP_INSPECT AUTONOMOUS_DB_BACKUP_DELETE	DeleteAutonomousDatabaseBackup
AUTONOMOUS_DB_BACKUP_INSPECT AUTONOMOUS_DB_BACKUP_CONTENT_READ AUTONOMOUS_DATABASE_CONTENT_WRITE	RestoreAutonomousDatabase
Required on the source and the target compartment: AUTONOMOUS_DATABASE_UPDATE AUTONOMOUS_DB_BACKUP_INSPECT AUTONOMOUS_DB_BACKU_CREATE AUTONOMOUS_DATABASE_CONTENT_WRITE Required in both the source and the target compartment when Private Endpoint is enabled: WNIC_ASSOCIATE_NETOWRK_SECURITY_GROUP NETWORK_SECURITY_GROUP_UPDATE_MEMBERS	ChangeAutonomousDatabaseCompartment

Permissions Required to Use Operation	API Operation

Three possible cases:
• If Workload is `NULL`:
`AUTONOMOUS_DATABASE_UPDATE`
• If Workload is not `NULL`:
`AUTONOMOUS_DATABASE_CREATE`
`AUTONOMOUS_DATABASE_UPDATE`
• If Tagging is enabled:
`AUTONOMOUS_DATABASE_UPDATE`
`AUTONOMOUS_DATABASE_INSPECT`
	`UpdateAutonomousDatabase`: Use this API for changes or updates for any of the following operations:
	• set admin password (`adminPassword`)
	• auto start/stop schedule (`scheduledOperations`)
	• manage customer contacts (`customerContacts`)
	• edit tool configuration (`dbToolsDetails`)
	• update BYOL license options (`licenseModel` and `byolComputeCountLimit`)
	• update display name (`displayName`)
	• join an elastic pool
	• update elastic pool options
	• manage encryption keys
	• update to autonomous data guard for disaster recovery (`isLocalDataGuardEnabled` and `disasterRecoveryType`)
	• change database operation mode read/write read-only (`openMode`)
	• update network access with ACLs (`whitelistedIps`)
	• update network access with private endpoint (`privateEndpointLabel`)
	• rename database (`dbName`)
	• scale compute limits (`computeCount`)
	• manage compute auto scaling option (`isAutoScalingEnabled`)
	• scale storage limits (`dataStorageSizeInTBs`)
	• manage storage auto scaling options (`isAutoScalingForStorageEnabled`)
	• change workload type (`dbWorkload`)

requires `changeAutonomousDatabaseSubscription`	`ChangeAutonomousDatabaseSubscription`
requires `getSaasAdminUser`	`SaasAdminUserStatus`

requires `updateSaasAdminUser`	`ConfigureSaasAdminUser`
	`ListAutonomousDatabaseCharacterSets`
	`ListAutonomousDatabaseMaintenanceWindows`

Cloning Permissions

General IAM permissions are supported for Autonomous Database. In addition you can use `target.autonomous-database.cloneType` with the supported permission values to control the level of access, as shown in the following table.

target.autonomous-database.cloneType Value	Description
`CLONE-FULL`	Allow full clone only.

target.autonomous-database.cloneType Value	Description
CLONE-METADATA	Allow metadata clone only.
CLONE-REFRESHABLE	Allow refreshable clone only.
/CLONE*/	Allow any kind of clone.

Example policies with the supported `target.autonomous-database.cloneType` permission values:

```
Allow group group-name to manage autonomous-databases in compartment id
compartment-ocid
    where all {request.permission =  'AUTONOMOUS_DATABASE_CREATE',
target.autonomous-database.cloneType = 'CLONE-FULL'}
```

```
Allow group group-name to manage autonomous-databases in compartment id
compartment-ocid
    where all {request.permission =  'AUTONOMOUS_DATABASE_CREATE',
target.autonomous-database.cloneType = 'CLONE-METADATA'}
```

```
Allow group group-name to manage autonomous-databases in compartment id
compartment-ocid
    where all {request.permission =  'AUTONOMOUS_DATABASE_CREATE',
target.autonomous-database.cloneType = 'CLONE-REFRESHABLE'}
```

```
Allow group group-name to manage autonomous-databases in compartment id
compartment-ocid
    where all {request.permission =  'AUTONOMOUS_DATABASE_CREATE',
target.autonomous-database.cloneType = /CLONE*/}
```

See Permissions for more information.

Provide Specific Privileges in IAM Policies to Manage Autonomous Database

Lists IAM policies that you can use with an authorization verb and a condition to grant more granular operations to a group.

For example, to allow the group `MyGroup` to start Autonomous Databases using the `StartAutonomousDatabase` API:

```
Allow MyGroup to manage autonomous-databases where request.operation =
'StartAutonomousDatabase'
```

See Verbs and Conditions for more information.

Authorization Verb List
autonomousDatabaseManualRefresh
changeAutonomousDatabaseCompartment

Authorization Verb List
changeAutonomousDatabaseSubscription
changeDisasterRecoveryConfiguration
configureAutonomousDatabaseVaultKey
configureSaasAdminUser
createAutonomousDatabase
createAutonomousDatabaseBackup
deleteAutonomousDatabase
deleteAutonomousDatabaseBackup
deregisterAutonomousDatabaseDataSafe
disableAutonomousDatabaseManagement
disableAutonomousDatabaseOperationsInsights
enableAutonomousDatabaseManagement
enableAutonomousDatabaseOperationsInsights
failOverAutonomousDatabase
generateAutonomousDatabaseWallet
getAutonomousDatabase
getAutonomousDatabaseBackup
getAutonomousDatabaseRegionalWallet
getAutonomousDatabaseWallet
listAutonomousDatabaseBackups
listAutonomousDatabaseCharacterSets
listAutonomousDatabaseClones
listAutonomousDatabaseMaintenanceWindows
listAutonomousDatabasePeers
listAutonomousDatabaseRefreshableClones
listAutonomousDatabases
registerAutonomousDatabaseDataSafe
resourcePoolShapes
restartAutonomousDatabase
restoreAutonomousDatabase
rotateAutonomousDatabaseEncryptionKey
SaasAdminUserStatus
shrinkAutonomousDatabase
startAutonomousDatabase
stopAutonomousDatabase
switchoverAutonomousDatabase
updateAutonomousDatabase
updateAutonomousDatabaseBackup
updateAutonomousDatabaseRegionalWallet
updateAutonomousDatabaseWallet

The Authorization Verb `updateAutonomousDatabase` groups together privileges to use several API operations.

Operation
DeregisterAutonomousDatabaseDataSafe
DisableAutonomousDatabaseOperationsInsights
DisableDatabaseManagement
EnableAutonomousDatabaseOperationsInsights
RegisterAutonomousDatabaseDataSafe

For example:

```
Allow MyGroup to manage autonomous-databases where request.operation =
'updateAutonomousDatabase'
```

Use Oracle Data Safe with Autonomous Database

Provides information on using Oracle Data Safe on Autonomous Database.

- About Oracle Data Safe with Autonomous Database
  Oracle Data Safe, which is included with Autonomous Database, provides a unified control center that helps you manage the day-to-day security and compliance requirements of Oracle Databases.
- Register Autonomous Database with Oracle Data Safe
  To use Oracle Data Safe you first need to register your database with Oracle Data Safe.
- Use Oracle Data Safe Features
  After you register Autonomous Database with Oracle Data Safe you can use the Data Safe features.

About Oracle Data Safe with Autonomous Database

Oracle Data Safe, which is included with Autonomous Database, provides a unified control center that helps you manage the day-to-day security and compliance requirements of Oracle Databases.

Data Safe helps you to evaluate security controls, assess user security, monitor user activity, mitigate risk from compromised accounts, and address data security compliance requirements for your database. Data Safe accomplishes this by evaluating the sensitivity of your data and assisting you when you need to mask sensitive data for non-production databases.

Oracle Data Safe provides features to assist you when:

- Your organization's policies require that you monitor your databases and retain audit records.
- You need to protect against common database attacks coming from risks such as compromised accounts.
- Your developers need to use copies of production data for work on a new application and you're wondering what kinds of sensitive information the production data contains.
- You need to make sure that staff changes haven't left dormant user accounts on your databases.

Oracle Data Safe provides the following:

- **Security Assessment**: Configuration errors and configuration drift are significant contributors to data breaches. Use security assessment to evaluate your database's configuration and compare it to Oracle and industry best practices. Security assessment provides reports on areas of risk and notifies you when configurations change.
- **User Assessment**: Many breaches start with a compromised user account. User Assessment helps you spot the riskiest database accounts, those accounts which if compromised could cause the most damage. User Assessment helps you take proactive steps to secure these accounts. User Assessment Baselines make it easy to know when new accounts are added, or when an account's privileges are modified. You can use Oracle Cloud Infrastructure Events to receive proactive notifications when a database deviates from its baseline.
- **Data Discovery**: Provides support to locate and to manage sensitive data in your applications. Data discovery scans your database for over 150 different types of sensitive data and helps you to understand what types and how much sensitive data you are storing. Use the data discovery reports to formulate audit policies, develop data masking templates, and create effective access control policies.
- **Data Masking** Minimize the amount of sensitive data your organization maintains to help you meet compliance requirements and satisfy data privacy regulations. Data masking helps you remove risk from your non-production databases by replacing sensitive information with masked data. With reusable masking templates, over 50 included masking formats, and the ability to easily create custom formats for your organization's unique requirements, data masking can streamline your application development and testing operations.
- **Activity Auditing** Activity auditing collects audit records from databases and helps you manage audit policies. Understanding and reporting on user activity, data access, and changes to database structures supports regulatory compliance requirements and can aid in post-incident investigations. Audit insights make it easy to identify inefficient audit policies, while alerts based on audit data proactively notify you of risky activity.

  > **Note:**
  >
  > One (1) million audit records per database per month are included for your Autonomous Database if using the audit collection for Activity Auditing in Oracle Data Safe.

- **SQL Firewall Management** Protect against risks such as SQL injection attacks or compromised accounts. Oracle SQL Firewall is a security capability available with Oracle Database 23ai that offers best-in-class protection against these risks. The SQL Firewall feature in Oracle Data Safe lets you centrally manage and monitor the SQL Firewall policies for your target databases. Data Safe lets you collect authorized SQL activities of a database user, generate and enable the policy with allowlists of approved SQL statements and database connection paths and provides a comprehensive view of any SQL Firewall violations across the fleet of your target databases.

See Oracle Data Safe Overview for more information.

Register Autonomous Database with Oracle Data Safe

To use Oracle Data Safe you first need to register your database with Oracle Data Safe.

To get started, register your database:

[1] SQL Firewall is only available with Oracle Database 23ai.

1. Apply the necessary Identity and Access Management (IAM) permissions to register your target database.

   See Permissions to register an Autonomous Database for more information.

2. If you are registering an Autonomous Database that is configured to use a private IP address, then you need to create an Oracle Data Safe private endpoint either before or during registration.

   See Create an Oracle Data Safe Private Endpoint for more information.

3. Use the Oracle Data Safe Wizard to register your Autonomous Database instance.

   See Register an Autonomous Database for details on the registration steps.

Use Oracle Data Safe Features

After you register Autonomous Database with Oracle Data Safe you can use the Data Safe features.

Data Safe Feature	More Information
Security Assessment	Security Assessments are automatically scheduled once a week. Start by reviewing the security assessment report for your database: View the latest assessment for a target database. See Security Assessment Overview for more information.
User Assessment	User Assessments are automatically scheduled once a week. Start by reviewing the user assessment report for your database: View the latest user assessment for a target database See User Assessment Overview for more information.
Data Discovery	Start by discovering sensitive data in your database: Create Sensitive Data Models See Data Discovery Overview for more information.
Data Masking	To determine where sensitive data is stored in your database, run Data Discovery. After you know where sensitive data is stored in your database, you can create a masking policy: Create Masking Policies For example, after you create a masking policy you can make a copy of a production database and apply the masking policy to the non-production database: Mask Sensitive Data on a Target Database See Data Masking Overview for more information.
Activity Auditing	To use activity auditing, start the audit trail for your target database in Data Safe: Start an Audit Trail After the audit trail is started you can monitor and analyze your audit data with pre-defined audit reports: View a Predefined or Custom Audit Report See Activity Auditing Overview for more information.

Data Safe Feature	More Information
SQL Firewall	SQL Firewall[1] in Oracle Data Safe lets you centrally manage the SQL Firewalls and provides a comprehensive view of SQL Firewall violations across the fleet of your target databases. Data Safe lets you collect authorized SQL activities of a database user you wish to protect, monitor the progress of the collection, generate and enable the policy with allowlists of approved SQL statements and database connection paths. Start by enabling the SQL Firewall in your 23ai target database: Enable SQL Firewall On Your Target Database Next, you need to generate and enable a SQL Firewall policy with allowlists for the database user you wish to protect: Generate and Enforce SQL Firewall Policies. After you start enforcing the SQL Firewall policy, you can monitor and analyze the violations in the predefined violation reports: View and Manage Violations Reports. You can find more details on SQL Firewall at SQL Firewall Overview.

Break Glass Access for SaaS on Autonomous Database

Autonomous Database supports break glass access for SaaS providers. Break glass access allows a SaaS operations team, when explicitly authorized by a SaaS customer, to access a customer's database to perform critical or emergency operations.

- About Break Glass Access on Autonomous Database
  Break glass access on Autonomous Database supports SaaS providers, where the SaaS organization defines procedures to permit a SaaS operations team member to access a customer's database when they are explicitly authorized by the customer.
- Enable Break Glass Access
  After authorization to access a database with `SAAS_ADMIN` is approved through procedures defined by your organization, use the Autonomous Database CLI or API to enable the `SAAS_ADMIN` user.
- Disable Break Glass Access
  Use the Autonomous Database CLI or API to disable `SAAS_ADMIN` user access.
- Notes for Break Glass Access
  Provides notes for break glass access.

About Break Glass Access on Autonomous Database

Break glass access on Autonomous Database supports SaaS providers, where the SaaS organization defines procedures to permit a SaaS operations team member to access a customer's database when they are explicitly authorized by the customer.

Break Glass Sample Use Case with Example.com

Consider a SaaS provider named `example.com` that is using Autonomous Database for their product. In usual operations the SaaS provider, `example.com`, creates an Autonomous Database instance for each SaaS customer. In this model a SaaS customer, for example a customer named Scott, is an end-user for the `example.com` product (and a SaaS customer whose data is stored in an Autonomous Database instance). The provider `example.com` creates and stores all of Scott's data in an Autonomous Database instance, and the customer, Scott, has no direct access to the database.

This SaaS model is summarized as follows:

- The Oracle customer creating Autonomous Database instances is the SaaS organization, (`example.com`).
- The SaaS provider is `example.com`.
- The SaaS customer is Scott.

If and when something goes wrong with regards to application performance, or there is some other critical issue that requires troubleshooting by the SaaS operations team, the customer Scott, can grant access so that the operations team can access Scott's database for troubleshooting. The SaaS operations team is only authorized to establish direct access to Scott's Autonomous Database instance through a SaaS defined approval process (in other words, after `example.com` receives permission from their customer, Scott).

Break Glass and the Autonomous Database SAAS_ADMIN User

When a SaaS invokes the break glass API on a customer's Autonomous Database instance, this enables the `SAAS_ADMIN` user. The SaaS operations team can then access the instance using the `SAAS_ADMIN` user with a specified set of roles, for a limited time.
By default the `SAAS_ADMIN` user is locked. Using an approval process defined by the SaaS organization, the `SAAS_ADMIN` user can be enabled to allow access to an Autonomous Database instance. The break glass name comes from manual fire alarms that require their users to break a small glass window pane before activating the alarm (the glass must be broken to prevent the alarm from being triggered by mistake). Similarly, the `SAAS_ADMIN` user doesn't normally access the database and access requires a predefined approval process.

Depending on the type of access granted, when enabled, the `SAAS_ADMIN` user can access the database to investigate issues or to make changes associated with an emergency or other unusual event. When break glass access expires or when access is explicitly disabled, the `SAAS_ADMIN` account password/secrets are immediately rotated and `SAAS_ADMIN` user access is revoked. All the actions that the `SAAS_ADMIN` user performs are audited.

The `SAAS_ADMIN` user is enabled with one of three access types:

- `read-only`: provides read only access to the instance. This is the default access type and includes default roles: `CREATE SESSION`, `SELECT ANY TABLE`, `SELECT ANY DICTIONARY`, `SELECT_CATALOG_ROLE`.
- `read/write`: provides read/write access to the instance. The default roles for this type are: `CREATE SESSION`, `SELECT ANY TABLE`, `SELECT ANY DICTIONARY`, `SELECT_CATALOG_ROLE`, `INSERT ANY TABLE`, and `UPDATE ANY TABLE`.
- `admin`: provides admin access to the instance. The default roles for this type are: `CREATE SESSION` and `PDB_DBA`.

Break Glass API

The `SAAS_ADMIN` user is enabled and disabled only through the Command Line Interface (CLI) or using the Autonomous Database REST APIs.

For information about using the REST APIs and signing requests, see REST APIs and Security Credentials.

For information about SDKs, see Software Development Kits and Command Line Interface.

Use these APIs for Break Glass operations:

- To enable or disable `SAAS_ADMIN`, use configureSaasAdminUser.
- To check if the `SAAS_ADMIN` user is enabled use getSaasAdminUserStatus.

Enable Break Glass Access

After authorization to access a database with `SAAS_ADMIN` is approved through procedures defined by your organization, use the Autonomous Database CLI or API to enable the `SAAS_ADMIN` user.

You must have the manage autonomous database privilege to enable the `SAAS_ADMIN` user.

Before you enable the `SAAS_ADMIN` user to access a database you must obtain values for the required parameters.

Parameter	Description
`isEnabled`	Specifies a Boolean value. Use `TRUE` to enable.
`password`	Specifies the password for the `SAAS_ADMIN` user. If you specify `secretId` you cannot specify `password`. The password you provide as a parameter must conform to the Autonomous Database password requirements. See About User Passwords on Autonomous Database for more information.
`secretId`	Specifies the value of a secret's Oracle Cloud Infrastructure Vault secret OCID. If you specify `password` you cannot specify `secretId`. See Overview of Vault for more information. The password you provide as a secret in the Oracle Cloud Infrastructure Vault must conform to the Autonomous Database password requirements. See About User Passwords on Autonomous Database for more information.
`secretVersionNumber`	Specifies the version number of the secret specified with the `secretId`. This parameter is optional. The default is to use the latest secret version. This parameter only applies when `secretId` is also specified.
`accessType`	One of: `read-only`, `read/write`, or `admin`. The default is `read-only`.
`duration`	Specifies the duration in hours, in the range of 1 hour to 24 hours. The default is 1 hour.

To enable the `SAAS_ADMIN` user on an Autonomous Database instance you must define the required access using OCI Identity and Access Management policy statements written by an administrator.

The following policy is required:

```
Allow group Group_Name to manage autonomous-databases in compartment
Compartment_Name
```

See IAM Policies for Autonomous Database and Getting Started with Policies for more information.

Topics

- Enable Break Glass Access with a Password
  Use the Autonomous Database CLI or API to enable `SAAS_ADMIN` with a password.
- Enable Break Glass Access with a Vault Secret
  Use the Autonomous Database CLI or API to enable `SAAS_ADMIN` with a `secretId`, when the secret is stored in Oracle Cloud Infrastructure Vault.

Enable Break Glass Access with a Password

Use the Autonomous Database CLI or API to enable `SAAS_ADMIN` with a password.

1. Use the API or the CLI and specify a value for the password to enable `SAAS_ADMIN` with a `password`.

   For example:

   ```
   POST https://dbaas-api.svc.ad3.us-phoenix-1/20160918/autonomousDatabases/
   ocid1.autonomousdatabase.oc1.phx.uniqueId/actions/configureSaasAdminUser

   {  "isEnabled": true,
      "password": password,
      "accessType": "READ_ONLY",
      "duration": 17
   }
   ```

   See configureSaasAdminUser for more information.

2. Verify that `SAAS_ADMIN` user is enabled.

   ```
   POST https://dbaas-api.svc.ad3.us-phoenix-1/20160918/autonomousDatabases/
   ocid1.autonomousdatabase.oc1.phx.uniqueId/actions/getSaasAdminUserStatus

   {    "isEnabled": true,
        "accessType": "READ_ONLY",
        "timeSaasAdminUserEnabled": "2023-11-23T01:59:07.032Z"
   }
   ```

   This response shows that the `SAAS_ADMIN` user is enabled and that access type is `READ_ONLY`. The timestamp shows the time when `SAAS_ADMIN` was enabled (time is in UTC).

   See getSaasAdminUserStatus for more information.

Enable Break Glass Access with a Vault Secret

Use the Autonomous Database CLI or API to enable `SAAS_ADMIN` with a `secretId`, when the secret is stored in Oracle Cloud Infrastructure Vault.

When you specify a `secretId`, in order for Autonomous Database to reach the secret in the Oracle Cloud Infrastructure Vault, the following conditions must apply:

- The secret must be in `current` or `previous` state.
- You must have the proper user group policy that allows `READ` access to the specific secret in a given compartment. For example:

  ```
  Allow userGroup1 to read secret-bundles in compartment training
  ```

To enable `SAAS_ADMIN` with a `secretId` with the secret stored in Oracle Cloud Infrastructure Vault:

1. Use the API or the CLI and specify an OCID value for the `secretId`.

For example:

```
POST https://dbaas-api.svc.ad3.us-phoenix-1/20160918/autonomousDatabases/
ocid1.autonomousdatabase.oc1.phx.uniqueId/actions/configureSaasAdminUser

{  "isEnabled": true,
   "secretId": "ocid1.key.co1.ap-
mumbai-1.example..aaaaaaaauq5ok5nq3bf2vwetkpqsoa",
   "accessType": "READ_ONLY",
   "duration": 20
}
```

Specifying a secret version is optional. If you specify a secret version in the API call with `secretVersionNumber`, the specified secret version is used. If you do not specify a secret version the call uses the latest secret version.

See configureSaasAdminUser for more information.

2. Verify that `SAAS_ADMIN` user is enabled.

   For example:

```
POST https://dbaas-api.svc.ad3.us-phoenix-1/20160918/autonomousDatabases/
ocid1.autonomousdatabase.oc1.phx.uniqueId/actions/getSaasAdminUserStatus

{    "isEnabled": true,
     "accessType": "READ_ONLY",
     "timeSaasAdminUserEnabled": "2023-11-23T01:59:07.032Z"
}
```

   This response shows that the `SAAS_ADMIN` user is enabled and the access type is `READ_ONLY`. The timestamp shows the time when the user was enabled (time is in UTC).

   See getSaasAdminUserStatus for more information.

Disable Break Glass Access

Use the Autonomous Database CLI or API to disable `SAAS_ADMIN` user access.

By default access expires after one hour if the `duration` parameter is not set when `SAAS_ADMIN` is enabled. If the `duration` parameter is set when `SAAS_ADMIN` is enabled, then access expires after the specified `duration` number of hours. As an alternative to letting the access expire based on the either the default expiration time or the duration you specify, you can use configureSaasAdminUser to explicitly disable `SAAS_ADMIN` user access.

To disable the `SAAS_ADMIN` user on an Autonomous Database instance you must define the required access using OCI Identity and Access Management policy statements written by an administrator.

The following policy is required:

```
Allow group Group_Name to manage autonomous-databases in compartment
Compartment_Name
```

See IAM Policies for Autonomous Database and Getting Started with Policies for more information.

1. Disable `SAAS_ADMIN` user access.

   For example:

   ```
   POST https://dbaas-api.svc.ad3.us-phoenix-1/20160918/autonomousDatabases/
   ocid1.autonomousdatabase.oc1.phx.uniqueId/actions/configureSaasAdminUser
   {
     "isEnabled": false
   }
   ```

   See configureSaasAdminUser API for more information.

2. Verify that `SAAS_ADMIN` user is disabled.

   For example:

   ```
   POST https://dbaas-api.svc.ad3.us-phoenix-1/20160918/autonomousDatabases/
   ocid1.autonomousdatabase.oc1.phx.uniqueId/actions/getSaasAdminUserStatus

   {    "isEnabled": false
   }
   ```

   This response indicates that the `SAAS_ADMIN` user is disabled.

   See getSaasAdminUserStatus for more information.

When you disable the `SAAS_ADMIN` user, access to the database is revoked and Autonomous Database rotates the password or the secret that was used to access the database.

Notes for Break Glass Access

Provides notes for break glass access.

Notes for break glass access:

- The `duration` you specify when you enable `SAAS_ADMIN` applies either until the specified time expires or until you explicitly disable the `SAAS_ADMIN` user. You cannot change this value after you enable the `SAAS_ADMIN` user.
- Always Free Autonomous Database does not support access with the `SAAS_ADMIN` user.
- The `DBA_CLOUD_CONFIG` view provides information on the `SAAS_ADMIN` user.

  For example, the use the following query to obtain information on the status for the `SAAS_ADMIN` user:

  ```
  SELECT PARAM_VALUE FROM DBA_CLOUD_CONFIG WHERE
              param_name ='saas_admin_access' order by 1;
  ```

  The presence of a value for `auth_revoker` means the access has been revoked and shows the user who revoked access.

  The `auth_end` shows a `planned` time. After authorization is revoked, if the authorization expired at the time end of the `duration` specified when `SAAS_ADMIN` user was enabled, the `planned` time will be the same as the `actual` time. If the `planned` and the `actual` time are different, this means that `SAAS_ADMIN` authorization was revoked before the `duration` expired.

For example, if `SAAS_ADMIN` is enabled with a duration of 2 hours, and after 1 hour access is disabled by calling the API configureSaasAdminUser to disable the `SAAS_ADMIN` user, the `auth_end planned` and `actual` times will be different.

If this query shows no rows, then the `SAAS_ADMIN` user does not exist. It may have been created and dropped or it was never created.

Encrypt Data While Exporting or Decrypt Data While Importing

For greater security you can encrypt data that you export to Object Storage. When data on Object Storage is encrypted you can decrypt the data you import or when you use the data in an external table.

- Encrypt Data While Exporting to Object Storage
  You can encrypt table data while exporting to Object Storage.
- Decrypt Data While Importing from Object Storage
  You can decrypt and load data from encrypted files stored in Object Storage. You can also decrypt encrypted data on Object Storage that you use in an external table.

Encrypt Data While Exporting to Object Storage

You can encrypt table data while exporting to Object Storage.

Use the `format` parameter and the `encryption` option with `DBMS_CLOUD.EXPORT_DATA` to encrypt data when you export from Autonomous Database to Object Storage.

Note the following when you export encrypted data to Object Storage:

- The `encryption` option is only supported when exporting data from Autonomous Database to Object Storage as CSV, JSON, or XML.
- When the export includes both encryption and compression, the order of operations is: first the data is compressed, next the data is encrypted, and then it is uploaded to Object Storage.
- There are two supported encryption methods:
  - Using a user-defined function.
  - Using a `DBMS_CRYPTO` specified encryption algorithm.

    See DBMS_CRYPTO for information on the cryptographic functions and procedures for encryption and decryption.

Topics

- Encrypt Data Using DBMS_CRYPTO Encryption Algorithms
  Shows the steps to encrypt data using `DBMS_CRYPTO` encryption algorithms while exporting to Cloud Object Storage.
- Encrypt Data with a User Defined Encryption Function
  Shows the steps to encrypt data using a user-defined encryption function while exporting to Cloud Object Storage.

Encrypt Data Using DBMS_CRYPTO Encryption Algorithms

Shows the steps to encrypt data using `DBMS_CRYPTO` encryption algorithms while exporting to Cloud Object Storage.

Perform the following steps to encrypt data while exporting to Cloud Object Storage (this example exports table data to a CSV file):

1. Connect to your Autonomous Database instance.

   See Connect to Autonomous Database for more information.

2. Store your Cloud Object Storage credential using `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'OBJ_STORE_CRED',
       username        => 'user1@example.com',
       password        => 'password'
     );
   END;
   /
   ```

   The values you provide for `username` and `password` depend on the Cloud Object Storage service you are using.

   See CREATE_CREDENTIAL Procedure for more information.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

3. Create a credential to store the encryption key (the encryption key to be used for encrypting data).

   When you encrypt data using `DBMS_CRYPTO` encryption algorithms you store the encryption key in a credential. The key is specified in the `password` field in a credential you create with `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL (
       credential_name => 'ENC_CRED_NAME',
       username        => 'Any_username',
       password        => 'password'
     );
   END;
   /
   ```

   As an alternative you can create a credential to store the key in a vault. For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL (
       credential_name  => 'ENC_CRED_NAME',
       params           => JSON_OBJECT ('username' value 'Any_username',
   ```

```
                                        'region'    value 'Region',
                                        'secret_id' value 'Secret_id_value'));
END;
/
```

> **Note:**
>
> The `username` parameter you specify in the credential that stores the key can be any string.

This creates the `ENC_CRED_NAME` credential which is a vault secret credential, where the secret (decryption/encryption key) is stored as a secret in Oracle Cloud Infrastructure Vault.

See CREATE_CREDENTIAL Procedure for more information.

4. Run `DBMS_CLOUD.EXPORT_DATA`.

   Use the `format` parameter with the `encryption` option. The `encryption` type specifies the `DBMS_CRYPTO` encryption algorithm to use to encrypt the table data and the `credential_name` value is credential that specifies the secret (encryption key).

   For example:

```
BEGIN
    DBMS_CLOUD.EXPORT_DATA (
        credential_name => 'OBJ_STORE_CRED',
        file_uri_list   => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/encrypted.csv',
        query           => 'SELECT * FROM ADMIN.employees',
        format          => json_object(
               'type' value 'csv',
               'encryption' value  json_object(
                      'type' value DBMS_CRYPTO.ENCRYPT_AES256 +
DBMS_CRYPTO.CHAIN_CBC + DBMS_CRYPTO.PAD_PKCS5,
                      'credential_name' value 'ENC_CRED_NAME'))
      );
END;
/
```

   This encrypts and exports the data from the `EMPLOYEES` table into a CSV file.

   See DBMS_CRYPTO Algorithms for more information on encryption algorithms.

   In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

   See EXPORT_DATA Procedure and DBMS_CLOUD Package Format Options for EXPORT_DATA for more information.

After you encrypt files with `DBMS_CLOUD.EXPORT_DATA`, when you use `DBMS_CRYPTO` encryption algorithms to encrypt the files, you have these options for using or importing the files you exported:

- You can use `DBMS_CLOUD.COPY_DATA` or `DBMS_CLOUD.COPY_COLLECTION` with the same encryption algorithm options and the key to decrypt the files.

See Decrypt and Load Data Using DBMS_CRYPTO Algorithms for more information.

- You can query the data in an external table by supplying the same encryption algorithm options and the key to decrypt the files, with any of the following procedures:
  - `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`
  - `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE`
  - `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE`

    For `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE` this option is only applicable to the Object Storage files.

  See Decrypt and Load Data Using DBMS_CRYPTO Algorithms for more information.
- On a system that is not an Autonomous Database you can use the `DBMS_CRYPTO` package with the same algorithm options and the key to decrypt the files.

  Note that the key is stored as a `VARCHAR2` in the credential in Autonomous Database but `DBMS_CRYPTO` uses `RAW` type for the key parameter.

  See DBMS_CRYPTO Algorithms for more information on encryption algorithms.

Encrypt Data with a User Defined Encryption Function

Shows the steps to encrypt data using a user-defined encryption function while exporting to Cloud Object Storage.

Perform the following steps to encrypt data while exporting to Cloud Object Storage (this example exports table data to a CSV file):

1. Connect to your Autonomous Database instance.

   See Connect to Autonomous Database for more information.
2. Store your Cloud Object Storage credential using `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'OBJ_STORE_CRED',
    username        => 'user1@example.com',
    password        => 'password'
  );
END;
/
```

   The values you provide for `username` and `password` depend on the Cloud Object Storage service you are using.

   See CREATE_CREDENTIAL Procedure for more information.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.
3. Create a user-defined callback function to encrypt data.

For example:

```
CREATE OR REPLACE FUNCTION encryption_func (data IN BLOB)
  RETURN BLOB
  IS
      l_encrypted_data BLOB;
       BEGIN
     DBMS_LOB.CREATETEMPORARY (l_encrypted_data, TRUE, DBMS_LOB.CALL);
     DBMS_CRYPTO.ENCRYPT (
         dst => l_encrypted_data,
         src => data,
         typ => DBMS_CRYPTO.ENCRYPT_AES256 + DBMS_CRYPTO.CHAIN_CBC +
DBMS_CRYPTO.PAD_PKCS5,
         key => 'encryption key'
      );
     RETURN l_encrypted_data;
END encryption_func;
/
```

This creates the `ENCRYPTION_FUNC` encryption function. This function encrypts data using a stream or block cipher with a user supplied key.

> **Note:**
>
> You must create an encryption key to be used as a value in the `KEY` parameter. See DBMS_CRYPTO Operational Notes for more information on generating the encryption key.

4. Run `DBMS_CLOUD.EXPORT_DATA` with the `format` parameter, include the `encryption` option and specify a `user_defined_function`.

   For example:

```
BEGIN
      DBMS_CLOUD.EXPORT_DATA (
        credential_name => 'OBJ_STORE_CRED',
        file_uri_list   => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/encrypted.csv',
        query           => 'SELECT * FROM ADMIN.emp',
        format          => json_object(
                               'type' value 'csv',
                               'encryption' value
json_object('user_defined_function' value 'admin.encryption_func'))
      );
END;
/
```

This encrypts the data from the specified query the on `EMP` table and exports the data as a CSV file on Cloud Object Storage. The `format` parameter with the `encryption` value specifies the user-defined encryption function to use to encrypt the data.

> **Note:**
>
> You must have `EXECUTE` privilege on the encryption function.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

See EXPORT_DATA Procedure and DBMS_CLOUD Package Format Options for EXPORT_DATA for more information.

Decrypt Data While Importing from Object Storage

You can decrypt and load data from encrypted files stored in Object Storage. You can also decrypt encrypted data on Object Storage that you use in an external table.

This option is useful when migrating from an on-premises database to an Autonomous Database if the data in your source files is encrypted.

> **Note:**
>
> This option is only supported for Object Storage files less than 4 GB.

This option is applicable for the following procedures:

- `DBMS_CLOUD.COPY_DATA`
- `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`
- `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE`
- `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE`

  For `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE` this option is only applicable to the Object Storage files.
- `DBMS_CLOUD.COPY_COLLECTION`

Topics

- Decrypt and Load Data Using DBMS_CRYPTO Algorithms
  Shows the steps to decrypt encrypted files from Object Storage and load the data into a table on Autonomous Database (the decrypt step uses `DBMS_CRYPTO` algorithms).
- Decrypt and Load Data with a User Defined Function
  Shows the steps to decrypt files in Object Storage and load the data into tables using a user-defined decryption function.

Decrypt and Load Data Using DBMS_CRYPTO Algorithms

Shows the steps to decrypt encrypted files from Object Storage and load the data into a table on Autonomous Database (the decrypt step uses `DBMS_CRYPTO` algorithms).

As a prerequisite you must have encrypted files and uploaded the files into Object Storage. This example uses a CSV file and it is assumed that the file is encrypted using

`DBMS_CRYPTO.ENCRYPT_AES256 + DBMS_CRYPTO.CHAIN_CBC + DBMS_CRYPTO.PAD_PKCS5` algorithm and uploaded to your Cloud Object Storage.

See ENCRYPT Function for more information on the `ENCRYPT` function.

See DBMS_CRYPTO Operational Notes for more information on generating an encryption key.

To decrypt and load data into an existing table on Autonomous Database from Object Storage:

1. Connect to your Autonomous Database instance.

   See Connect to Autonomous Database for more information.

2. Store your Cloud Object Storage credential using `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'OBJ_STORE_CRED',
       username        => 'user1@example.com',
       password        => 'password'
     );
   END;
   /
   ```

   The values you provide for `username` and `password` depend on the Cloud Object Storage service you are using.

   See CREATE_CREDENTIAL Procedure for more information.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

3. Create a credential to store the key using `DBMS_CLOUD.CREATE_CREDENTIAL`. For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL (
       credential_name => 'ENC_CRED_NAME',
       username        => 'Any_username',
       password        => 'password'
     );
   END;
   /
   ```

   As an alternative you can create credentials to store the key in a vault. For example:

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL (
       credential_name  => 'ENC_CRED_NAME',
       params           => JSON_OBJECT ('username' value 'Any_username',
                                        'region'    value 'Region',
                                        'secret_id' value 'Secret_id_value'));
   END;
   /
   ```

> **Note:**
>
> The `username` parameter you specify in the credential that stores the key can be any string.

This creates the `ENC_CRED_NAME` credential which is a vault secret credential, where the secret (decryption/encryption key) is stored as a secret in the Oracle Cloud Infrastructure Vault.

See CREATE_CREDENTIAL Procedure for more information.

4. Run `DBMS_CLOUD.COPY_DATA` and specify `DBMS_CRYPTO` encryption algorithm as decryption method.

```
BEGIN
 DBMS_CLOUD.COPY_DATA (
   table_name       => 'CSV_COPY_DATA',
   credential_name  => 'OBJ_STORE_CRED',
   file_uri_list    => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/encrypted.csv',
   format           => json_object( 'type' value 'csv', 'encryption' value
json_object('type' value DBMS_CRYPTO.ENCRYPT_AES256 +
DBMS_CRYPTO.CHAIN_CBC + DBMS_CRYPTO.PAD_PKCS5,'credential_name' value
'ENC_CRED_NAME'))
  );
END;
/
```

This decrypts the `ENCRYPTED.CSV` file in the Object Storage. The data is then loaded into the `CSV_COPY_DATA` table. The `format` parameter `encryption` option value specifies a `DBMS_CRYPTO` encryption algorithm to be used to decrypt data.

See DBMS_CRYPTO Algorithms for more information on encryption algorithms.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

For detailed information about the parameters, see COPY_DATA Procedure.

For detailed information about the available `format` parameters, you can use with `DBMS_CLOUD.COPY_DATA`, see DBMS_CLOUD Package Format Options.

Decrypt and Load Data with a User Defined Function

Shows the steps to decrypt files in Object Storage and load the data into tables using a user-defined decryption function.

As a prerequisite for these steps you must have encrypted files and uploaded the files into Object Storage. This example uses a CSV file and it is assumed that the file is encrypted using `DBMS_CRYPTO.ENCRYPT_AES256 + DBMS_CRYPTO.CHAIN_CBC + DBMS_CRYPTO.PAD_PKCS5` algorithm and uploaded to your Cloud Object Storage.

See ENCRYPT Function for more information on the `ENCRYPT` function.

See DBMS_CRYPTO Operational Notes for more information on generating an encryption key.

To decrypt and load data into an existing table on Autonomous Database from Object Storage:

1. Connect to your Autonomous Database instance.

   See Connect to Autonomous Database for more information.

2. Store your Cloud Object Storage credential using `DBMS_CLOUD.CREATE_CREDENTIAL`.

   For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'OBJ_STORE_CRED',
    username        => 'user1@example.com',
    password        => 'password'
  );
END;
/
```

   The values you provide for `username` and `password` depend on the Cloud Object Storage service you are using.

   See CREATE_CREDENTIAL Procedure for more information.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

3. Create a user-defined function decryption callback function.

   For example:

```
CREATE OR REPLACE FUNCTION decryption_func (data IN BLOB)
RETURN BLOB
IS
  l_decrypted_data     BLOB;
  response             DBMS_CLOUD_TYPES.resp;
  l_response_json      JSON_ARRAY_T;
  l_response_json_obj  JSON_OBJECT_T;
  l_secret_id          VARCHAR2(4000) := '<secret_ocid>';
  l_key                VARCHAR2(4000);
BEGIN
   response := DBMS_CLOUD.send_request (
               credential_name => 'OCI$RESOURCE_PRINCIPAL',
               uri             => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/' ||
l_secret_id,
               method          => 'GET');
  l_response_json_obj := JSON_OBJECT_T.parse
(DBMS_CLOUD.get_response_text(response));
  l_key := UTL_RAW.cast_to_varchar2
(UTL_ENCODE.base64_decode(UTL_RAW.cast_to_raw
(l_response_json_obj.get_object('secretBundleContent').get_string('content'
))));

  DBMS_LOB.createtemporary (l_decrypted_data, TRUE, DBMS_LOB.CALL);
  DBMS_CRYPTO.decrypt (
      dst => l_decrypted_data,
      src => data,
```

```
      typ => DBMS_CRYPTO.ENCRYPT_AES256 + DBMS_CRYPTO.CHAIN_CBC +
DBMS_CRYPTO.PAD_PKCS5,
      key => l_key
    );
  RETURN l_decrypted_data;
END decryption_func;
/
```

This creates the `DECRYPTION_FUNC` decryption function. This function decrypts data using a stream or block cipher with a user supplied key. The user supplied key in the example is stored in Oracle Cloud Infrastructure Vault and is retrieved dynamically by making a REST call to Oracle Cloud Infrastructure Vault service.

4. Run `DBMS_CLOUD.COPY_DATA` and specify the `format` option `encryption` and specify the user-defined function you created to decrypt the data.

```
BEGIN
 DBMS_CLOUD.COPY_DATA (
    table_name      => 'CSV_COPY_DATA',
    credential_name => 'OBJ_STORE_CRED',
    file_uri_list   => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/encrypted.csv',
    format          => json_object(
                         'type' value 'csv',
                         'encryption' value
json_object('user_defined_function' value 'admin.decryption_func'))
  );
end;
/
```

This decrypts the `ENCRYPTED.CSV` file in the Object Storage. The data is then loaded into the `CSV_COPY_DATA` table. The `format` parameter `encryption` option value specifies a user-defined function name to use to decrypt data.

> **Note:**
>
> You must have the `EXECUTE` privilege on the user-defined function.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

For detailed information about the parameters, see COPY_DATA Procedure.

For detailed information about the available `format` parameters, you can use with `DBMS_CLOUD.COPY_DATA`, see DBMS_CLOUD Package Format Options.

Manage Oracle Cloud Operator Access

Describes how to grant temporary access to your database for an Oracle Cloud Operator.

Oracle Cloud Operators do not have authorization to access your data or any other information in your database schemas. When access to your database schemas is required to

troubleshoot or mitigate an issue, you can allow a cloud operator to access your Autonomous Database schemas for a limited time.

You allow a cloud operator to access the database schemas by running the procedure `DBMS_CLOUD_ADMIN.ENABLE_OPERATOR_ACCESS`. This means if you file a service request with Oracle Cloud Support and Oracle Cloud Operators need to access your database schemas, you must also enable operator access by running `DBMS_CLOUD_ADMIN.ENABLE_OPERATOR_ACCESS`.

Topics

- Enable Cloud Operator Access
  Describes the steps to enable Cloud Operator access on an Autonomous Database instance.
- Disable Cloud Operator Access
  Describes the steps to disable Cloud Operator access on an Autonomous Database instance.

Enable Cloud Operator Access

Describes the steps to enable Cloud Operator access on an Autonomous Database instance.

1. Enable operator access.

   For example:

   ```
   BEGIN
      DBMS_CLOUD_ADMIN.ENABLE_OPERATOR_ACCESS(
          auth_duration => 12 );
   END;
   /
   ```

   The `auth_duration` parameter specifies the number of Hours for which the Cloud Operator is granted access. The value must be a whole number in the range of 1 to 24 and the default is 1 (Hour).

   See ENABLE_OPERATOR_ACCESS Procedure for more information.

2. Verify that operator access is enabled.

   For example:

   ```
   SELECT param_name, param_value FROM DBA_CLOUD_CONFIG
         WHERE LOWER(param_name) = 'operator_access';


   PARAM_NAME PARAM_VALUE
   ---------------
   ------------------------------------------------------------------------
   -
   operator_access {"auth_grantor":"\"ADMIN\"","auth_begin":"26-FEB-24
   07.34.37.144846 PM UTC",
   "auth_duration":"12 hour", "planned_auth_end":"27-FEB-24 07.34.37.146297
   AM UTC"}
   ```

`DBMS_CLOUD_ADMIN.ENABLE_OPERATOR_ACCESS` allows access only to the Cloud Operator and does not enable access for any other user.

All operations performed by the Cloud Operator are stored in the view `DBA_OPERATOR_ACCESS`. See View Oracle Cloud Infrastructure Operations Actions for more information.

Disable Cloud Operator Access

Describes the steps to disable Cloud Operator access on an Autonomous Database instance.

1. Disable operator access.

   For example

   ```
   BEGIN
       DBMS_CLOUD_ADMIN.DISABLE_OPERATOR_ACCESS;
   END;
   /
   ```

   See DISABLE_OPERATOR_ACCESS Procedure for more information.

2. Verify that operator access is disabled.

   For example:

   ```
   SELECT param_name, param_value FROM DBA_CLOUD_CONFIG
         WHERE LOWER(param_name) = 'operator_access';


   No data found
   ```

Performance Monitor and Management

Oracle Autonomous Database Serverless includes several features that automatically monitor, analyze and optimize the performance of your Autonomous Database.

- Use OCI Ops Insights on Autonomous Database
  OCI Ops Insights is a cloud-native service that enables users to make informed, data-driven, Autonomous Database resource and performance management decisions.
- Manage Optimizer Statistics on Autonomous Database
  Describes Autonomous Database commands to run when you need to gather optimizer statistics or enable optimizer hints.
- Manage Automatic Indexing on Autonomous Database
  Automatic indexing automates the index management tasks in Autonomous Database. Auto Indexing is disabled by default in Autonomous Database.
- Manage Automatic Partitioning on Autonomous Database

Use OCI Ops Insights on Autonomous Database

OCI Ops Insights is a cloud-native service that enables users to make informed, data-driven, Autonomous Database resource and performance management decisions.

See Get Started with Ops Insights for information about enabling and disabling Ops Insights. In particular, see Enable Autonomous Databases & Full Feature Support for Autonomous Database configuration options.

Manage Optimizer Statistics on Autonomous Database

Describes Autonomous Database commands to run when you need to gather optimizer statistics or enable optimizer hints.

There are differences in the commands to run to gather optimizer statistics or enable optimizer hints, depending on your workload: Data Warehouse, Transaction Processing, or JSON Database.

- Manage Optimizer Statistics and Hints with Data Warehouse Workloads
  Describes Autonomous Database commands to run when you need to gather optimizer statistics or enable optimizer hints with Data Warehouse workloads.
- Manage Optimizer Statistics and Hints with Transaction Processing and JSON Database Workloads
  Describes Autonomous Database commands to run when you need to gather optimizer statistics or enable optimizer hints.

Manage Optimizer Statistics and Hints with Data Warehouse Workloads

Describes Autonomous Database commands to run when you need to gather optimizer statistics or enable optimizer hints with Data Warehouse workloads.

Manage Optimizer Statistics with Data Warehouse Workloads

Autonomous Database with Data Warehouse workloads gathers optimizer statistics automatically for tables loaded with direct path operations issued in SQL (direct path load operations that bypass the SQL data processing, such as SQL*Loader direct path, do not collect statistics). For example, for loads using the `DBMS_CLOUD` package the database gathers optimizer statistics automatically.

If you have tables modified using conventional DML operations you can run commands to gather optimizer statistics for those tables. For example, for the `SH` schema you can gather statistics for all tables in the schema using the following command:

```
BEGIN
  DBMS_STATS.GATHER_SCHEMA_STATS('SH', options=>'GATHER AUTO');
END;
/
```

This example gathers statistics for all tables that have stale statistics in the `SH` schema.

For more information about direct-path loads see *Loading Tables*.

For more information on optimizer statistics see *Database Concepts*.

Manage Optimizer Hints with Data Warehouse Workloads

Autonomous Database with Data Warehouse ignores optimizer hints and `PARALLEL` hints in SQL statements by default. If your application relies on hints you can enable optimizer hints by setting the parameter `OPTIMIZER_IGNORE_HINTS` to `FALSE` at the session or system level using `ALTER SESSION` or `ALTER SYSTEM`. For example, the following command enables hints in your session:

```
ALTER SESSION
   SET OPTIMIZER_IGNORE_HINTS=FALSE;
```

You can also enable `PARALLEL` hints in your SQL statements by setting `OPTIMIZER_IGNORE_PARALLEL_HINTS` to `FALSE` at the session or system level using `ALTER SESSION` or `ALTER SYSTEM`. For example, the following command enables `PARALLEL` hints in your session:

```
ALTER SESSION
   SET OPTIMIZER_IGNORE_PARALLEL_HINTS=FALSE;
```

Manage Optimizer Statistics and Hints with Transaction Processing and JSON Database Workloads

Describes Autonomous Database commands to run when you need to gather optimizer statistics or enable optimizer hints.

Manage Optimizer Statistics with Transaction Processing and JSON Database Workloads

Autonomous Database gathers optimizer statistics automatically so that you do not need to perform this task manually and this helps to ensure your statistics are current. Automatic statistics gathering is enabled in Autonomous Database and runs in a standard maintenance window.

> **Note:**
>
> The automatic statistics gathering maintenance window is different than the maintenance window on the Oracle Cloud Infrastructure console. The Oracle Cloud Infrastructure maintenance window shows system patching information.

For more information on automatic statistics gathering maintenance window times and automatic optimizer statistics collection, see *Database Administrator's Guide*.

For more information on optimizer statistics see *SQL Tuning Guide*.

Manage Optimizer Hints with Transaction Processing and JSON Database Workloads

Autonomous Database with Transaction Processing and JSON Database workloads honors optimizer hints and `PARALLEL` hints in SQL statements by default. You can disable optimizer hints by setting the parameter `OPTIMIZER_IGNORE_HINTS` to `TRUE` at the session or system level using `ALTER SESSION` or `ALTER SYSTEM`. For example, the following command disables hints in your session:

```
ALTER SESSION
   SET OPTIMIZER_IGNORE_HINTS=TRUE;
```

You can also disable `PARALLEL` hints in your SQL statements by setting `OPTIMIZER_IGNORE_PARALLEL_HINTS` to `TRUE` at the session or system level using `ALTER SESSION` or `ALTER SYSTEM`.

```
ALTER SESSION
   SET OPTIMIZER_IGNORE_PARALLEL_HINTS=TRUE;
```

Manage Automatic Indexing on Autonomous Database

Automatic indexing automates the index management tasks in Autonomous Database. Auto Indexing is disabled by default in Autonomous Database.

Creating indexes manually requires deep knowledge of the data model, application, and data distribution. In the past, DBAs were responsible for making choices about which indexes to create, and then sometimes the DBAs did not revise their choices or maintain indexes as the conditions changed. As a result, opportunities for improvement were lost, and use of unnecessary indexes could be a performance liability. The automatic indexing feature in Autonomous Database monitors the application workload and creates and maintains indexes automatically.

To enable automatic indexing:

1. Use the `DBMS_AUTO_INDEX.CONFIGURE` procedure to enable automatic indexing:

   ```
   EXEC DBMS_AUTO_INDEX.CONFIGURE('AUTO_INDEX_MODE','IMPLEMENT');
   ```

   This enables automatic indexing in a database and creates any new auto indexes as visible indexes, so that they can be used in SQL statements.
2. Use the `DBMS_AUTO_INDEX` package to report on the automatic task and to set automatic indexing preferences.

> **Note:**
>
> When automatic indexing is enabled, index compression for auto indexes is enabled by default.

To disable automatic indexing:

1. The following statement disables automatic indexing in a database so that no new auto indexes are created (existing auto indexes remain enabled):

   ```
   EXEC DBMS_AUTO_INDEX.CONFIGURE('AUTO_INDEX_MODE','OFF');
   ```

When you use SODA with Autonomous Database the following restrictions apply:

- Automatic indexing is not supported for SQL and PL/SQL code that uses the SQL/JSON function `json_exists`. See SQL/JSON Condition JSON_EXISTS for more information.
- Automatic indexing is not supported for SODA query-by-example (QBE).

See Managing Auto Indexes for more information.

Manage Automatic Partitioning on Autonomous Database

Automatic partitioning analyzes and automates partition creation for tables and indexes of a specified schema to improve performance and manageability in Autonomous Database.

Automatic partitioning, when applied, is transparent and does not require any user interaction or maintenance.

> **Note:**
>
> Automatic partitioning does not interfere with existing partitioning strategies and is complementary to manual partitioning in Autonomous Database. Manually partitioned tables are excluded as candidates for automatic partitioning.

- About Automatic Partitioning
  Automatic partitioning in Autonomous Database analyzes the application workload and automatically applies partitioning to tables and their indexes to improve performance or to allow better management of large tables.
- How Automatic Partitioning Works
- Configure Automatic Partitioning
  Use the `DBMS_AUTO_PARTITION.CONFIGURE` procedure to configure automatic partitioning options.
- Use Automatic Partitioning
  Describes the flow and general processes for using and managing automatic partitioning in Autonomous Database.
- Generate Automatic Partitioning Reports
  Generate automatic partitioning reports using the `REPORT_ACTIVITY` and `REPORT_LAST_ACTIVITY` functions of the `DBMS_AUTO_PARTITION` package.
- Example Automatic Partitioning Scenarios
  Example scenarios for automatic partitioning using the `DBMS_AUTO_PARTITION` API procedures and functions.
- Data Dictionary Views for Automatic Partitioning
  There are two new views and one updated view in the data dictionary for information about the automatic partitioning configuration and recommendations in your database.

About Automatic Partitioning

Automatic partitioning in Autonomous Database analyzes the application workload and automatically applies partitioning to tables and their indexes to improve performance or to allow better management of large tables.

Finding appropriate partitioning strategies requires deep knowledge of the application workload and the data distribution. When you perform manual partitioning, you must analyze your workload and make choices about how to apply partitioning to tables and indexes to improve the performance of applications. Automatic partitioning enables Autonomous Database users to benefit from partitioning without performing manual schema and workload analysis.

Automatic partitioning uses a single-column partition key combined with single-level partitioning. Automatic partitioning does not support more complex partitioning strategies such as multi-column partitioned tables or composite partitioning.

Automatic partitioning chooses from the following partition methods:

- AUTOMATIC INTERVAL: This choice is best suited for ranges of partition key values.
- LIST AUTOMATIC: This partitioning method applies to distinct partition key values.
- HASH: Applies partitioning on the partition key's hash values.

See Partitioning Concepts for more information.

Automatic partitioning provides the following functionality:

- Analyzes application workload and finds the optimal partitioning strategy to improve query performance for tables eligible for automatic partitioning.
- Provides PL/SQL APIs for configuring automatic partitioning in a database, generating reports about partitioning findings, and generating and applying an identified partitioning strategy for a given workload.

> **Note:**
>
> Automatic partitioning requires explicit calls to the `DBMS_AUTO_PARTITION` PL/SQL APIs to recommend and apply automatic partitioning to an Autonomous Database.

How Automatic Partitioning Works

Unlike automatic indexing, automatic partitioning does not run periodically as a background task. Automatic partitioning only runs when you invoke it using the `DBMS_AUTO_PARTITION.RECOMMEND_PARTITION_METHOD` function.

When invoked, automatic partitioning identifies candidate tables for automatic partitioning, evaluates partition schemes, and implements a partitioning strategy.

When you invoke automatic partitioning it performs the following tasks:

1. Identifies candidate tables for automatic partitioning by analyzing the workload for selected candidate tables.

   By default, automatic partitioning uses the workload information collected in an Autonomous Database for analysis. Depending on the size of the workload, a sample of queries might be considered.

2. Evaluates partition schemes based on workload analysis, quantification, and verification of the performance benefits:
   a. Candidate empty partition schemes with synthesized statistics are created internally and analyzed for performance.
   b. The candidate scheme with the highest estimated IO reduction is chosen as the optimal partitioning strategy and is internally implemented to test and verify performance.
   c. If a candidate partition scheme does not improve performance beyond specified performance and regression criteria, automatic partitioning is not recommended.
3. Implements the optimal partitioning strategy, if configured to do so, for the tables analyzed by the automatic partitioning procedures.

Configure Automatic Partitioning

Use the `DBMS_AUTO_PARTITION.CONFIGURE` procedure to configure automatic partitioning options.

Enable and implement recommendations

```
EXEC DBMS_AUTO_PARTITION.CONFIGURE('AUTO_PARTITION_MODE','IMPLEMENT');
```

Enable recommendations, but do not implement those recommendations

```
EXEC DBMS_AUTO_PARTITION.CONFIGURE('AUTO_PARTITION_MODE','REPORT ONLY');
```

Disable new recommendations and implementation of those new recommendations

```
EXEC DBMS_AUTO_PARTITION.CONFIGURE('AUTO_PARTITION_MODE','OFF');
```

> **Note:**
>
> This mode does not disable existing automatically partitioned tables.

Manage schemas and tables for automatic partitioning

Use the `AUTO_PARTITION_SCHEMA` and `AUTO_PARTITION_TABLE` settings to specify schemas and tables considered for automatic partitioning.

> **Note:**
>
> When automatic partitioning is invoked, all schemas and tables in user-managed schemas are considered for automatic partitioning if both the inclusion and exclusion lists are empty.

- Assuming the inclusion list and the exclusion list are empty, add the `HR` schema and the `SH.SALES` table to the exclusion list, preventing only those objects from automatic partitioning analysis.

```
BEGIN
    DBMS_AUTO_PARTITION.CONFIGURE(
         PARAMETER_NAME  => 'AUTO_PARTITION_SCHEMA',
         PARAMETER_VALUE => 'HR',
         ALLOW           => FALSE);

    DBMS_AUTO_PARTITION.CONFIGURE(
         PARAMETER_NAME  => 'AUTO_PARTITION_TABLE',
         PARAMETER_VALUE => 'SH.SALES',
         ALLOW           => FALSE);
END;
/
```

- After the previous example runs, use the following to remove the `HR` schema from the exclusion list.

```
BEGIN
    DBMS_AUTO_PARTITION.CONFIGURE(
         PARAMETER_NAME  => 'AUTO_PARTITION_SCHEMA',
         PARAMETER_VALUE => 'HR',
         ALLOW           => NULL);
```

```
END;
/
```

- Use the following command to remove all schemas from the exclusion list.

```
BEGIN
    DBMS_AUTO_PARTITION.CONFIGURE(
         PARAMETER_NAME  => 'AUTO_PARTITION_SCHEMA',
         PARAMETER_VALUE => NULL,
         ALLOW           => TRUE);
END;
/
```

- Assuming the inclusion and exclusion lists are empty, the following example adds the `HR` schema to the inclusion list. As soon as the inclusion list is no longer empty, only schemas in the inclusion list are considered.

  With this example, only the `HR` schema is a candidate for automatic partitioning.

```
BEGIN
    DBMS_AUTO_PARTITION.CONFIGURE(
         PARAMETER_NAME  => 'AUTO_PARTITION_SCHEMA',
         PARAMETER_VALUE => 'HR',
         ALLOW           => TRUE);
END;
/
```

Manage Automatic Partitioning Report Retention Period

Set the retention period for automatic partitioning reports to 365 days.

```
BEGIN
    DBMS_AUTO_PARTITION.CONFIGURE(
         PARAMETER_NAME  => 'AUTO_PARTITION_REPORT_RETENTION',
         PARAMETER_VALUE => '365');
END;
/
```

See CONFIGURE Procedure for more information.

Use Automatic Partitioning

Describes the flow and general processes for using and managing automatic partitioning in Autonomous Database.

1. Choose the database for automatic partitioning.
   In general, Oracle recommends using automatic partitioning in cloned or manually created databases rather than production databases. The analysis and verification of automatic partitioning using `RECOMMEND_PARTITION_METHOD` is potentially a resource-intensive and long running operation that can add undesirable processing to your database.

   To use a secondary database for automatic partitioning analysis, the database must have information about your workload in an internally managed SQL workload repository.

   a. Use a cloned database for automatic partitioning.

Autonomous Database automatically collects workload information over time in an internally managed SQL workload repository maintained in the SQL Tuning Set (`SYS_AUTO_STS`). If you clone your production database after having run the workload for a while, the clone will have the necessary workload information. You can use automatic partitioning with such clones without any additional actions.

See Clone an Autonomous Database Instance

b. Use other databases for automatic partitioning.
You can run your workload manually to collect the necessary workload information. If you manually run your workload prior to using automatic partitioning, any Autonomous Database that contains your desired schemas and data can be used for automatic partitioning after your workload was run, regardless of whether it is cloned or manually created.

2. Recommend automatic partitioning.
Use `RECOMMEND_PARTITION_METHOD` to analyze your database, specific schemas, or specific tables to identify the optimal partitioning strategy, if any. The recommendation analyzes your workload and schemas verifying performance benefits by running your workload against an internally created auxiliary table. This can be a resource-intensive and long running operation, requiring CPU and IO to create the auxiliary table and verify performance. You will also temporarily need additional space, of 1 - 1.5 times, your largest candidate table.

3. Apply the recommendation.
Any recommendation can be implemented with the `APPLY_RECOMMENDATION` procedure in the database where the recommendation analysis occurred. Alternatively, any recommendation can be extracted from the database used for analysis and applied to any database, such as a production system. The script needed for manual modification is stored in column `MODIFY_TABLE_DDL` in the `DBA_AUTO_PARTITION_RECOMMENDATION` view.

   Oracle recommends applying automatic partitioning to your database at off-peak time. While your tables will be modified to automatically partitioned tables, the conversion adds additional resource requirements to your system, such as additional CPU and IO. Automatic partitioning requires as much as 1.5 times the size of the table to being modified as additional free space, depending on concurrent ongoing DML operations on those tables.

Generate Automatic Partitioning Reports

Generate automatic partitioning reports using the `REPORT_ACTIVITY` and `REPORT_LAST_ACTIVITY` functions of the `DBMS_AUTO_PARTITION` package.

Generate a report, in plain text format, of automatic partitioning operations for a specific period

This example generates a report containing typical information about the automatic partitioning operations for the last 24 hours. The report is generated in plain text format by default.

```
DECLARE
  Report clob := NULL
BEGIN
    Report := DBMS_AUTO_PARTITION.REPORT_ACTIVITY();
END;
/
```

Generate a report, in HTML format, of automatic partitioning operations for MAY 2021

This example generates a report containing basic information about the automatic partitioning operations for the month of MAY 2021. The report is generated in the HTML format, and it includes only a summary of automatic partitioning operations.

```
DECLARE
  Report clob := NULL
BEGIN
    Report := DBMS_AUTO_PARTITION.REPORT_ACTIVITY(
                      ACTIVITY_START => TO_TIMESTAMP('2021-05-01', 'YYYY-MM-
DD'),
                      ACTIVITY_END   => TO_TIMESTAMP('2021-06-01', 'YYYY-MM-
DD'),
                      TYPE           => 'HTML',
                      SECTION        => 'SUMMARY',
                      LEVEL          => 'BASIC' );
END;
/
```

Generate a report, in plain text format, of the last automatic partitioning operation

This example generates a report containing typical information about the last automatic partitioning operation. The report is generated in the plain text format by default.

```
DECLARE
  Report clob := NULL
BEGIN
    Report := DBMS_AUTO_PARTITION.REPORT_LAST_ACTIVITY();
END;
/
```

See REPORT_ACTIVITY Function for more information.

See REPORT_LAST_ACTIVITY Function for more information.

Example Automatic Partitioning Scenarios

Example scenarios for automatic partitioning using the `DBMS_AUTO_PARTITION` API procedures and functions.

Generate a recommendation for a single table and manually apply the recommendation

1. Set `AUTO_PARTITION_MODE` parameter to `REPORT ONLY` to enable an automatic partitioning recommendation to be made and verified. The recommendation is not applied to the table.

   ```
   BEGIN
       DBMS_AUTO_PARTITION.CONFIGURE(
           PARAMETER_NAME  => 'AUTO_PARTITION_MODE',
           PARAMETER_VALUE => 'REPORT ONLY');
   END;
   /
   ```

2. Validate that `TPCH.LINEITEM` table is a candidate for automatic partitioning. This step is optional and recommended when you are selectively targeting single tables.

```
SELECT DBMS_AUTO_PARTITION.VALIDATE_CANDIDATE_TABLE(
           TABLE_OWNER => 'TPCH',
           TABLE_NAME  => 'LINEITEM')
FROM DUAL;
```

   If the table is a valid candidate, when you invoke automatic partitioning for a recommendation analysis it returns as `VALID`. Otherwise, the violation criteria is shown.

   See VALIDATE_CANDIDATE_TABLE Function for a list of criteria for eligible candidate tables.

3. Invoke the `DBMS_AUTO_PARTITION` API to generate a recommendation for the `TPCH.LINEITEM` table.

```
-- DEFINE SQLPLUS BIND VARIABLE FOR RECOMMENDATION ID
VARIABLE RECOMMENDATION_ID VARCHAR2(32);
BEGIN
  :RECOMMENDATION_ID := DBMS_AUTO_PARTITION.RECOMMEND_PARTITION_METHOD(
                            TABLE_OWNER => 'TPCH',
                            TABLE_NAME  => 'LINEITEM');
END;
/
```

   The recommendation analysis and verification that you perform with `DBMS_AUTO_PARTITION.RECOMMEND_PARTITION_METHOD` can be a resource-intensive and long running operation and might take considerable time. You should perform this step on a database that is not your primary production system. Oracle recommends giving the verification operation sufficient resources by choosing the HIGH service.

4. Check the recommendation. The view `DBA_AUTO_PARTITION_RECOMMENDATIONS` contains the information on the recommendation. In this example, check the recommended partition key and partition method.

```
SELECT PARTITION_METHOD, PARTITION_KEY
  FROM DBA_AUTO_PARTITION_RECOMMENDATIONS
  WHERE RECOMMENDATION_ID = :RECOMMENDATION_ID;
```

   Additionally, query the same view to get the performance analysis report generated for the workload after the table was partitioned according to the recommendation.

```
SELECT REPORT
    FROM DBA_AUTO_PARTITION_RECOMMENDATIONS
    WHERE RECOMMENDATION_ID = :RECOMMENDATION_ID;
```

5. After manual validation of the recommendation, apply the recommendation. If you are applying the recommendation in the database where the recommendation analysis has taken place, apply the recommendation by executing the `APPLY_RECOMMENDATION` procedure.

```
BEGIN
    DBMS_AUTO_PARTITION.APPLY_RECOMMENDATION(
         RECOMMENDATION_ID => :RECOMMENDATION_ID);
```

```
END;
/
```

If you want to apply the recommendation to a different database, such as your production environment, extract the modification DDL. Then, run the extracted modification DDL in your target database. The query to extract the modification DDL is as follows:

```
SELECT MODIFY_TABLE_DDL
    FROM DBA_AUTO_PARTITION_RECOMMENDATIONS
    WHERE RECOMMENDATION_ID = :RECOMMENDATION_ID;
```

Example output of modification DDL:

```
BEGIN
    -- DBMS_AUTO_PARTITION RECOMMENDATION_ID
C3F7A59E085C2F25E05333885A0A55EA
    --   FOR TABLE "TPCH"."LINEITEM"
    --   GENERATED AT 06/04/2021 20:52:29
    DBMS_AUTO_PARTITION.BEGIN_APPLY(EXPECTED_NUMBER_OF_PARTITIONS => 10);

    EXECUTE IMMEDIATE
    'ALTER TABLE "TPCH"."LINEITEM"
    MODIFY PARTITION BYLIST(SYS_OP_INTERVAL_HIGH_BOUND
        ("L_SHIPDATE", INTERVAL ''10'' MONTH, TIMESTAMP ''1992-01-01
00:00:00''))
    AUTOMATIC /* SCORE=23533.11;  */
    (PARTITION P_NULL VALUES(NULL))
    AUTO ONLINE PARALLEL';

    DBMS_AUTO_PARTITION.END_APPLY;
EXCEPTION WHEN OTHERS THEN
    DBMS_AUTO_PARTITION.END_APPLY;
    RAISE;
END;
```

6. Verify that the table was automatically partitioned, query the catalog views.

```
SELECT T.AUTO, T.PARTITIONING_TYPE, C.COLUMN_NAME
    FROM DBA_PART_TABLES T, DBA_PART_KEY_COLUMNS C
    WHERE T.OWNER = 'TPCH' AND T.TABLE_NAME = 'LINEITEM'
        AND T.OWNER = C.OWNER AND T.TABLE_NAME = C.NAME;
```

Use this query to identify when automatic partitioning was applied to a given table.

```
SELECT APPLY_TIMESTAMP_START, APPLY_TIMESTAMP_END
    FROM DBA_AUTO_PARTITION_RECOMMENDATIONS
    WHERE TABLE_OWNER = 'TPCH' AND TABLE_NAME = 'LINEITEM';
```

See CONFIGURE Procedure for information.

See VALIDATE_CANDIDATE_TABLE Function for information.

See RECOMMEND_PARTITION_METHOD Function for information.

See APPLY_RECOMMENDATION Procedure for information.

Generate a recommendation for eligible tables and manually apply the recommendation

1. Set `AUTO_PARTITION_MODE` parameter to `REPORT ONLY` to enable an automatic partitioning recommendation to be made and verified. The recommendation is not applied to existing tables.

```
BEGIN
    DBMS_AUTO_PARTITION.CONFIGURE(
        PARAMETER_NAME  => 'AUTO_PARTITION_MODE',
        PARAMETER_VALUE => 'REPORT ONLY');
END;
/
```

2. Invoke the `DBMS_AUTO_PARTITION` API to generate a recommendation table.

```
-- DEFINE SQLPLUS BIND VARIABLE FOR RECOMMENDATION ID
VARIABLE RECOMMENDATION_ID VARCHAR2(32);
BEGIN
  :RECOMMENDATION_ID := DBMS_AUTO_PARTITION.RECOMMEND_PARTITION_METHOD();
END;
/
```

   The recommendation analysis and verification is a resource-intensive and long running operation and might take considerable time. On secondary, non-production databases, Oracle recommends giving the verification sufficient resources by choosing service HIGH.

3. Query the `DBA_AUTO_PARTITION_RECOMMENDATIONS` view to see which tables were analyzed.

```
SELECT TABLE_OWNER, TABLE_NAME, PARTITION_METHOD, PARTITION_KEY
    FROM DBA_AUTO_PARTITION_RECOMMENDATIONS
    WHERE RECOMMENDATION_ID = :RECOMMENDATION_ID
    ORDER BY RECOMMENDATION_SEQ;
```

4. Use this query to drill-down in the report for a specific table that was analyzed in the run, the `TPCH.LINEITEM` table in this example.

```
SELECT REPORT
    FROM DBA_AUTO_PARTITION_RECOMMENDATIONS
    WHERE RECOMMENDATION_ID = :RECOMMENDATION_ID
        AND TABLE_OWNER = 'TPCH'
        AND TABLE_NAME  = 'LINEITEM';
```

5. Apply the recommendation by executing the `APPLY_RECOMMENDATION` procedure.

```
BEGIN
    DBMS_AUTO_PARTITION.APPLY_RECOMMENDATION(
         RECOMMENDATION_ID => :RECOMMENDATION_ID);
END;
/
```

Alternately, apply recommendations for a specific table that was analyzed, the `TPCH.LINEITEM` table in this example.

```
BEGIN
    DBMS_AUTO_PARTITION.APPLY_RECOMMENDATION(
        RECOMMENDATION_ID => :RECOMMENDATION_ID,
        TABLE_OWNER       => 'TPCH',
        TABLE_NAME        => 'LINEITEM');
END;
/
```

See CONFIGURE Procedure for information.

See RECOMMEND_PARTITION_METHOD Function for information.

See APPLY_RECOMMENDATION Procedure for information.

> **Note:**
>
> Recommendations of automatic partitioning generated by the `RECOMMEND_PARTITION_METHOD` function have a time limit, specified by the `TIME_LIMIT` parameter, with a default of 1 day. If you are analyzing a large system with many candidate tables, a single invocation may not generate a recommendation for all tables that can be partitioned. You can safely invoke the recommendation for auto partitioning repeatedly to generate recommendations for additional tables. When the function is invoked and zero rows are in `DBA_AUTO_PARTITION_RECOMMENDATIONS` for the `RECOMMENDATION_ID`, then the function did not find any additional candidate tables for automatic partitioning.

Generate and automatically apply recommendations for eligible tables

1. Set `AUTO_PARTITION_MODE` parameter to `REPORT ONLY` to enable an automatic partitioning recommendation to be made and verified. The recommendation is not applied to existing tables.

   ```
   BEGIN
       DBMS_AUTO_PARTITION.CONFIGURE(
           PARAMETER_NAME  => 'AUTO_PARTITION_MODE',
           PARAMETER_VALUE => 'IMPLEMENT');
   END;
   /
   ```

2. Invoke the `DBMS_AUTO_PARTITION` API to generate a recommendation table.

   ```
   -- DEFINE SQLPLUS BIND VARIABLE FOR RECOMMENDATION ID
   VARIABLE RECOMMENDATION_ID VARCHAR2(32);
   BEGIN
     :RECOMMENDATION_ID := DBMS_AUTO_PARTITION.RECOMMEND_PARTITION_METHOD();
   END;
   /
   ```

   The recommendation analysis and verification is a resource-intensive and long running operation and might take considerable time. On secondary, non-production databases, Oracle recommends giving the verification sufficient resources by choosing service HIGH.

3. Query the `DBA_AUTO_PARTITION_RECOMMENDATIONS` view to see which tables were analyzed.

```
SELECT TABLE_OWNER, TABLE_NAME, PARTITION_METHOD, PARTITION_KEY
    FROM DBA_AUTO_PARTITION_RECOMMENDATIONS
    WHERE RECOMMENDATION_ID = :RECOMMENDATION_ID
    ORDER BY RECOMMENDATION_SEQ;
```

4. Use the `REPORT_LAST_ACTIVITY` function to retrieve the report on the actions taken during the last run.

```
SELECT DBMS_AUTO_PARTITION.REPORT_LAST_ACTIVITY() FROM DUAL;
```

See CONFIGURE Procedure for information.

See RECOMMEND_PARTITION_METHOD Function for information.

See REPORT_LAST_ACTIVITY Function for information.

Data Dictionary Views for Automatic Partitioning

There are two new views and one updated view in the data dictionary for information about the automatic partitioning configuration and recommendations in your database.

- DBMS_AUTO_PARTITION DBA_AUTO_PARTITION_CONFIG View
  Displays the current configuration parameter settings for automatic partitioning.
- DBMS_AUTO_PARTITION DBA_AUTO_PARTITION_RECOMMENDATIONS View
  When you run `CONFIGURE` or `RECOMMEND_PARTITION`, the results from those procedures is stored in this view. The `RECOMMENDATION_ID` is used in several procedures and functions.
- DBMS_AUTO_PARTITION Updates to Existing Views
  Discusses the changes to existing views as a result of the implementation of automatic partitioning.

DBMS_AUTO_PARTITION DBA_AUTO_PARTITION_CONFIG View

Displays the current configuration parameter settings for automatic partitioning.

Column	Description
`PARAMETER_NAME`	Name of the configuration parameter
`PARAMETER_VALUE`	Value of the configuration parameter
`LAST_MODIFIED`	Time, in UTC, at which the parameter value was last modified.
`MODIFIED_BY`	User who last modified the parameter value

DBMS_AUTO_PARTITION DBA_AUTO_PARTITION_RECOMMENDATIONS View

When you run `CONFIGURE` or `RECOMMEND_PARTITION`, the results from those procedures is stored in this view. The `RECOMMENDATION_ID` is used in several procedures and functions.

Column	Description
`TABLE_OWNER`	Owner of the table
`TABLE_NAME`	Owner of the table

Column	Description
PARTITION_METHOD	Recommended partition method. See CONFIGURE Procedure
PARTITION_KEY	Recommended partition key. NULL means that analysis complete and recommendation is not to partition the table.
GENERATE_TIMESTAMP	Time, in UTC, when this recommendation was generated.
RECOMMENDATION_ID	ID used with DBMS_AUTO_PARTITION APIs to get additional information about this recommendation.
RECOMMENDATION_SEQ	When a recommendation ID has recommendations for multiple tables, provides the order in which the recommendations were generated. Performance reports are generated assuming that earlier recommendations have been applied. For example, the report for RECOMMENDATION_SEQ = 2 assumes recommendations have been applied for both RECOMMENDATION_SEQ = 1 and RECOMMENDATION_SEQ = 2.
MODIFY_TABLE_DDL	DDL that would be, or was, used to apply the recommendation.
APPLY_TIMESTAMP_START	Time, in UTC, when application of this recommendation was started. NULL if recommendation was not applied.
APPLY_TIMESTAMP_END	Time, in UTC, when application of this recommendation was finished. NULL if recommendation was not applied or if application has not finished.
REPORT	SQL Performance Analyzer report from SQL execution on database after recommendation is applied.

DBMS_AUTO_PARTITION Updates to Existing Views

Discusses the changes to existing views as a result of the implementation of automatic partitioning.

In *_PART_TABLES, the AUTO column (VARCHAR2(3)) was added. Its values are as follows:

- YES - If the table was partitioned by DBMS_AUTO_PARTITION.
- NO - If the table was not partitioned by DBMS_AUTO_PARTITION.

Using Oracle Graph with Autonomous Database

Oracle Graph with Autonomous Database enables you to create graphs from data in your Autonomous Database. With graphs you can analyze your data based on connections and relationships between data entities.

As an Analyst or a Developer you can use graph algorithms and graph pattern queries for ranking, clustering, and path analysis in a graph model of your data. You can use graph features to detect anomalous patterns, identify communities, and find new connections in your data. Then, you can use graphs in your applications, for example, for fraud detection in banking, improved traceability in smart manufacturing, building linked data applications, and more; all while gaining enterprise-grade security, ease of data ingestion, and support for a wide range of workloads.

Autonomous Database includes Graph Studio, which automates graph data management and simplifies modeling, analysis, and visualization across the graph analytics lifecycle. Autonomous Database includes all the graph capabilities from Oracle Database, except the following:

- When using RDF graphs, querying external data sources using SPARQL federated queries is not supported.
- Triple level security in RDF graphs using Oracle Label Security are not supported. See Fine-Grained Access Control for RDF Data for more information.

Database roles must be set to use Graph Studio on an Autonomous Database instance. See Required Roles to Access Tools from Database Actions for more information.

- About Oracle Graph Studio on Autonomous Database
- Building Applications Using Graph APIs with Autonomous Database
  If you are developing an application using graphs, you can also access Property Graph APIs and RDF Graph APIs with Autonomous Database without Graph Studio.

About Oracle Graph Studio on Autonomous Database

Graph Studio allows you to work with the two popular graph models, property graphs and RDF graphs.

Graph Studio features include automated modeling to create graphs from database tables, an integrated notebook to run graph queries and analytics, and native graph and other visualizations. You can easily create and manage either of the graph models. You can validate the graphs by executing queries and exploring their properties.

See Graph Studio: Interactive, Self-Service User Interface and Access Graph Studio Using Oracle Cloud Infrastructure Console for more information on Graph Studio.

Notes for Graph Studio on Autonomous Database:

- Graph Studio is supported with Data Warehouse and Transaction Processing workloads.

Building Applications Using Graph APIs with Autonomous Database

If you are developing an application using graphs, you can also access Property Graph APIs and RDF Graph APIs with Autonomous Database without Graph Studio.

- **Property Graph**: Create and work with Property Graphs with Python and Java developer APIs and the Property Graph Query Language (PGQL). You can use developer APIs to create a graph, run graph queries, and execute graph algorithms.

  You can use these APIs by deploying Oracle Graph Server and Client from OCI Marketplace.

  See Oracle Database Graph for more information.
- **RDF Graph**: Using RDF Graph Server and Query UI you can manage, store, query, and perform inferencing on RDF graphs in Autonomous Database. You can use a REST API, a SPARQL endpoint, a web UI to run SPARQL queries, and developer APIs for advanced RDF graph data management operations.

  You can use these APIs by deploying Oracle RDF Graph Server and Query UI from OCI Marketplace.

  See Using RDF Graph Server and Query UI for more information.

Use Oracle Machine Learning with Autonomous Database

Oracle Machine Learning components are integrated into Autonomous Database.

One of the Oracle Machine Learning components, Oracle Machine Learning Notebooks is a notebook style application designed for advanced SQL users. Oracle Machine Learning Notebooks provides interactive data analysis that lets you develop, document, share, and automate reports based on sophisticated analytics and data models.
Key features of Oracle Machine Learning Notebooks:

- Allows collaboration among data scientists, developers, business users
- Leverages the scalability and performance of Oracle Platform and its Cloud Services

To use Oracle Machine Learning Notebooks with Autonomous Database you need to add or create Oracle Machine Learning users and then access Oracle Machine Learning Notebooks:

- **OML User Management** lets the **Admin** (user with administrative privileges) create and modify Oracle Machine Learning user accounts. See Create and Update User Accounts for Oracle Machine Learning Components on Autonomous Database for details on accessing **OML User Management**.
- **OML Application** – application **users** access Oracle Machine Learning Notebooks to create, view, and share notebooks for data analytics, data visualization, and other Oracle Machine Learning Notebooks tasks. Users access Oracle Machine Learning Notebooks by clicking the Oracle Machine Learning Home icon on the User Administration page or from the account details mailed to a new user. See Work with Oracle Machine Learning User Interface for Data Access, Analysis, and Discovery for details on using the OML application.

See Machine Learning with Autonomous Database for more information.

- Machine Learning with Autonomous Database
  Oracle Machine Learning enables you to solve key enterprise business problems and accelerates the development and deployment of data science and machine learning-based solutions using scalable, automated, and secure machine learning for model building, evaluation, and deployment.

Machine Learning with Autonomous Database

Oracle Machine Learning enables you to solve key enterprise business problems and accelerates the development and deployment of data science and machine learning-based solutions using scalable, automated, and secure machine learning for model building, evaluation, and deployment.

Certain Oracle Machine Learning products are available on specific Oracle Database platforms.

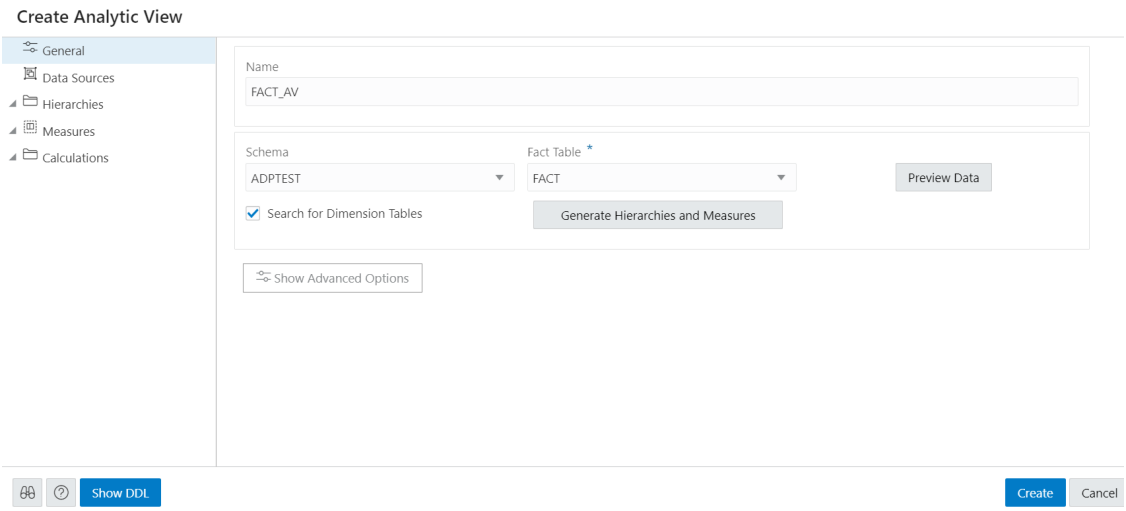
Product Name	Description	Available on:
OML Notebooks	Oracle Machine Learning Notebooks is a collaborative user interface for data scientists, business/data analysts, and SQL, R, and Python users who explore, prepare, and transform data, and who perform machine learning in Oracle Autonomous Database.	Oracle Autonomous Database

Product Name	Description	Available on:
OML AutoML User Interface	AutoML User Interface (AutoML UI) is an Oracle Machine Learning interface that provides you no-code automated machine learning. Business users without extensive data science background can use AutoML UI to create and deploy machine learning models, while data scientists can use it as a productivity tool for quickly assessing algorithms and hyperparameters.	Oracle Autonomous Database
OML for SQL	Oracle Machine Learning for SQL (OML4SQL) provides scalable in-database machine learning algorithms for model building and data scoring.	Oracle Database (on premises and Database Cloud Service) and Oracle Autonomous Database
OML for Python	Oracle Machine Learning for Python (OML4Py) is a Python API for performing statistical and machine learning analysis on data in your Oracle Database. It provides a Python interface to the in-database algorithms and supports deployment of user-defined Python functions with optimal system-supported data-parallel and task-parallel invocation. OML4Py also provides Python functions supporting automated machine learning (AutoML).	Oracle Database and Oracle Database Cloud Service 19c, 21c, 23ai, and Oracle Autonomous Database
Oracle Spatial AI	Oracle Spatial AI is a set of capabilities in OML4Py for detecting patterns and making predictions from geospatial data. It includes features for end-to-end workflows, including spatial machine learning pipelines.	Oracle Autonomous Database
OML for R	Oracle Machine Learning for R (OML4R) is an R API for performing statistical and machine learning analysis on data in your Oracle database. It provides an R interface to the in-database algorithms and supports deployment of user-defined R functions with optional system-supported data-parallel and task-parallel invocation.	Oracle Database and Oracle Database Cloud Service, and Oracle Autonomous Database

Product Name	Description	Available on:
OML Services	The REST API for Oracle Machine Learning Services provides a REST API supporting model management and deployment, data and model monitoring, and cognitive text capabilities. It supports lightweight scoring for real-time and streaming use cases for both in-database models and ONNX-format models.	Oracle Autonomous Database
Oracle Data Miner	An extension to Oracle SQL Developer that enables data scientists and citizen data scientists to explore and prepare data, easily build and compare multiple machine learning models, make predictions, and accelerate model deployment.	Oracle Database and Oracle Autonomous Database

Creating Analytic Views

You can create Analytic Views and view information about them. You can also edit and perform other actions on them.



When you create an Analytic View, you identify a fact table that contains the data to inspect. The **Generate Hierarchies and Measures** button looks at the contents of that table, identifies any hierarchies in the fact table, and searches for other tables that may contain related hierarchies.

While creating an Analytic View, you can enable or disable the following advanced options:

- Autonomous Aggregate Cache, which uses the dimensional metadata of the Analytic View to manage a cache and that improves query response time.

- Analytic View Transparency Views, which presents Analytic Views as regular database views and enables you to use your analytic tools of choice while gaining the benefits of Analytic Views.
- Analytic View Base Table Query Transformation, which enables you to use your existing reports and tools without requiring changes to them.

Create Analytic View

To create Analytic View, click **Create** from the Data Analysis home page and select **Create Analytic View** to begin the process.

Click **Cancel** to cancel the creation of the Analytic View at any time.

Specify Attributes of the Analytic View

On the **General** tab of the **Create Analytic View** pane, specify the following:

- The name for the Analytic View
- The fact table for the view
- Advanced options

You can also preview the data of the fact table and see statistics about that data.

In the **Name** field, specify a name of your choice.

The **Schema** field has the current user's schema. You can only create an Analytic View in that schema.

In the **Fact Table** field, expand the drop-down list and click **More Sources**. The **Select Sources** dialog box has a list of the available tables and views. Select a table or view from the list.

To filter the list, begin typing characters in the Filter field. As you type, the list changes to show the tables or views that contain the characters. Clear the field to show the complete list again. After you select a table or view, click OK.

To enable or disable the advanced options, on the **Create Analytic View** pane, click the **Show Advanced Options** icon at the bottom left. Select or deselect options as desired.

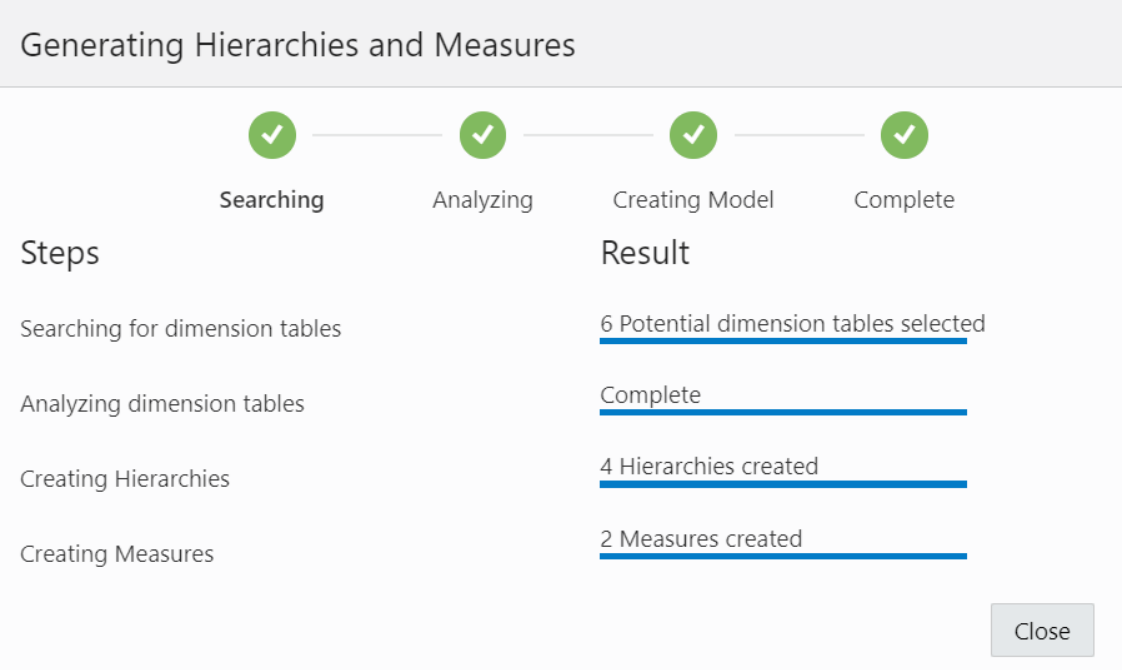
To view the data in the fact table and statistics about the data, click the **Preview Data** button. In the Preview and Statistics pane, the Preview tab shows the columns of the table and the data in the columns.

The Statistics tab shows the size of the table and the number of rows and columns. The statistics may take a few moments to appear, during which time the message, "No statistics available..." may appear. The statistics include the names of the columns, their data types, the number of distinct values and null values, the maximum and minimum values, and other information. The bar graph displays the top unique column values and the number of their occurrences for the selected column. Point to a bar in the graph to see the number of occurrences of the unique value.

Click **Close** to close the Preview and Statistics pane and return to the Create Analytic View pane.

Click on **Generate Hierarchies and Measures** icon.

The Generating Hierarchies and Measures dialog box displays the progress of searching for dimension tables, analyzing the dimension tables and identifying and creating the data sources, joins, hierarchies, and measures to use. When the process completes, click **Close**.



The **Search for Dimension Tables** check box when selected, enables you to search for dimension tables while generating hierarchies and measures.

After the hierarchies and measures are generated, they are displayed under their respective tabs. Review the hierarchies and measures created for you.

Specify the **Name**, **Fact Table** and select **Advanced Options** in the **General** tab of **Create Analytic View** pane. Click **Create** to generate an Analytic View.

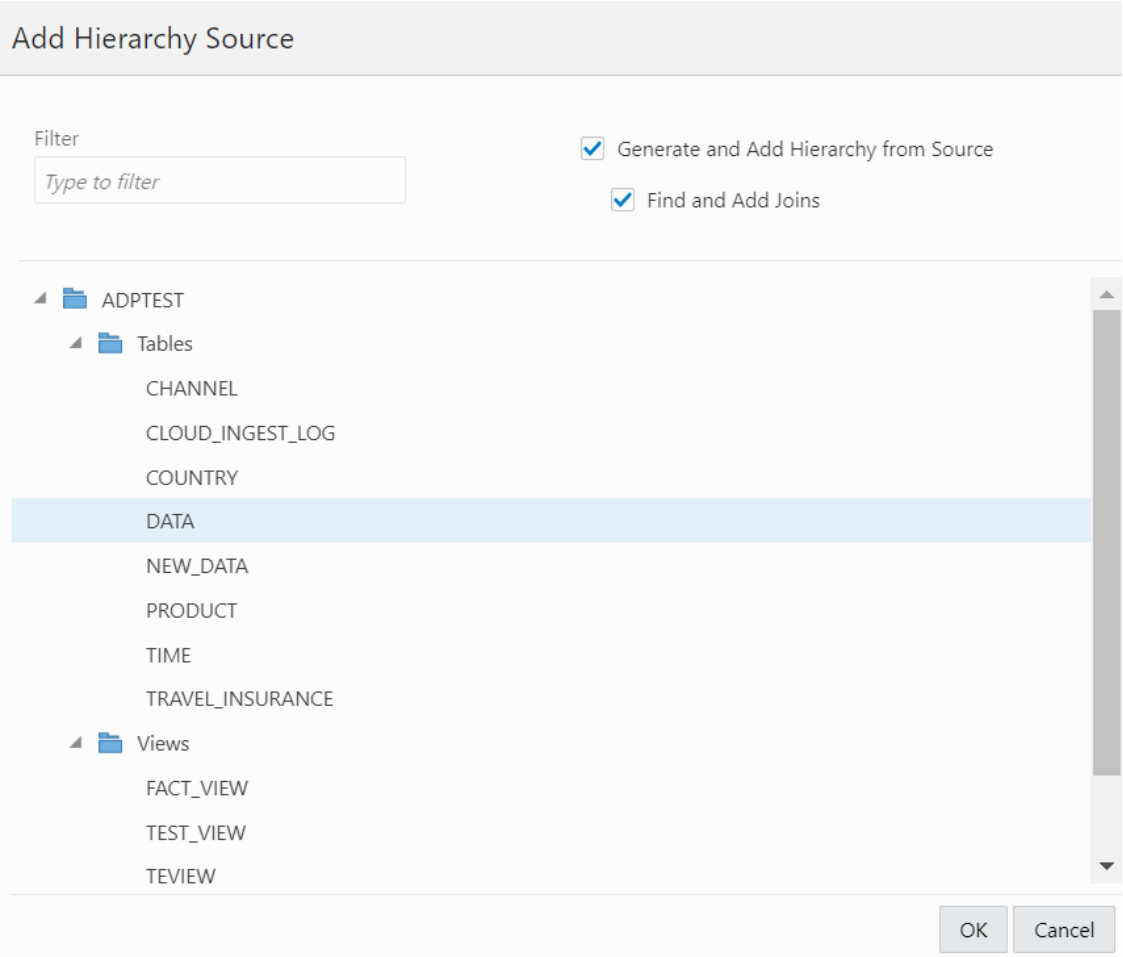
View Data Sources

The Data Sources tab displays the sources of the data and the relationships among them. It has a graphical display of the fact table and the related dimension tables. For example, a fact table of health insurance data might have columns for geography identifiers, income codes, and gender codes. The Data Sources tab would display items for the fact table and for the geography, income, and gender dimension tables.

You can add hierarchies from data sources even after generating hierarchies from the existing fact table. You can add one or more hierarchies to your new or existing analytic view. Multiple hierarchies can be defined and used in an analytical view, however only one will be used by default.

Right-click the Data Sources tab and select **Add Hierarchy Sources** or select **Add Hierarchy Sources**.

Selecting **Add Hierarchy Sources** launches an **Add Hierarchy Source** dialog box.



You can view all the fact tables and views associated with the analytic view.

In the filter field, you can either manually look for the source or start typing to search for the fact table or views from the list of available fact tables and views. After typing the full name of the source, the tool automatically matches the fact table or view.

Select **Generate and Add hierarchy from Source** to generate analysis and hierarchies associated with the source data you select.

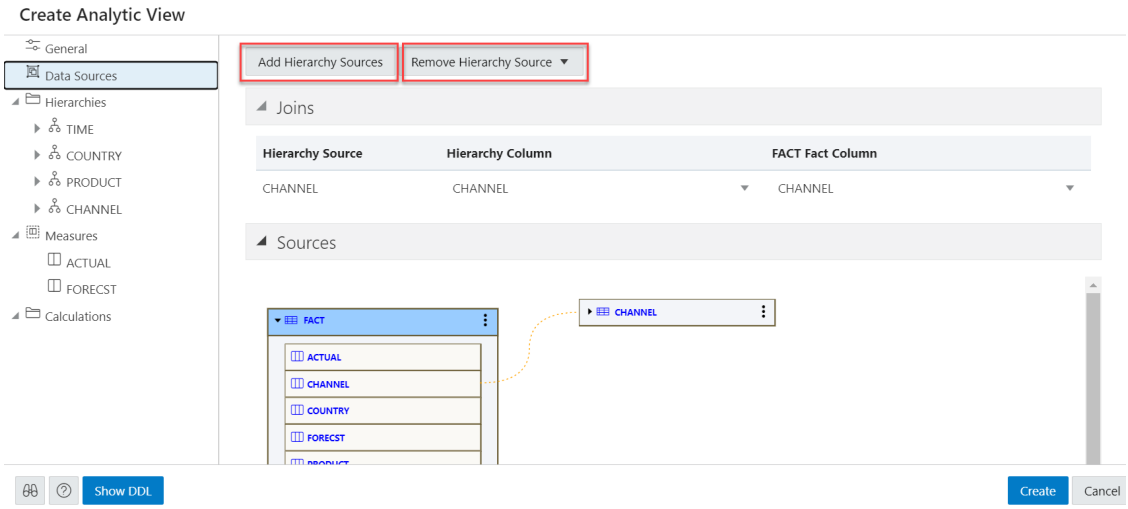
Select **Find and Add Joins** to link all the data sources with the fact table. You can add multiple join entries for a single hierarchy.

Click **OK** to select the source.

The Generating Hierarchies and Measures dialog box displays the progress of analyzing the dimension tables and creating the hierarchies. When the process completes, click **Close**.

> **Note:**
>
> When you add a hierarchy from the data source, you see the new hierarchy in the list of hierarchies in the Hierarchies tab. You can navigate between the Data Sources tab, the Hierarchies tab, the Measures tab, the Calculations tab. You can add a hierarchy from a source that is not connected by navigating back to the Data Sources tab.

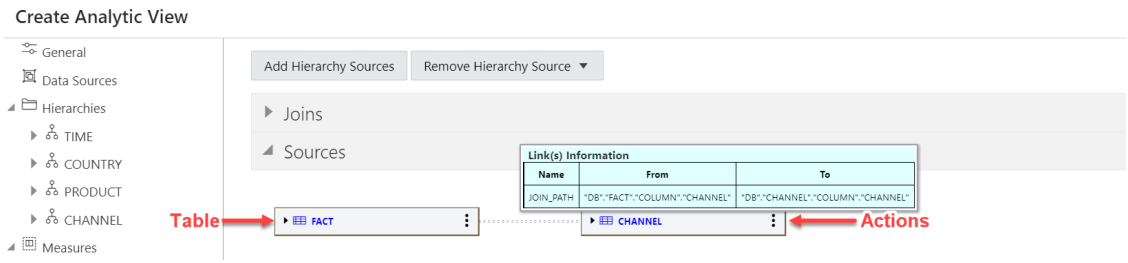
Select **Remove Hierarchy Source** to remove the hierarchies you create from the data sources. You cannot remove hierarchies generated from the fact table you select from this option.



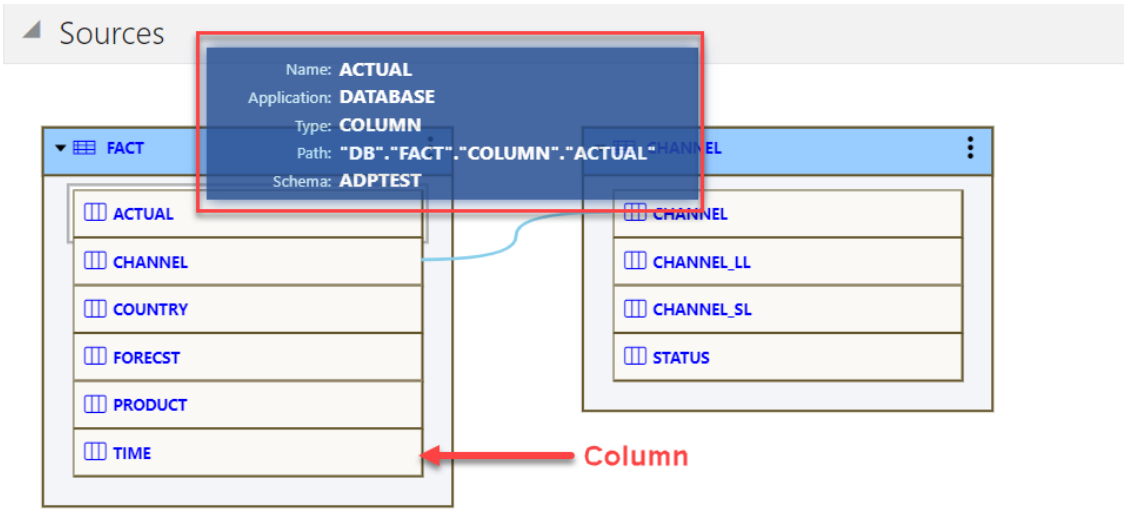
Expand **Joins** to view the **Hierarchy Source**, **Hierarchy Column** and the **Fact column** mapped with the Analytic View. The **Joins** is visible only when the hierarchy table differs from the fact table. You can add multiple join entries for a single hierarchy.

Expand **Sources** to view the fact table associated with the Analytic View. The data model expands to include the data from the source that you added.

Pointing to an item displays the name, application, type, path and the schema of the table. Click the **Actions** (three vertical dots) icon at the right of the item to display a menu to expand or collapse the view of the table.



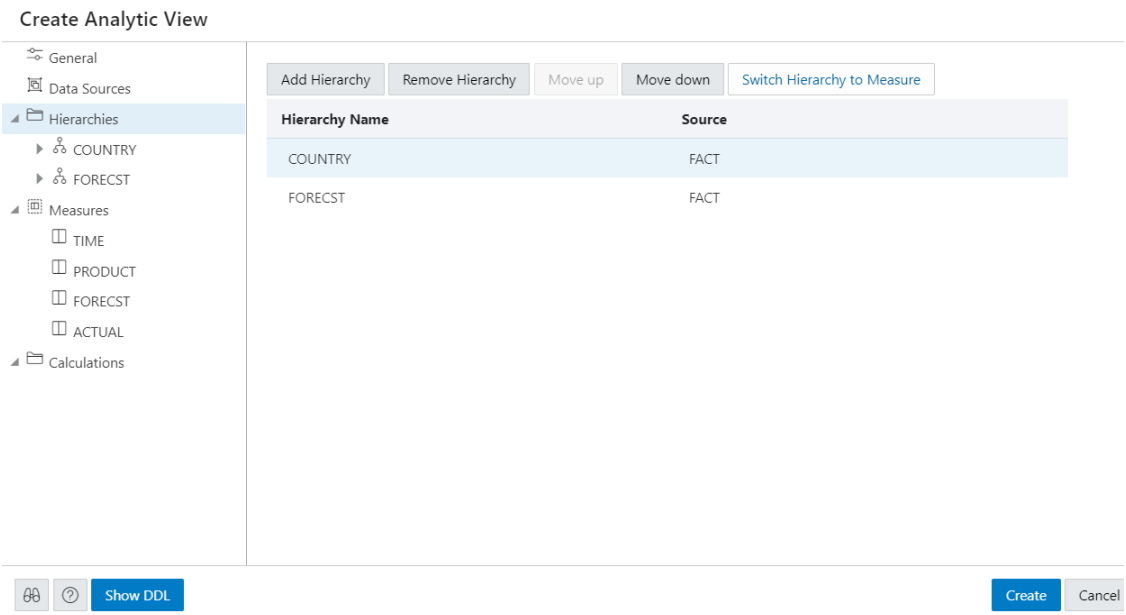
An expanded item displays the columns of the table. Pointing to a column displays the name, application, type, path, and schema of the column.



The lines that connect the dimension tables to the fact table indicate the join paths between them. Pointing to a line displays information about the join paths of the links between the tables. If the line connects a table that is collapsed, then the line is dotted. If the line connects two expanded tables, then the line is solid and connects the column in the dimension table to the column in the fact table.

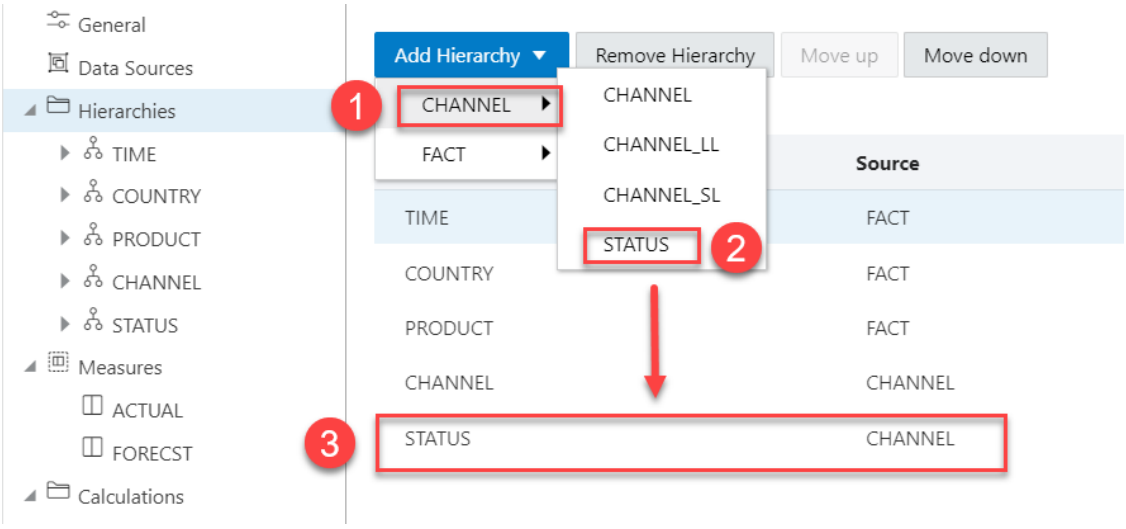
View and Manage Hierarchies

The Hierarchies tab displays the hierarchies generated by the Analytic View creation tool. The display includes the name of the hierarchy and the source table.



An analytic view must include at least one hierarchy.

To add a Hierarchy, click **Add Hierarchy**. This results in a display as a list of column in that table. Select a column that operates as the detailed level of the hierarchy and be the join-key to the fact table.



To remove the hierarchy, select the hierarchy you want to remove from the list and click **Remove Hierarchy**

Select **Move Up** or **Move Down** to position the order of the Hierarchy in the resulting view.

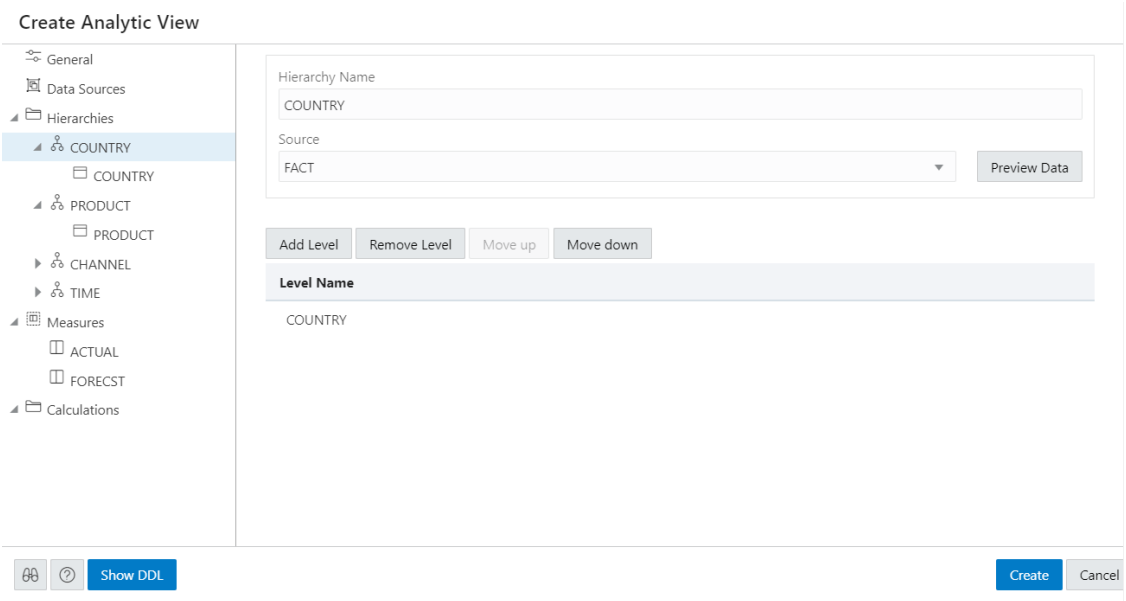
Click **Switch Hierarchy to Measure** to change the hierarchy you select to a measure in the Measures list.

You can also **Add Hierarchy** and **Add Hierarchy From Table** by right-clicking the Hierarchy tab.

If you click on a hierarchy name, a dialog box displays the Hierarchy Name and Source.

To change the source, select a different source from the drop-down list.

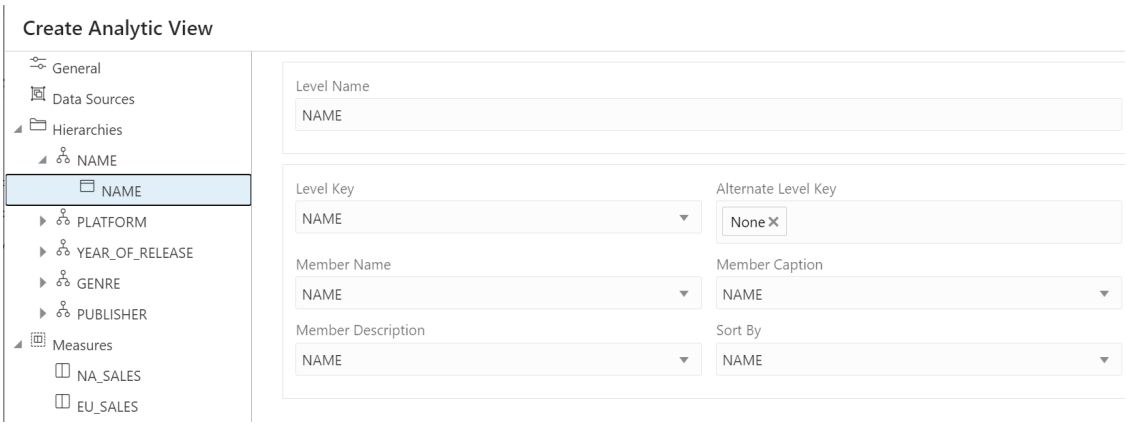
Select **Add Level** to add a level to the hierarchy. Click **Remove Level** to remove the selected level from the hierarchy.



To view the data in the fact table and statistics about the data, click the **Preview Data** button. In the Preview and Statistics pane, the Preview tab shows the columns of the table and the data in the columns. The Statistics tab shows the size of the table and the number of rows and columns.

If you click on a particular level in the Hierarchy tab, a dialog box displays it's respective Level Name, Level Key, Alternate Level Key, Member Name, Member Caption, Member Description, source, and Sort By drop-down. To change any of the field values, enter the value in the appropriate field.

> **Note:**
>
> You can enter multiple level keys Member Name, Member Caption, Member Description and Sort By.



Member Captions and Member Descriptions generally represent detailed labels for objects. These are typically end-user-friendly names. For example, you can caption a hierarchy representing geography areas named GEOGRAPHY_HIERARCHY as "Geography" and specify its description as "Geographic areas such as cities, states, and countries."

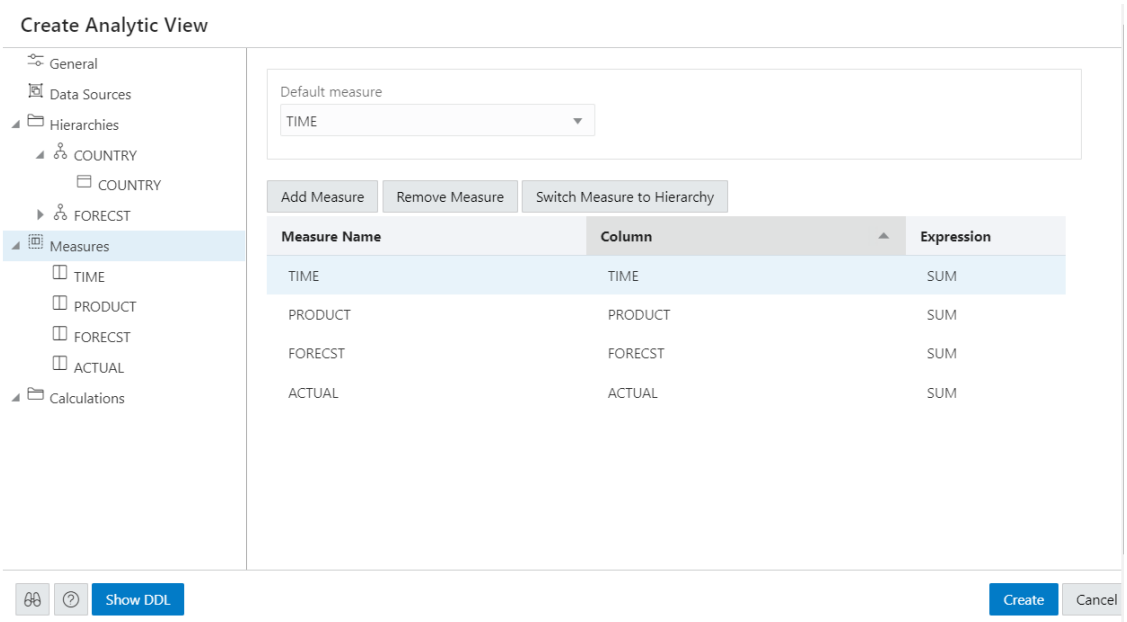
To see the measures for the Analytic View, click **Measures** tab. To immediately create the Analytic View, click **Create**. To cancel the creation, click **Cancel**.

View and Manage Measures

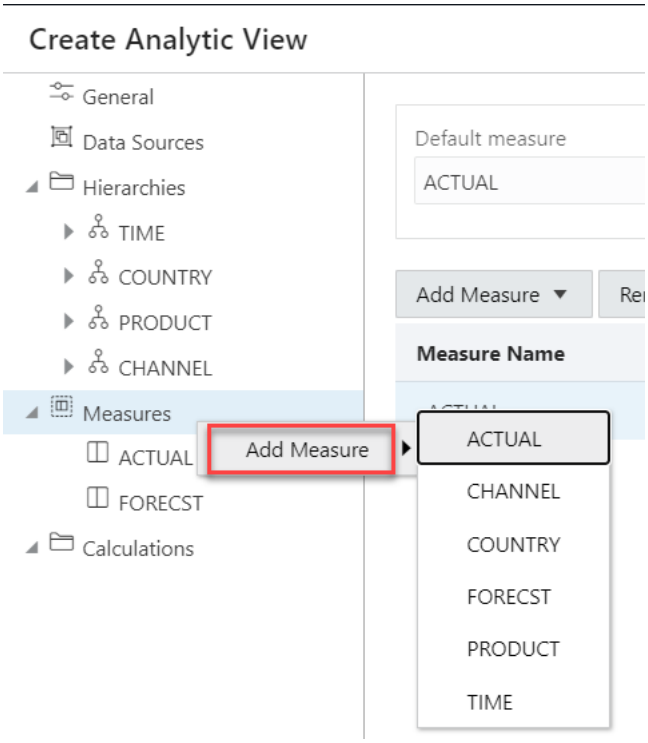
The Measures tab displays the measures suggested for the Analytic View. It displays the Measure Name, Column, and operator Expression for each measure.

The measures specify fact data and the calculations or other operations to perform on the data.

To add measures, click **Add Measure**. You can view a new measure at the bottom of the measures list. To remove the measure, select the measure you want to remove from the list and click **Remove Measure**.



To alternatively add a measure from the data source, right- click the Measures tab. This pops up a list of columns that can be used as measures. Select one measure from the list.



You can exclude a column from the measures on right-clicking the Measures tab and selecting Remove Measure.

Click **Switch Measure to Hierarchy** to change the measure you select to hierarchy in the Hierarchies list.

You must specify a measure as the default measure for the analytic view; otherwise, the first measure in the definition is the default. Select **Default Measure** from the drop-down.

To add a measure, right-click the Measures tab and select **Add Measure**. To remove a measure, select the particular measure you want to remove, right-click on it and select **Remove Measure**.

You can select a different column for a measure from the Column drop-down list. You can select a different operator from the Expression drop-down list.

In creating an analytic view, you must specify one or more hierarchies and a fact table that has at least one measure column and a column to join to each of the dimension tables outside of the fact table.

Create new calculated measures

You can add measure calculations to a query of an analytic view.

The measures and hierarchies associated with the analytic views enable us to create new calculated measures.

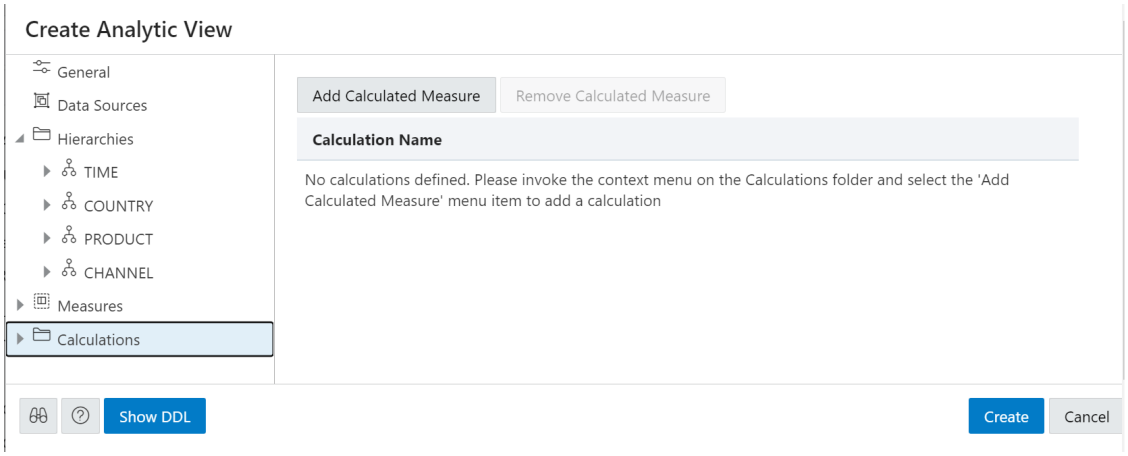
Calculated measures return values from data stored in one or more measures. You compute these measures at run time.

> **Note:**
>
> You can create the measures without increasing the size of the database since the calculated measures do not store the data. However, they may slow performance. You need to decide which measures to calculate on demand.

The Analytic Views provides easy-to-use templates for creating calculated measures.

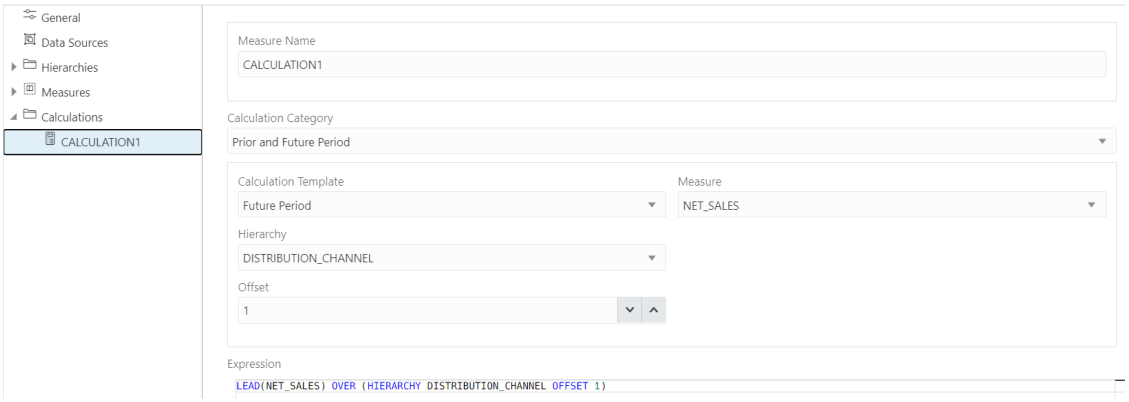
Once you create a calculated measure, it appears in the list of measures of the Analytic View .You can create a calculated measure at any time which is available for querying in SQL.

The Data Analysis tool provides easy-to-use templates for creating calculated measures.



Click **Add Calculated Measure** to add calculations to the measures. You can view the new calculation with system generated name in the **Calculations** tab.

Click the newly created calculated measure.



In the **Measure Name** field, enter the name of the calculated measure.

You can select preferred category of calculation from a list of options such as Prior and Future Period, Cumulative Aggregates, Period To Date, Parallel Period, Moving Aggregates, Share, Qualified Data Reference, and Ranking using the **Calculation Category** drop-down.

Your choice of category of calculation dynamically changes the **Calculation Template**.

For more details on how to use Calculation templates, see Using Calculation Templates.

Select the **Measure** and **Hierarchy** on which you want to base the calculated measures.

Select **Offset** value by clicking the up or the down arrow. The number specifies the number of members to move either forward or backward from the current member. The ordering of members within a level is dependent on the definition of the attribute dimension used by the hierarchy. The default value is 0 which represents POSITION FROM BEGINNING.

The Expression field lists the expressions which the calculated measure uses.

On the creation of the Analytic view, the calculated measure appears in the navigation tree in the Calculated Measures folder.

Click **Create**. A confirmation dialog box appears that asks for your confirmation. Select **Yes** to proceed with the creation of Analytic View.

After creating the Analytic View, you will view a success message informing you of its creation.

On editing the Analytic View you create, you can view the calculated measure in the navigation tree in the Calculations folder.

Click the **Tour** icon for a guided tour of the worksheet highlighting salient features and providing information if you are new to the interface.

Click the **help** icon to open the contextual or online help for the page you are viewing.

Click **Show DDL** to generate Data Definition Language statements for the analytic view.

Edit Analytic View

You might want to edit an Analytic View to make changes to the data sources, the hierarchies, or the measures.

To edit an Analytic View, click the **Action** icon on the Analytic View item, then click **Edit Analytic View**. On the Edit Analytic View screen, select a tab and make changes as desired.

When you have completed the changes, click **Update**.

Use Select AI for Natural Language Interaction with your Database

Select AI enables SQL access to generative AI using Large Language Models (LLMs) and embedding models. This includes support for natural language to SQL query generation and retrieval augmented generation, among other features.

- About Select AI
  Use natural language to interact with your database and LLMs through SQL to enhance user productivity and develop AI-based applications. Select AI simplifies and automates using generative AI, whether generating, running, and explaining SQL from a natural language prompt, using retrieval augmented generation with vector stores, generating synthetic data, or chatting with the LLM.
- Select AI Concepts
  Explores the concepts and terms related to Select AI.
- Select AI Use Cases
  Select AI enhances data interaction and enables developers to build AI-driven applications directly from SQL, transforming natural language prompts to SQL queries and text responses, supporting chat interaction with LLMs, enhancing response accuracy with current data using RAG, and generating synthetic data.
- Getting Started with Select AI
  To get started, review the prerequisites and the tasks that you need to perform to use Select AI.
- Manage AI Profiles
  You can create and manage your AI profiles through `DBMS_CLOUD_AI` package.
- Use AI Keyword to Enter Prompts
  Use `AI` as the keyword in a `SELECT` statement for interacting with the database using natural language prompts.
- Enable Conversations to Enhance User Interaction
  Conversations in Select AI refer to the interactive dialogue between the user and the system, where a sequence of user-provided natural language prompts are used to query or interact with the database.
- Select AI with Retrieval Augmented Generation (RAG)
  Select AI with RAG augments your natural language prompt by retrieving content from your specified vector store using semantic similarity search. This reduces hallucinations by using your specific and up-to-date content and provides more relevant natural language responses to your prompts.
- Synthetic Data Generation
  Generate synthetic data using random generators, algorithms, statistical models, and Large Language Models (LLMs) to simulate real data for developing and testing solutions effectively.
- Examples of Using Select AI
  Explore integrating Oracle's Select AI with various supported AI providers to generate, run, and explain SQL from natural language prompts or chat with the LLM.

Related Topics

- DBMS_CLOUD_AI Package
  The `DBMS_CLOUD_AI` package, with Select AI, facilitates and configures the translation of natural language prompts to generate, run, explain SQL statements. Also, enables retrieval

augmented generation and natural language-based interactions, including chatting with LLMs.

About Select AI

Use natural language to interact with your database and LLMs through SQL to enhance user productivity and develop AI-based applications. Select AI simplifies and automates using generative AI, whether generating, running, and explaining SQL from a natural language prompt, using retrieval augmented generation with vector stores, generating synthetic data, or chatting with the LLM.

When you use Select AI, Autonomous Database manages the process of converting natural language into SQL. This means you can provide a natural language prompt instead of SQL code to interact with your data. Select AI serves as a productivity tool for SQL users and developers and enables non-expert SQL users to derive useful insights from their data, without having to understand data structures or technical languages.

Select AI also automates the retrieval augmented generation (RAG) process from generating vector embeddings to retrieving relevant content based on your prompt through semantic similarity search using your vector store. Other features include synthetic data generation, supporting chat history for conversations, and other features, all from a SQL interface.

The `DBMS_CLOUD_AI` package enables integration with a user-specified LLM for generating SQL code using natural language prompts. For natural language to SQL generation, this package provides an augmented prompt to the LLM containing the relevant database schema metadata. This enables generating, running, and explaining SQL queries based on natural language prompts. It also facilitates retrieval augmented generation using vector stores, synthetic data generation, and allows for chatting with the LLM. The `DBMS_CLOUD_AI` package works with AI providers listed in Select your AI Provider and LLMs.

> **Note:**
>
> - You must have an account with the AI provider and provide the credentials through `DBMS_CLOUD_AI` objects that the Autonomous Database uses.
> - You can submit prompts in multiple languages. The quality of the result depends on the abilities of the specific LLM or the embedding model (transformer) being used. Check your LLM or embedding model documentation for multi-language support.

Topics

- Usage Guidelines
  Provides usage guidelines to assist in the use of Select AI for natural language to SQL generation..
- Supported Platforms
  Select AI is supported on Autonomous Database Serverless and Autonomous Database on Dedicated Exadata Infrastructure and Cloud at Customers.
- Select your AI Provider and LLMs
  Choose an AI provider and LLM that meets your security standards and aligns with your specific needs, like text or code generation.

Usage Guidelines

Provides usage guidelines to assist in the use of Select AI for natural language to SQL generation..

Intended Use

This feature generates, runs, and explains SQL queries from user-provided natural language prompts. It automates tasks that users would otherwise perform manually using their schema metadata and a large language model (LLM) of their choice. Additionally, it facilitates retrieval augmented generation with vector stores and enables chatting with the LLM.

Depending on the Select AI action you specify, you provide a prompt, whether for natural language to SQL generation, RAG, or pass-through chat, and Select AI automates the interaction with LLMs and your database using SQL and PL/SQL interfaces. Specifically, it generates SQL queries from natural language based on metadata from the specified schema and tables. Additionally, it facilitates chat-based generative AI, optionally enhanced with content from vector stores through retrieval augmented generation (RAG) for improved response quality. It also explains SQL queries based on natural language prompts and supports synthetic data generation for one or multiple schema tables. Select AI enables submitting general requests with the `chat` action.

Prompt Augmentation Data

For SQL query generation, the database augments the user-specified prompt with database metadata to mitigate hallucinations from the LLM. The augmented prompt is then sent to the user-specified LLM to produce the query. When using vector stores with retrieval augmented generation (RAG), content from the vector store is retrieved using semantic similarity search with the provided prompt. This content becomes part of the augmented prompt that is sent to the LLM.

The database augments the prompt with schema metadata only. This metadata may include schema definitions, table and column comments, and content available from the data dictionary. For the purposes of SQL generation, the database does not provide table or view contents (actual row or column values) when augmenting the prompt.

The `narrate` action, however, does provide to the LLM either:

- the result of a natural language to SQL query, which contains database data, or
- the result of semantic similarity search as retrieved from the vector store supporting retrieval augmented generation (RAG).

The LLM uses these results to generate a natural language text response.

> **WARNING:**
>
> Large language models (LLMs) have been trained on a broad set of text documentation and content, typically from the Internet. As a result, LLMs may have incorporated patterns from invalid or malicious content, including SQL injection. Thus, while LLMs are adept at generating useful and relevant content, they also can generate incorrect and false information including SQL queries that produce inaccurate results and/or compromise security of your data.
>
> The queries generated on your behalf by the user-specified LLM provider will be run in your database. Your use of this feature is solely at your own risk, and, notwithstanding any other terms and conditions related to the services provided by Oracle, constitutes your acceptance of that risk and express exclusion of Oracle's responsibility or liability for any damages resulting from that use.

Supported Platforms

Select AI is supported on Autonomous Database Serverless and Autonomous Database on Dedicated Exadata Infrastructure and Cloud at Customers.

- Autonomous Database Serverless
- Autonomous Database on Dedicated Exadata Infrastructure
- Autonomous Database on Dedicated Exadata Infrastructure Region
- Autonomous Database Cloud@Customer

Select your AI Provider and LLMs

Choose an AI provider and LLM that meets your security standards and aligns with your specific needs, like text or code generation.

Different LLMs excel at various tasks based on their training data and intended purpose. Some models are excellent for text generation but may not perform well in code generation, while others are specifically optimized for coding tasks. Choose an LLM that best suits your needs.

AI Provider	LLMs	Embedding Model for RAG	Purpose
OCI Generative AI	• meta.llama-3.3-70b-instruct (default) • meta.llama-3.2-90b-vision-instruct • meta.llama-3.2-11b-vision-instruct • meta.llama-3.1-70b-instruct • meta.llama-3.1-405b-instruct • cohere.command-r-08-2024 • cohere.command-r-plus-08-2024 • cohere.command-r-16k (deprecated) • cohere.command–r-plus (deprecated) See About the Chat Models in Generative AI.	• cohere.embed-english-v3.0 (default) • cohere.embed-multilingual-v3.0 • cohere.embed-english-light-v3.0 • cohere.embed-multilingual-light-v3.0 See About the Embedding Models in Generative AI.	The OCI Generative AI Chat models are supported for all `SELECT AI` actions such as `runsql`, `showsql`, `explainsql`, `narrate`, and `chat`. The OCI Generate text models are supported only for `SELECT AI chat` action.

Note: For models that accept images, use `meta`

AI Provider	LLMs	Embedding Model for RAG	Purpose

```
.llama-3.2-90b-vision-instruct. This model is specifical
```

AI Provider	LLMs	Embedding Model for RAG	Purpose
			y t r a i n e d f o r v i s i o n a n d i m a g e s . W h i l e i t c a n b e u s e d f o r t e x t a n d S Q

AI Provider	LLMs	Embedding Model for RAG	Purpose
			Lgeneration,themodelisbeststuitedforimages.Tolearnmore,s

AI Provider	LLMs	Embedding Model for RAG	Purpose

```
ee ChatinOCIGenerativeAI. The meta.llama-3.2-11b-vision-in
```

AI Provider	LLMs	Embedding Model for RAG	Purpose
			struct model provides robust multimodal capabilities. To configure your profile attributes, see Profile Attributes.

AI Provider	LLMs	Embedding Model for RAG	Purpose
Azure OpenAI Service	• GPT-4o • GPT-4 • GPT-4 Turbo with Vision • GPT-3.5-Turbo	text-embedding-ada-002	Best suited for generating SQL from natural language prompts, `chat` action, and Select AI RAG.
OpenAI	• gpt-3.5-turbo (default) • gpt-4o • gpt-4o-mini • gpt-4 • gpt-4-0613 • gpt-4-32k • gpt-4-32k-0613 • gpt-3.5-turbo-0613 • gpt-3.5-turbo-16k • gpt-3.5-turbo-16k-0613	text-embedding-ada-002	Best suited for generating SQL from natural language prompts, `chat` action, and Select AI RAG.
Cohere	• command (default) • command-nightly (experimental) • command-r • command-r-plus • command-light • command-light-nightly (experimental) • custom models	embed-english-v2.0	Best suited for `chat` action.
Google	• gemini-1.5-flash (default) • gemini-1.5-pro • gemini-1.0-pro	text-embedding-004 (default)	Best suited for generating SQL from natural language prompts, `chat` action, and Select AI RAG.
Anthropic	• claude-3-5-sonnet-20240620 (default) • claude-3-opus-20240229 • claude-3-sonnet-20240229 • claude-3-haiku-20240307	NA	Best suited for generating SQL from natural language prompts, `chat` action, and Select AI RAG.
Hugging Face	• Mixtral-8x7B-Instruct-v0.1 (default) • Meta-Llama-3-70B-Instruct • Qwen1.5-1.8B • other chat models	NA	Best suited for generating SQL from natural language prompts, `chat` action, and Select AI RAG.

> **Note:**
>
> Embedding models are also known as transformer models.

Select AI Concepts

Explores the concepts and terms related to Select AI.

- Actions
- AI Profile
- AI Provider
- Conversations
- Database Credentials
- Hallucination in LLM
- IAM
- Large Language Model (LLM)
- Metadata
- Metadata Clone
- Natural Language Prompts
- Network Access Control List (ACL)
- Retrieval Augmented Generation (RAG)
- Semantic Similarity Search
- Vector Distance
- Vector Index
- Vector Store

Actions

An action in Select AI is a keyword that instructs Select AI to perform different behavior when acting on the prompt. By specifying an action, users can instruct Select AI to process their natural language prompt to generate SQL code, to respond to a chat prompt, narrate the output, display the SQL statement, or explain the SQL code, leveraging the LLMs to efficiently interact with the data within their database environment.

The following are the supported actions for Select AI:

- `runsql`: Generates the SQL statement for a natural language prompt and runs the underlying SQL query to return a rowset. This is the default action and does not require specifying this parameter.
- `showsql`: Displays the SQL statement for a natural language prompt.
- `narrate`: Sends the result of the SQL query run by the database back to the LLM to generate a natural language description of that result.

  When a vector index is specified in the AI profile to enable RAG, the system uses the specified transformer model to create a vector embedding from the prompt for semantic

similarity search against the vector store.. The system then adds the retrieved content from the vector store to the user prompt and sends it to the LLM to generate a response based on this information.

- `chat`: Passes the user prompt directly to the LLM to generate a response, which is provided to the user.
- `explainsql`: Explains the generated SQL from the prompt into natural language. This option sends the generated SQL to the AI provider, which then produces a natural language explanation.

For more details on how you can use these actions, see Examples of Using Select AI.

AI Profile

An AI profile is a specification that includes the AI provider to use and other details regarding metadata and database objects required for generating responses to natural language prompts. See CREATE_PROFILE Procedure and Profile Attributes.

AI Provider

An AI Provider in Select AI refers to the service provider that supplies the LLM or transformer or both for processing and generating responses to natural language prompts. These providers offer models that can interpret and convert natural language for the use cases highlighted under the LLM concept. See Select your AI Provider and LLMs for the supported providers.

Conversations

Conversations in Select AI represent an interactive exchange between the user and the system, enabling users to query or interact with the database through a series of natural language prompts. Select AI incorporates up to 10 previous prompts into the current request, creating an augmented prompt sent to the LLM. See Enable Conversations to Enhance User Interaction.

Database Credentials

Database credentials are authentication credentials used to access and interact with databases. They typically consist of a user name and a password, sometimes supplemented by additional authentication factors like security tokens. These credentials are used to establish a secure connection between an application or user and a database, such that only authorized individuals or systems can access and manipulate the data stored within the database.

Hallucination in LLM

Hallucination in the context of Large Language Models refers to a phenomenon where the model generates text that is incorrect, nonsensical, or unrelated to the input prompt. Despite being a result of the model's attempt to generate coherent text, these responses can contain information that is fabricated, misleading, or purely imaginative. Hallucination can occur due to biases in training data, lack of proper context understanding, or limitations in the model's training process.

IAM

Oracle Cloud Infrastructure Identity and Access Management (IAM) lets you control who has access to your cloud resources. You can control what type of access a group of users have

and to which specific resources. To learn more, see Overview of Identity and Access Management.

Large Language Model (LLM)

A Large Language Model (LLM) refers to an advanced type of artificial intelligence model that is trained on massive amounts of text data to support a range of use cases depending on their training data. This includes understanding and generating human-like language as well as software code and database queries. These models are capable of performing a wide range of natural language processing tasks, including text generation, translation, summarization, question answering, sentiment analysis, and more. LLMs are typically based on sophisticated deep learning neural network models that learn patterns, context, and semantics from the input data, enabling them to generate coherent and contextually relevant text.

Metadata

Database metadata refers to the data that describes the structure, organization, and properties of the database tables and views.

Metadata Clone

A metadata clone or an Autonomous Database clone creates a copy of a metadata defining the database or schema, containing only the structure, not the actual data. This clone includes tables, indexes, views, statistics, procedures, and triggers without any data rows. Developers, testers, or those building database templates find this useful. To learn more, see Clone, Move, or Upgrade an Autonomous Database Instance.

Natural Language Prompts

Natural Language Prompts are human-readable instructions or requests provided to guide generative AI models, such as Large Language Models. Instead of using specific programming languages or commands, users can interact with these models by entering prompts in a more conversational or natural language form. The models then generate output based on the provided prompt.

Network Access Control List (ACL)

A Network Access Control List is a set of rules or permissions that define what network traffic is allowed to pass through a network device, such as a router, firewall, or gateway. ACLs are used to control and filter incoming and outgoing traffic based on various criteria such as IP addresses, port numbers, and protocols. They play a crucial role in network security by enabling administrators to manage and restrict network traffic to prevent unauthorized access, potential attacks, and data breaches.

Retrieval Augmented Generation (RAG)

Retrieval Augmented Generation (RAG) is a technique that involves retrieving relevant information for a user's query and supplying that information to a large language model (LLM) to improve responses and reduce hallucination.
Most commonly, RAG involves vector search, but more generally, includes augmenting a prompt of database content (either manually or automatically) such as schema metadata for SQL generation or database content explicitly queried. Other forms of augmentation can involve technologies such as graph analytics and traditional machine learning.

Semantic Similarity Search

Semantic similarity search identifies and retrieves data points that closely match a given query by comparing feature vectors in a vector store.

Vector Distance

Vector distance measure the similarity or dissimilarity between feature vectors by calculating the distance between them in a multidimensional space.

Vector Index

A vector index organizes and stores vectors to enable efficient similarity search and retrieval of related data.

Vector Store

A vector store includes systems that store, manage, and enable semantic similarity search involving vector embeddings. This includes standalone vector databases and Oracle Database 23ai AI Vector Search.

Select AI Use Cases

Select AI enhances data interaction and enables developers to build AI-driven applications directly from SQL, transforming natural language prompts to SQL queries and text responses, supporting chat interaction with LLMs, enhancing response accuracy with current data using RAG, and generating synthetic data.

Use cases include:

- Generate SQL from natural language prompts

  Developer productivity: Select AI significantly enhances developer productivity by providing "starter" SQL queries quickly. Developers can input natural language prompts, and Select AI generates SQL based on your database schema tables and views. This reduces the time and effort needed to write complex queries from scratch, allowing developers to focus on refining and optimizing the generated queries for their specific needs.

  Natural language queries for end-users: Select AI empowers end-users to interact with your application's underlying data tables and views using natural language queries. This functionality allows users without SQL expertise to ask questions and retrieve data directly, making data access more intuitive and user-friendly relative to the capabilities of LLM being used and the quality of the schema metadata available.

  Other capabilities for SQL generation: The following highlighted capabilities are also supported for natural language to SQL generation:

  – Specify schema or tables or views: Select AI enables you to specify an object list consisting of schema and optionally tables or views within that schema.
  – Automatically detect relevant table metadata: Select AI automatically detects relevant tables and send metadata only for those specific tables, relevant to the query, in Oracle Database 23ai.
  – Restrict table access: Select AI enables you to restrict table access by considering only the tables listed in the AI profile attributes for SQL generation.

  - Specify case sensitivity for columns: Select AI enables user to specify case sensitivity such that the LLM produces case-insensitive responses from the database and LLM.

- Conversations

  Enable chatbot-like features with Select AI, enabling users to have natural conversations for querying data and performing actions. These chats can keep track of context, giving follow-up answers that clarify or expand on original questions. This scenario boosts engagement and makes complicated queries easier through conversation.

- Customized media generation

  Select AI can be used to generate personalized media content such as emails tailored to individual customer details. For example, in your prompt you could instruct the LLM to create a friendly and upbeat email encouraging a customer to try a set of recommended products. These recommendations could be based on customer demographics or other specific information available in your database. This level of customization enhances customer engagement by delivering relevant and appealing content directly to the customer.

- Code generation

  With the Select AI `chat` action, you can use Select AI to ask your specified LLM to generate code from natural language prompts. This feature supports various programming languages such as SQL, Python, R, and Java. Examples include:

  - **Python Code**: "Write the Python code to compute a confusion matrix over a DataFrame with columns ACTUAL and PREDICTED."
  - **SQL DDL**: "Write the DDL for a SQL table with columns name, age, income, and country."
  - **SQL Query**: "Write the SQL query that will use the Oracle Machine Learning in-database model named CHURN_DT_MODEL to predict which customers will churn and with what probability."

- Retrieval Augmented Generation (RAG)

  Use vector store content for semantic similarity search to enhance prompt accuracy and relevance in LLM responses.

- Synthetic data generation

  Generate synthetic data using LLMs that conforms to your schema for solution testing, proofs of concept, and other uses. Synthetic data can support better testing of your applications in the absence of real data, leading to overall quality of your application.

  Synthetic data generation can also be used to populate an Autonomous Database clone or a metadata clone. Select AI supports generating synthetic data for such clones. Using synthetic data helps to protect sensitive data while enabling development, testing, and validating user experiences. It's also useful for AI and machine learning projects needing sample data for model training or test data for scoring.

Getting Started with Select AI

To get started, review the prerequisites and the tasks that you need to perform to use Select AI.

- Configure your system to use Select AI
- Create an AI profile and enable the AI profile
- Use AI keyword in SELECT statement and enter natural language prompts

- Example of using Select AI with OCI Generative AI
- Examples of Using Select AI
- Manage AI profile, configure vector index, or drop your AI profile by reviewing DBMS_CLOUD_AI package
- Customize your AI profiles by supplying the profile attributes
- Display AI profiles, AI profile attributes, and vector index details by querying DBMS_CLOUD_AI views
- See Autonomous Database Supplied Package Reference

Manage AI Profiles

You can create and manage your AI profiles through `DBMS_CLOUD_AI` package.

- Use DBMS_CLOUD_AI to Configure AI Profiles
  Autonomous Database uses AI profiles to facilitate and configure access to an LLM and to setup for generating, running, and explaining SQL based on natural language prompts. It also facilitates retrieval augmented generation using embedding models and vector indexes and allows for chatting with the LLM.
- Perform Prerequisites for Select AI
  Before you use Select AI, here are the steps to enable `DBMS_CLOUD_AI`.
- Create and Set an AI Profile
  Describes the steps to create and enable an AI profile.

> **See Also:**
>
> Use Select AI for Natural Language Interaction with your Database

Use DBMS_CLOUD_AI to Configure AI Profiles

Autonomous Database uses AI profiles to facilitate and configure access to an LLM and to setup for generating, running, and explaining SQL based on natural language prompts. It also facilitates retrieval augmented generation using embedding models and vector indexes and allows for chatting with the LLM.

AI profiles include database objects that are the target for natural language queries. Metadata used from these targets can include database table names, column names, column data types, and comments. You create and configure AI profiles using the `DBMS_CLOUD_AI.CREATE_PROFILE` and `DBMS_CLOUD_AI.SET_PROFILE` procedures.

In addition to specifying tables and views in the AI profile, you can also specify tables mapped with external tables, including those described in Query External Data with Data Catalog. This enables you to query data not just inside the database, but also data stored in a data lake's object store.

Perform Prerequisites for Select AI

Before you use Select AI, here are the steps to enable `DBMS_CLOUD_AI`.

The following are required to use `DBMS_CLOUD_AI`:

- Access to an Oracle Cloud Infrastructure cloud account and to an Autonomous Database instance.
- A paid API account of a supported AI provider, one of:

AI Provider	API Keys
OpenAI	See Use OpenAI to get your API keys.
Cohere	See Use Cohere to get your secret API keys.
Azure OpenAI Service	See Use Azure OpenAI Service for more information on how to configure Azure OpenAI Service.
OCI Generative AI	See Use OCI Generative AI.
Google	See Use Google to get your API keys.
Anthropic	See Use Anthropic to get your API keys.
Hugging Face	See Use Hugging Face to get your API keys.

- Network ACL privileges to access your external AI provider.

> **Note:**
>
> Network ACL privileges are not required for OCI Generative AI.

- A credential that provides access to the AI provider.

Enable `DBMS_CLOUD_AI`

To configure `DBMS_CLOUD_AI`:

1. Grant the `EXECUTE` privilege on the `DBMS_CLOUD_AI` package to the user who wants to use Select AI.

   By default, only the system administrator has `EXECUTE` privilege. The administrator can grant `EXECUTE` privilege to other users.

2. Grant `EXECUTE` privilege on `DBMS_CLOUD_PIPELINE` to the user who wants to use Select AI with RAG.

> **Note:**
>
> If the user already has the `DWROLE` role, this privilege is included and additional grant is not required.

3. Grant network ACL access to the user who wants to use Select AI and for the AI provider endpoint.

   The system administrator can grant network ACL access. See APPEND_HOST_ACE Procedure for more information.

4. Create a credential to enable access to your AI provider.
   See CREATE_CREDENTIAL Procedure for more information.

5. Grant quotas in tablespace to manage the amount of space in a specific tablespace to the user who wants to use Select AI with RAG.

Examples of Privileges to Run Select AI

The following example grants the `EXECUTE` privilege to `ADB_USER`:

```
GRANT execute on DBMS_CLOUD_AI to ADB_USER;
```

The following example grants `EXECUTE` privilege for the `DBMS_CLOUD_PIPELINE` package required for RAG:

```
GRANT EXECUTE on DBMS_CLOUD_PIPELINE to ADB_USER;
```

To check the privileges granted to a user for the `DBMS_CLOUD_AI` and `DBMS_CLOUD_PIPELINE` packages, an administrator can run the following:

```
SELECT table_name AS package_name, privilege
 FROM DBA_TAB_PRIVS
 WHERE grantee = '<username>'
 AND   (table_name = 'DBMS_CLOUD_PIPELINE'
        OR table_name = 'DBMS_CLOUD_AI');
```

The following example grants `ADB_USER` the privilege to use the *api.openai.com* endpoint.

> **Note:**
>
> This procedure is not applicable to OCI Generative AI.

```
BEGIN
    DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
         host => 'api.openai.com',
         ace  => xs$ace_type(privilege_list => xs$name_list('http'),
                             principal_name => 'ADB_USER',
                             principal_type => xs_acl.ptype_db)
   );
END;
/
```

The parameters are:

- `host`: The host, which can be the name or the IP address of the host. You can use a wildcard to specify a domain or an IP subnet. The host or domain name is not case sensitive.

AI Provider	Host
OpenAI	*api.openai.com*
Cohere	*api.cohere.ai*
Azure OpenAI Service	*<azure_resource_name>.openai.azure.com* See Profile Attributes to know more about `azure_resource_name`.

AI Provider	Host
Google	*generativelanguage.googleapis.com*
Anthropic	*api.anthropic.com*
Hugging Face	*api-inference.huggingface.co*

- `ace`: The access control entries (ACE). The `XS$ACE_TYPE` type is provided to construct each ACE entry for the ACL. For more details, see Creating ACLs and ACEs.

The following example creates a credential to enable access to OpenAI.

```
EXEC
DBMS_CLOUD.CREATE_CREDENTIAL(
credential_name   => 'OPENAI_CRED',
username          =>  'OPENAI',
password          =>  '<your_api_token>');
```

The parameters are:

- `credential_name`: The name of the credential to be stored. The `credential_name` parameter must conform to Oracle object naming conventions.
- `username`: The `username` and `password` arguments together specify your AI provider credentials.

  The `username` is a user-specified user name.
- `password`: The `username` and `password` arguments together specify your AI provider credentials.

  The `password` is your AI provider secret API key, and depends on the provider:

AI Provider	API Keys
OpenAI	See Use OpenAI to get your API keys.
Cohere	See Use Cohere to get your API keys.
Azure OpenAI Service	See Use Azure OpenAI Service to get your API keys and to configure the service. **Note:** If you are using the Azure OpenAI Service principal to authenticate, you can skip the `DBMS_CLOUD.CREATE_CREDENTIAL` procedure. See Examples of Using Select AI for an example of authenticating using Azure OpenAI Service principal.
OCI Generative AI	See Use OCI Generative AI to generate API signing keys.

AI Provider	API Keys
Google	See Use Google to generate your API keys.
Anthropic	See Use Anthropic to generate your API keys.
Hugging Face	See Use Hugging Face to generate your API keys.

The following example grants quotas on tablespace to the `ADB_USER` to use Select AI with RAG:

```
ALTER USER ADB_USER QUOTA 1T ON <tablespace_name>;
```

To the check the tablespace quota granted to a user, run the following:

```
SELECT TABLESPACE_NAME, BYTES, MAX_BYTES
FROM DBA_TS_QUOTAS
WHERE USERNAME = '<username>' AND
      TABLESPACE_NAME LIKE 'DATA%';
```

The parameters are:

- `TABLESPACE_NAME`: The tablespace for which the quota is assigned. In Autonomous Database, tablespaces are managed automatically and have `DATA` as a prefix.
- `BYTES`: The amount of space currently used by the user in the tablespace.
- `MAX_BYTES`: The maximum quota assigned (in bytes). If `MAX_BYTES` is -1, it means the user has unlimited quota on the tablespace. The database user creating the vector index must have `MAX_BYTES` sufficiently larger than bytes to accommodate the vector index, or `MAX_BYTES` must be -1 for unlimited quota.

Topics

- Use OpenAI
  To enable OpenAI to generate SQL and text responses to your natural language prompts, obtain API keys from your OpenAI paid account.
- Use Cohere
  To enable Cohere to generate SQL and text responses to your natural language prompts, obtain API keys from your Cohere paid account.
- Use Azure OpenAI Service
  To enable Azure OpenAI Service to generate SQL and text responses to your natural language prompts, configure and provide access to the AI provider.
- Use OCI Generative AI
  To enable OCI Generative AI to generate SQL and text responses to your natural language prompts, generate an API signing key.
- Use Google
  To enable Google AI Studio to generate SQL and text responses to your natural language prompts, obtain API keys from your Google AI Studio paid account.
- Use Anthropic
  To enable Anthropic Developer Console to generate SQL and text responses to your natural language prompts, obtain API keys from your Anthropic Developer Console paid account.

- Use Hugging Face
  To enable Hugging Face as your AI provider to generate SQL and text responses to your natural language prompts, obtain API keys from your Hugging Face paid account.

Use OpenAI

To enable OpenAI to generate SQL and text responses to your natural language prompts, obtain API keys from your OpenAI paid account.

You can find your secret API key from your profile dashboard under API keys.

Use Cohere

To enable Cohere to generate SQL and text responses to your natural language prompts, obtain API keys from your Cohere paid account.

1. Login to Cohere's website with your credentials. Cohere Dashboard appears by default.
2. Alternately,click **Dashboard**.
3. Click **API Keys** on the left navigation. Copy the default API key or create another key. See API-Keys for more information.

Use Azure OpenAI Service

To enable Azure OpenAI Service to generate SQL and text responses to your natural language prompts, configure and provide access to the AI provider.

To use Azure OpenAI Service, perform the following steps:

1. Obtain your secret API keys. You can find your API keys in the Resource Management section of your Azure portal. On your Azure OpenAI Service Resource page, click **Keys and Endpoint**. You can copy either KEY1 or KEY2.
2. Create an Azure OpenAI Service resource and deploy a model: Create and deploy an Azure OpenAI Service resource.

   > **Tip:**
   >
   > - Make note of the resource name and deployment name as those parameters are used to provide network access permission and to create your Azure OpenAI Service profile using the `DBMS_CLOUD_AI.CREATE_PROFILE` procedure.
   > - To know about rate limits for token per minute on a model, see Azure OpenAI Service quotas and limits.

3. Allow access to Azure OpenAI Service:
   - You can use your secret API key to allow access to Azure OpenAI Service. To know more, see the example in Examples of Using Select AI.
   - Allow service principal to access Azure OpenAI Service and grant the required permissions to the service principal: Use Azure Service Principal to Access Azure Resources and Role-based access control for Azure OpenAI Service. To know more, see the example in Examples of Using Select AI.

Use OCI Generative AI

To enable OCI Generative AI to generate SQL and text responses to your natural language prompts, generate an API signing key.

Use the Console or command line to generate an API signing key for any Oracle Database instance. See How to Generate the API Signing Key.

Use Google

To enable Google AI Studio to generate SQL and text responses to your natural language prompts, obtain API keys from your Google AI Studio paid account.

1. Go to Google AI Studio.
2. Click Sign In to Google AI Studio.
3. Click Get API key on the prompt screen.
4. Select all the applicable options on the next page.
5. Click Create API key.
6. Click Create API key in new project.

   The screen displays the progress and generates an API key. Copy the key and save it.

Use Anthropic

To enable Anthropic Developer Console to generate SQL and text responses to your natural language prompts, obtain API keys from your Anthropic Developer Console paid account.

1. Go to Anthropic Developer Console.
2. Sign up for an account if you don't have one already.
3. Once logged in, navigate to the API section or the dashboard.
4. Look for an option to generate or view API keys.
5. Click to create a new API key.
6. Copy the generated API key and save it.

   The Claude API is a paid service. You'll need to add credits to your account before you can use the API key.

Use Hugging Face

To enable Hugging Face as your AI provider to generate SQL and text responses to your natural language prompts, obtain API keys from your Hugging Face paid account.

1. Go to Hugging Face.
2. Sign up for an account if you don't have one already.
3. Navigate to your account settings.
4. In the navigation menu locate the **Access Tokens**.
5. Click to create a new API key.
6. Copy the generated API key and save it.

Create and Set an AI Profile

Describes the steps to create and enable an AI profile.

Use `DBMS_CLOUD_AI.CREATE_PROFILE` to create an AI profile. Run `DBMS_CLOUD_AI.SET_PROFILE` to enable the AI profile so that you can use `SELECT AI` with a natural language prompt.

> **Note:**
>
> You must run `DBMS_CLOUD_AI.SET_PROFILE` in each new stateful database session (connection) before you use `SELECT AI`. If you are using a stateless connection, you must use the `DBMS_CLOUD_AI.GENERATE` function which enables you to specify the profile name in each invocation.

The following example with the OpenAI provider creates an `AI` profile called `OPENAI` and sets the `OPENAI` profile for the current user session.

```
-- Create AI profile
--BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
  profile_name   => 'OPENAI',
  attributes     =>'{"provider": "openai",
            "credential_name": "OPENAI_CRED",
            "object_list": [{"owner": "SH", "name": "customers"},
                      {"owner": "SH", "name": "sales"},
                      {"owner": "SH", "name": "products"},
                      {"owner": "SH", "name": "countries"}]
       }');
END;
/

PL/SQL procedure successfully completed.

--
-- Enable AI profile in current session
--
EXEC DBMS_CLOUD_AI.set_profile('OPENAI');

PL/SQL procedure successfully completed.
```

Use AI Keyword to Enter Prompts

Use `AI` as the keyword in a `SELECT` statement for interacting with the database using natural language prompts.

The `AI` keyword in a `SELECT` statement instructs the SQL execution engine to use the LLM identified in the active AI profile to process natural language and to generate SQL.

You can use the `AI` keyword in a query with Oracle clients such as SQL Developer, OML Notebooks, and third-party tools, to interact with database in natural language.

> **Note:**
>
> You cannot run PL/SQL statements, DDL statements, or DML statements using the `AI` keyword.

Syntax

The syntax for running AI prompt is:

```
SELECT AI action natural_language_prompt
```

Parameters

The following are the parameters available for the *action* parameter:

Parameter	Description
`runsql`	Runs the underlying SQL command for the natural language prompt. This is the default action and it is optional to specify this parameter.
`showsql`	Displays the SQL statement for a natural language prompt.
`explainsql`	Explains the generated SQL from the prompt in a natural language. This option sends the generated SQL to the AI provider to produce a natural language explanation.
`narrate`	Sends the result of the SQL query run by the database back to the LLM to generate a natural language description of that result. When a vector index is specified in the AI profile to enable RAG, the system uses the specified transformer (or default transformer) model to create a vector embedding from the prompt for semantic similarity search against the vector store. The system then adds the retrieved content from the vector store to the user prompt and sends it to the LLM to generate a response based on this information. If you do not want table data or vector search documents to be sent to an LLM, a user with administrator privileges can disable such access for all users of the given database. This, in effect, disables the `narrate` action.
`chat`	Passes the user prompt directly to the LLM to generate a response, which is provided to the user. If `conversation` in the `DBMS_CLOUD_AI.CREATE_PROFILE` function is set to `true`, this option includes content from prior interactions or prompts, potentially including schema metadata.

Usage Notes

- The `AI` keyword for Select AI is not supported in Database Actions or APEX Service. You can use only the `DBMS_CLOUD_AI.GENERATE` function.

  Do not use `DBMS_CLOUD_AI.SET_PROFILE` in Database Actions or APEX Service. Instead, set the AI profile in the `profile_name` argument of `DBMS_CLOUD_AI.GENERATE`.
- The `AI` keyword is supported only in a `SELECT` statement.
- You cannot run PL/SQL statements, DDL statements, or DML statements using the `AI` keyword.
- The sequence is `SELECT` followed by `AI`. These keywords are not case-sensitive. After setting your AI profile using `DBMS_CLOUD_AI.SET_PROFILE`, and optional action, the text

after `SELECT AI` is a natural language prompt. If an AI profile is not set, `SELECT AI` reports the following error when running a `SELECT AI` statement:

```
ORA-00923: FROM keyword not found where expected
00923. 00000 -  "FROM keyword not found where expected"
```

- Special character usage rules apply according to Oracle guidelines. For example, use single quotes twice if you are using an apostrophe in a sentence.

```
select ai how many customers in SF don''t own their own home
```

- LLMs are subject to *hallucinations* and results are not always correct:
  - It is possible that `SELECT AI` may not be able to run the generated SQL query for a specific natural language prompt.
  - It is possible that `SELECT AI` may not be able to generate a SQL query for a specific natural language prompt.
  - It is possible that the LLM may not generate a SQL query that produces a correct result given your natural language prompt.

  In such a scenario, `SELECT AI` may respond with information to assist you in generating valid SQL.
- For better results with natural language to SQL generation, use database views or tables with contextual column names or consider adding column comments explaining values stored in the columns.
- Use the `explainsql` action, as in `SELECT AI explainsql`, to get a more detailed explanation of the SQL statement, as compared to `SELECT AI showsql`.
- To access DBA or USER views, see DBMS_CLOUD_AI Views.

Enable Conversations to Enhance User Interaction

Conversations in Select AI refer to the interactive dialogue between the user and the system, where a sequence of user-provided natural language prompts are used to query or interact with the database.

Select AI includes past prompts (up to 10) in the current call as part of the augmented prompt sent to the LLM. Conversations enable a more intuitive and accessible way to work with your data, making it easier for users to extract insights and perform tasks without needing deep technical knowledge of SQL.

For example, conversations provide an intuitive way to analyze data:

- Ask a question such as "What are the total number of customers"
- Follow up with context aware questions like:
  - "break out count of customers by country"
  - "What age group is most common"
  - "keep the top 5 customers and their country by their purchases and include a rank in the result"

To enable conversations, set the `conversation` attribute to `true` in your AI profile, see Profile Attributes. See Example: Enable Conversations in Select AI for a complete example.

Select AI with Retrieval Augmented Generation (RAG)

Select AI with RAG augments your natural language prompt by retrieving content from your specified vector store using semantic similarity search. This reduces hallucinations by using your specific and up-to-date content and provides more relevant natural language responses to your prompts.

Select AI automates the Retrieval Augmented Generation (RAG) process. This technique retrieves data from enterprise sources using AI vector search and augments user prompts for your specified large language model (LLM). By leveraging information from enterprise data stores, RAG reduces hallucinations and generates grounded responses.

RAG uses AI vector search on a vector index to find semantically similar data for the specified question. Vector store processes vector embeddings, which are mathematical representations of various data points like text, images, and audio. These embeddings capture the meaning of the data, enabling efficient processing and analysis. For more details on vector embeddings and AI vector search, see Overview of AI Vector Search.

Select AI integrates with AI vector search available in Oracle Autonomous Database 23ai for similarity search using vector embeddings.

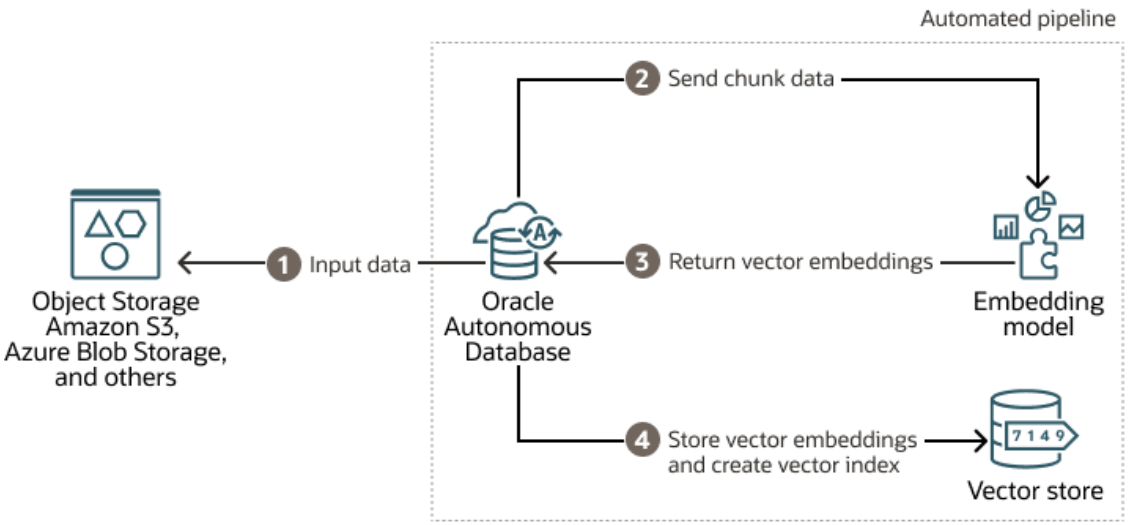
Topics

- Build your Vector Store
- Use DBMS_CLOUD_AI to Create and Manage Vector Indexes
  Use the `DBMS_CLOUD_AI` package to create and manage vector indexes and configure vector database JSON parameters.
- Use In-database Transformer Models
  Select AI RAG enables you to use pretrained ONNX transformer models that are imported into your database in Oracle Database 23ai instance for generating embedding vectors from document chunks and user prompts.
- Benefits of Select AI RAG
  Simplify querying, enhance response accuracy with current data, and gain transparency by reviewing sources used by the LLM.

Build your Vector Store

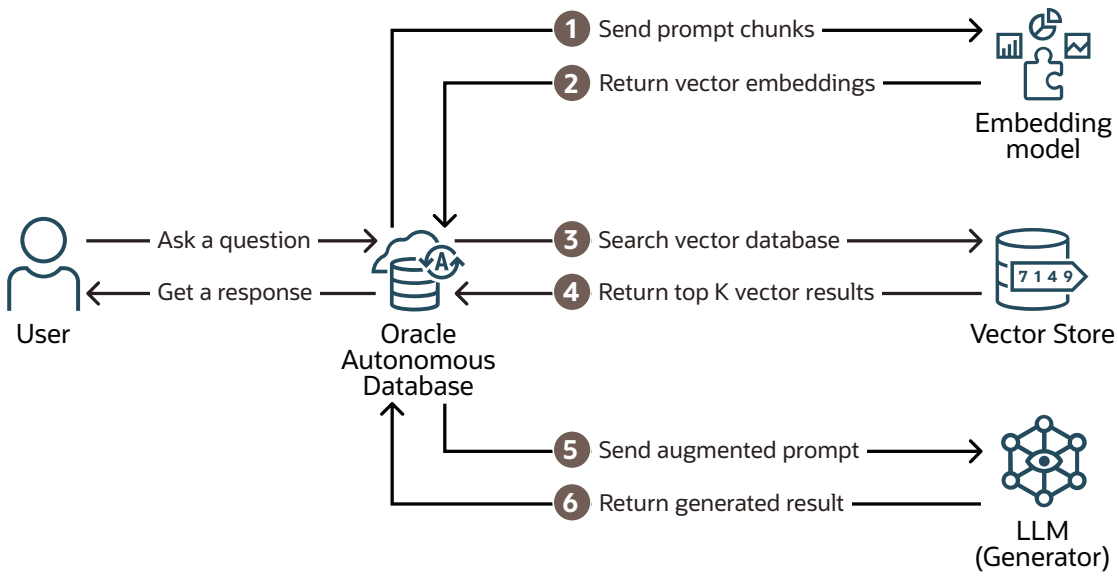
Select AI automates the creation and population of vector store by converting input documents (for example, PDF, DOC, JSON, XML, or HTML) from your object store to plain text. Oracle Text supports around 150 file types. For a complete list of all the supported document formats, see Supported Document Formats.
Select AI automatically processes documents to chunks, generates embeddings, stores them in the specified vector store, and updates the vector index as new data arrives.

Here is how the input from Object Storage is used with Select AI RAG:



1. Input: Data is initially stored in an Object Storage.
2. Oracle Autonomous Database retrieves the input data or the document, chunks it and sends the chunks to an embedding model.
3. The embedding model processes the chunk data and returns vector embeddings.
4. The vector embeddings are then stored in a vector store for use with RAG. As content is added, the vector index is automatically updated.

RAG retrieves relevant pieces of information from the enterprise database to answer a user's question. This information is provided to the specified large language model along with the user prompt. Select AI uses this additional enterprise information to enhance the prompt, improving the LLM's response. RAG can enhance response quality with update-to-date enterprise information from the vector store.



Select AI implements RAG as follows:

1. Input: User asks a question (specifies a prompt) using Select AI `narrate` action.

2. Select AI generates vector embeddings of the prompt using the embedding model specified in the AI profile.
3. The vector search index uses the vector embedding of the question to find matching content from the customer's enterprise data (searching the vector store) which has been indexed.
4. The vector search returns top K texts similar to the input to your Autonomous Database instance.
5. Autonomous Database then sends these top K query results with user question to the LLM.
6. The LLM returns its response to your Autonomous Database instance.
7. Autonomous Database Select AI provides the response to the user.

Use DBMS_CLOUD_AI to Create and Manage Vector Indexes

Use the `DBMS_CLOUD_AI` package to create and manage vector indexes and configure vector database JSON parameters.

After you create credentials and provide network access to the vector database and the AI provider, your Autonomous Database instance uses AI profiles to configure access to LLMs. See Example: Set Up and Use Select AI with RAG for a complete example on setting it up and using it in Select AI statements.

> **Note:**
>
> If you do not want table data or vector search documents to be sent to an LLM, a user with administrator privileges can disable such access for all users of the given database. This, in effect, disables the `narrate` action for RAG.

You can configure AI profiles for providers listed in Select your AI Provider and LLMs through the `DBMS_CLOUD_AI` package.

> **See Also:**
>
> - Create a vector index: CREATE_VECTOR_INDEX Procedure.
> - Manage vector index profiles and other AI profiles: Summary of DBMS_CLOUD_AI Subprograms
> - Query vector index views: DBMS_CLOUD_AI Views.

Use In-database Transformer Models

Select AI RAG enables you to use pretrained ONNX transformer models that are imported into your database in Oracle Database 23ai instance for generating embedding vectors from document chunks and user prompts.

To learn more about importing a pretrained ONNX model into your Oracle Database 23ai instance, see

- Import Pretrained Models in ONNX Format for Vector Generation Within the Database
- Alternate Method to Import ONNX Models
- Pre-built Embedding Generation model for Oracle Database 23ai

> **Note:**
>
> You must import a pretrained ONNX-format transformer model into Oracle Database 23ai instance to use Select AI RAG with imported in-database transformer model. You can also use other transformer models from supported AI providers.

See Example: Select AI with In-database Transformer Models to explore the feature.

Benefits of Select AI RAG

Simplify querying, enhance response accuracy with current data, and gain transparency by reviewing sources used by the LLM.

Select AI RAG offers the following benefits:

- Simplify data querying and increase response accuracy: Enable users to query enterprise data using natural language and provide LLMs with detailed context from enterprise data to generate more accurate and relevant responses, reducing instances of LLM hallucinations.
- Up-to-date information: Provide LLMs access to current enterprise information using vector stores, eliminating the need for costly, time-consuming fine-tuning of LLMs trained on static data sets.
- Seamless integration: Integrate with Oracle AI Vector Search for streamlined data handling and enhanced performance.
- Automated data orchestration: Automate orchestration steps with a fully managed Vector Index pipeline, ensuring efficient processing of new data.
- Understandable contextual results: Has access and retrieves the sources used by the LLM from vector stores, ensuring transparency and confidence in results. Views and extracts data in natural language text or JSON format for easier integration and application development.

Synthetic Data Generation

Generate synthetic data using random generators, algorithms, statistical models, and Large Language Models (LLMs) to simulate real data for developing and testing solutions effectively.

Synthetic data can be a powerful tool when developing and testing solutions, especially when actual data doesn't yet exist or isn't allowed to be used. Synthetic, or artificially generated, data can have many of the characteristics of real data. Synthetic data is typically created using random generators, algorithms, or statistical models to simulate the characteristics and distributions of real data. However, this can be complex to produce or rely on tools with features of varying sophistication. With the availability of Large Language Models (LLMs), more relevant and schema-specific data may be generated that considers characteristics expressed in natural language.

Topics

- Benefits of Synthetic Data Generation
  Synthetic data generation enables populating database metadata clones, supporting development, testing, and machine learning projects without using sensitive data from original tables.
- Generate Synthetic Data
  Use `DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA` function to generate synthetic data and query the data using Select AI actions.
- Monitor and Troubleshoot Synthetic Data Generation
  When generating large amounts of data across many tables, Select AI splits synthetic data generation tasks into smaller chunks and runs tasks in parallel. The status of each chunk is tracked in the `SYNTHETIC_DATA$<operation_id>_STATUS` table.

Benefits of Synthetic Data Generation

Synthetic data generation enables populating database metadata clones, supporting development, testing, and machine learning projects without using sensitive data from original tables.

Synthetic Data Generation offers the following benefits:

- **Populating metadata clones with synthetic data**: A metadata clone replicates the structure of a database or schema without including actual data. Select AI allows synthetic data generation to populate these clones, protecting sensitive data while enabling development, testing, and creating templates. This approach supports performance and scalability testing.
- **Starting new projects**: When starting a new project, actual data may not be available. Synthetic data provides realistic samples to help demonstrate concepts and gain support for project proposals.
- **Validating user experience**: Synthetic data aids in testing user interfaces by providing diverse data sets to uncover design flaws, performance, and scalability issues.
- **Supporting AI and machine learning projects**: Synthetic data is useful for training AI and machine learning models when real data is unavailable or restricted. LLMs can generate data with specific patterns to facilitate model training and scoring.

Generate Synthetic Data

Use `DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA` function to generate synthetic data and query the data using Select AI actions.

To learn more, see GENERATE_SYNTHETIC_DATA Function and Example: Generate Synthetic Data.

Monitor and Troubleshoot Synthetic Data Generation

When generating large amounts of data across many tables, Select AI splits synthetic data generation tasks into smaller chunks and runs tasks in parallel. The status of each chunk is tracked in the `SYNTHETIC_DATA$<operation_id>_STATUS` table.

Synthetic data generation operations are logged in the tables `DBA_LOAD_OPERATIONS` and `USER_LOAD_OPERATIONS`. Use these tables to monitor the

DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA operation. See Track DBMS_CLOUD Load Operations for more details.

After running synthetic data generation in a given session, you can get the latest <*operation_id*> from USER_LOAD_OPERATION using the following:

```
SELECT max(id) FROM user_load_operations;
```

To view the synthetic data operations running in a different session, use the DBA_LOAD_OPERATIONS view.

View Status of Synthetic Data Generation

The status table of synthetic data generation operation shows the progress of each table and its corresponding chunk. The STATUS_TABLE column in USER_LOAD_OPERATIONS or DBA_LOAD_OPERATIONS shows the status table name. The table name is SYNTHETIC_DATA$<*operation_id*>_STATUS and it has following columns:

Name	Datatype	Description
ID	NUMBER	Unique identifier of the record.
NAME	VARCHAR2	Qualified name of the table, such as "ADB_USER"."EMPLOYEES"
BYTES	NUMBER	Desired number of records for this data generation task
ROWS_LOADED	NUMBER	Actual number of records generated.
CHECKSUM	VARCHAR2	Starting value for the primary key during this data generation task.
LAST_MODIFIED	TIMESTAMP WITH TIME ZONE	Timestamp indicating when the record was last modified.
STATUS	VARCHAR2	Status of the data generation task. The valid values are: • COMPLETED • FAILED • PENDING • SKIPPED
ERROR_CODE	NUMBER	Error code, if the data generation task fails.
ERROR_MESSAGE	VARCHAR2	Error message provided if the task fails.
END_TIME	TIMESTAMP WITH TIME ZONE	Timestamp marking the end of the data generation task.

Example: Check the Number of Records Generated for Each Table
To check the number of records generated for each table, issue the following:

```
SELECT name, SUM(rows_loaded) FROM synthetic_data$<operation_id>_status group
by name;
```

Query `ROWS_LOADED` to confirm how many number of rows are loaded for each chunk, and `SUM(ROWS_LOADED)` for rows for each table.

```
BEGIN
    DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA(
        profile_name => 'GENAI',
        object_list => '[{"owner": "ADB_USER", "name":
"DIRECTOR","record_count":150},
                        {"owner": "ADB_USER", "name":
"MOVIE_ACTOR","record_count":300},
                        {"owner": "ADB_USER", "name": "CLASSES",
"user_prompt":"all in fall semester","record_count":5},
                        {"owner": "ADB_USER", "name":
"ACTOR","record_count":220},
                        {"owner": "ADB_USER", "name":
"MOVIE","record_count":50}]'
    );
END;
/


-- Check loaded rows for each chunk
SQL> SELECT name, rows_loaded FROM synthetic_data$141_status order by name;
NAME                           ROWS_LOADED
----------------------------------------
"ADB_USER"."ACTOR"                      188
"ADB_USER"."ACTOR"                       32
"ADB_USER"."CLASSES"                      5
"ADB_USER"."DIRECTOR"                   150
"ADB_USER"."MOVIE"                       50
"ADB_USER"."MOVIE_ACTOR"                 38
"ADB_USER"."MOVIE_ACTOR"                114
"ADB_USER"."MOVIE_ACTOR"                148



-- Check loaded rows for each table
SQL> SELECT name, SUM(rows_loaded) FROM synthetic_data$141_status group by
name;

NAME                       SUM(ROWS_LOADED)
----------------------------------------
"ADB_USER"."DIRECTOR"                   150
"ADB_USER"."MOVIE_ACTOR"                300
"ADB_USER"."CLASSES"                      5
"ADB_USER"."ACTOR"                      220
"ADB_USER"."MOVIE"                       50
```

Examples of Using Select AI

Explore integrating Oracle's Select AI with various supported AI providers to generate, run, and explain SQL from natural language prompts or chat with the LLM.

- Example: Select AI Actions
  These examples illustrate common Select AI actions.

- Example: Select AI with OCI Generative AI
  These examples show how you can access OCI Generative AI using your OCI API key or Resource Principal, create an AI profile, and generate, run, and explain SQL from natural language prompts or chat using the OCI Generative AI LLMs.
- Example: Select AI with OpenAI
  This example shows how you can use OpenAI to generate SQL statements from natural language prompts.
- Example: Select AI with Cohere
  This example shows how you can use Cohere to generate SQL statements from natural language prompts.
- Example: Select AI with Azure OpenAI Service
  The following examples shows how you can enable access to Azure OpenAI Service using your API key or use Azure OpenAI Service Principal, create an AI profile, and generate SQL from natural language prompts.
- Example: Select AI with Google
- Example: Select AI with Anthropic
- Example: Select AI with Hugging Face
- Example: Enable Conversations in Select AI
  This example illustrates enabling conversations in Select AI.
- Example: Set Up and Use Select AI with RAG
  This example guides you through setting up credentials, configuring network access, and creating a vector index for integrating OCI Generative AI vector store cloud services with OpenAI using Oracle Autonomous Database.
- Example: Select AI with In-database Transformer Models
  This example demonstrates how you can import a pretrained transformer model that is stored in Oracle object storage into your Oracle Database 23ai instance and then use the imported in-database model in Select AI profile to generate vector embeddings for document chunks and user prompts.
- Example: Improve SQL Query Generation
  These examples demonstrate how comments, annotations, foreign key, and referential integrity constraints in database tables and columns can improve the generation of SQL queries from natural language prompts.
- Example: Generate Synthetic Data
- Example: Enable or Disable Data Access
  This example illustrates how administrators can control data access and prevent Select AI from sending actual schema tables to the LLM.
- Example: Restrict Table Access in AI Profile
  This example demonstrates how to restrict table access and instruct the LLM to use only the tables specified in the `object_list` of the AI profile.
- Example: Specify Case Sensitivity for Columns
  This example shows how you can set case sensitivity for columns in AI profile.

Example: Select AI Actions

These examples illustrate common Select AI actions.

The following example illustrates actions such as `runsql` (the default), `showsql`, `narrate`, `chat`, and `explainsql` that you can perform with `SELECT AI`. These examples use the `sh` schema with AI provider and profile attributes set in the `DBMS_CLOUD_AI.CREATE_PROFILE` function.

```
SQL> select ai how many customers exist;

CUSTOMER_COUNT
--------------
         55500

SQL> select ai showsql how many customers exist;

RESPONSE
--------------------------------------------------
SELECT COUNT(*) AS total_customers
FROM SH.CUSTOMERS


SQL> select ai narrate how many customers exist;

RESPONSE
--------------------------------------------------
There are a total of 55,500 customers in the database.

SQL> select ai chat how many customers exist;

RESPONSE
--------------------------------------------------------------------------------
--
It is impossible to determine the exact number of customers that exist as it
con
stantly changes due to various factors such as population growth, new
businesses
, and customer turnover. Additionally, the term "customer" can refer to
individu
als, businesses, or organizations, making it difficult to provide a specific
num
ber.


SQL> select ai explainsql how many customers in San Francisco are married;

RESPONSE
--------------------------------------------------------------------------------
--
SELECT COUNT(*) AS customer_count
FROM SH.CUSTOMERS AS c
WHERE c.CUST_STATE_PROVINCE = 'San Francisco' AND c.CUST_MARITAL_STATUS =
'Married';

Explanation:
```

```
- We use the 'SH' table alias for the 'CUSTOMERS' table for better
readability.
- The query uses the 'COUNT(*)' function to count the number of rows that
match the given conditions.
- The 'WHERE' clause is used to filter the results:
  - 'c.CUST_STATE_PROVINCE = 'San Francisco'' filters customers who have 'San
Francisco' as their state or province.
  - 'c.CUST_MARITAL_STATUS = 'Married'' filters customers who have 'Married'
as their marital status.
The result of this query will give you the count of customers in San
Francisco who are married, using the column alias 'customer_count' for the
result.

Remember to adjust the table and column names based on your actual schema if
they differ from the example.

Feel free to ask if you have more questions related to SQL or database in
general.
```

Example: Select AI with OCI Generative AI

These examples show how you can access OCI Generative AI using your OCI API key or Resource Principal, create an AI profile, and generate, run, and explain SQL from natural language prompts or chat using the OCI Generative AI LLMs.

> **Note:**
>
> If you do not specify the `model_name` parameter, OCI Generative AI uses the default model as per the table in Select your AI Provider and LLMs. To learn more about the parameters, see Profile Attributes.

```
-- Create Credential with OCI API key
--
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'GENAI_CRED',
    user_ocid       => 'ocid1.user.oc1..aaaa...',
    tenancy_ocid    => 'ocid1.tenancy.oc1..aaaa...',
    private_key     => '<your_api_key>',
    fingerprint     => '<your_fingerprint>'
  );
END;
/

--
-- Create AI profile
--
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
      profile_name
=>'GENAI',
      attributes   =>'{"provider":
"oci",
```

```
        "credential_name": "GENAI_CRED",
        "object_list": [{"owner": "SH", "name": "customers"},
                        {"owner": "SH", "name":
"countries"},
                        {"owner": "SH", "name":
"supplementary_demographics"},
                       {"owner": "SH", "name":
"profits"},
                       {"owner": "SH", "name":
"promotions"},
                       {"owner": "SH", "name": "products"}]
       }');
END;
/

PL/SQL procedure successfully completed.

--
-- Enable AI profile in current session
--
EXEC DBMS_CLOUD_AI.SET_PROFILE('GENAI');

PL/SQL procedure successfully completed.

--
-- Get Profile in current session
--
SELECT DBMS_CLOUD_AI.get_profile() from dual;

DBMS_CLOUD_AI.GET_PROFILE()
--------------------------------------------------------------------------------
--
"GENAI"

--
-- Use AI
--
SQL> select ai how many customers exist;
Number of Customers
-------------------
              55500

SQL> select ai how many customers in San Francisco are married;
COUNT(DISTINCTC."CUST_ID")
--------------------------
                        28

SQL> select ai showsql how many customers in San Francisco are married;
RESPONSE


--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
----------------------------------------
SELECT COUNT(DISTINCT c."CUST_ID")
FROM "SH"."CUSTOMERS" c
```

```
JOIN "SH"."COUNTRIES" co ON c."COUNTRY_ID" = co."COUNTRY_ID"
WHERE c."CUST_CITY" = 'San Francisco' AND c."CUST_MARITAL_STATUS" = 'married'

SQL> select ai explainsql how many customers in San Francisco are married;
RESPONSE


Here is the Oracle SQL query to find the number of customers in San Francisco
who are married:
```
SELECT COUNT(*)
FROM "SH"."CUSTOMERS" c
WHERE c."CUST_CITY" = 'San Francisco'
AND c."CUST_MARITAL_STATUS" = 'Married';
```
Explanation:
* We use the `COUNT(*)` aggregate function to count the number of rows that
match the conditions.
* We specify the table alias `c` for the `"SH"."CUSTOMERS"` table to make the
query more readable.
* We use the `WHERE` clause to filter the rows based on two conditions:
    + `c."CUST_CITY" = 'San Francisco'`: We filter the rows where the
customer's city is San Francisco.
    + `c."CUST_MARITAL_STATUS" = 'Married'`: We filter the rows where the
customer's marital status is Married.
* The double quotes around the table and column names are used to make the
query case-sensitive, as required in Oracle SQL.
Note: The column names and table names are assumed to be case-sensitive, so
we enclose them in double quotes. If the column names and table names are not
case-sensitive, you can remove the double quotes.
Collapse
 This snippet was truncated for display; see it in full

SQL> select ai narrate how many customers in San Francisco are married;

RESPONSE
-----------------------------------------------------------
There are 28 customers in San Francisco who are married.


SQL> select ai chat what is Autonomous Database;

RESPONSE



An Autonomous Database is a type of database that uses artificial
intelligence (AI) and machine learning (ML) to automate many of the
administrative and maintenance tasks typically performed by a database
administrator (DBA). This allows the database to manage itself, without human
intervention, to a large extent.
Autonomous databases are designed to be self-driving, self-securing, and self-
repairing, which means they can:
1. **Automate administrative tasks**: Such as provisioning, patching,
upgrading, and tuning, which frees up DBAs to focus on higher-level tasks.
2. **Optimize performance**: By automatically adjusting parameters, indexing,
```

```
and caching to ensure optimal performance and efficiency.
3. **Detect and respond to security threats**: By using AI-powered security
tools to identify and respond to potential security threats in real-time.
4. **Heal itself**: By automatically detecting and repairing errors,
corruption, or other issues that may arise.
5. **Scale up or down**: To match changing workload demands, without the need
for manual intervention.
The benefits of Autonomous Databases include:
1. **Increased efficiency**: By automating routine tasks, DBAs can focus on
more strategic activities.
2. **Improved performance**: Autonomous databases can optimize performance in
real-time, leading to faster query response times and better overall system
performance.
3. **Enhanced security**: AI-powered security tools can detect and respond to
threats more quickly and effectively than human administrators.
4. **Reduced costs**: By minimizing the need for manual intervention,
Autonomous Databases can help reduce labor costs and improve resource
utilization.
5. **Improved reliability**: Autonomous Databases can detect and repair
errors more quickly, reducing downtime and improving overall system
reliability.
Oracle Autonomous Database is a popular example of an Autonomous Database,
which was introduced in 2018. Other vendors, such as Amazon, Microsoft, and
Google, also offer Autonomous Database services as part of their cloud
offerings.
In summary, Autonomous Databases are designed to be self-managing, self-
optimizing, and self-healing, which can lead to improved performance,
security, and efficiency, while reducing costs and administrative burdens.

--
--Drop the profile
--
EXEC DBMS_CLOUD_AI.DROP_PROFILE('GENAI');

PL/SQL procedure successfully completed.
```

Example: Select AI with OCI Generative AI Resource Principal

To use resource principal with OCI Generative AI, Oracle Cloud Infrastructure tenancy administrator must grant access for Generative AI resources to a dynamic group. See Perform Prerequisites to Use Resource Principal with Autonomous Database to provide access to a dynamic group.

Set the required policies to obtain access to all Generative AI resources. See Getting Access to Generative AI to know more about Generative AI policies.

- To get access to all Generative AI resources in the entire tenancy, use the following policy:

  ```
  allow group <your-group-name> to manage generative-ai-family in tenancy
  ```

- To get access to all Generative AI resources in your compartment, use the following policy:

  ```
  allow group <your-group-name> to manage generative-ai-family in
  compartment <your-compartment-name>
  ```

Connect as an administrator and enable OCI resource principal. See ENABLE_PRINCIPAL_AUTH Procedure to configure the parameters.

> **Note:**
>
> OCI Generative AI uses `meta.llama-3-70b-instruct` as the default model if you do not specify the `model`. To learn more about the parameters, see Profile Attributes.

```
-- Connect as Administrator user and enable OCI resource principal.
BEGIN
  DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(provider  => 'OCI');
END;
/

--
-- Create AI profile
--
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
      profile_name
=>'GENAI',
      attributes =>'{"provider":
"oci",
        "credential_name": "OCI$RESOURCE_PRINCIPAL",
        "object_list": [{"owner": "SH", "name": "customers"},
                        {"owner": "SH", "name":
"countries"},
                        {"owner": "SH", "name":
"supplementary_demographics"},
                       {"owner": "SH", "name":
"profits"},
                       {"owner": "SH", "name":
"promotions"},
                       {"owner": "SH", "name": "products"}]
       }');
END;
/

PL/SQL procedure successfully completed.

--
-- Enable AI profile in current session
--
EXEC DBMS_CLOUD_AI.SET_PROFILE('GENAI');

PL/SQL procedure successfully completed.

--
-- Get Profile in current session
--
SELECT DBMS_CLOUD_AI.get_profile() from dual;

DBMS_CLOUD_AI.GET_PROFILE()
```

```
--------------------------------------------------------------------------------
--
"GENAI"

--
-- Use AI
--
SQL> select ai how many customers exist;
Number of Customers
-------------------
              55500

SQL> select ai how many customers in San Francisco are married;
COUNT(DISTINCTC."CUST_ID")
--------------------------
                        28

SQL> select ai showsql how many customers in San Francisco are married;
RESPONSE


--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
----------------------------------------
SELECT COUNT(DISTINCT c."CUST_ID")
FROM "SH"."CUSTOMERS" c
JOIN "SH"."COUNTRIES" co ON c."COUNTRY_ID" = co."COUNTRY_ID"
WHERE c."CUST_CITY" = 'San Francisco' AND c."CUST_MARITAL_STATUS" = 'married'

SQL> select ai explainsql how many customers in San Francisco are married;
RESPONSE


Here is the Oracle SQL query to find the number of customers in San Francisco
who are married:
```
SELECT COUNT(*)
FROM "SH"."CUSTOMERS" c
WHERE c."CUST_CITY" = 'San Francisco'
AND c."CUST_MARITAL_STATUS" = 'Married';
```
Explanation:
* We use the `COUNT(*)` aggregate function to count the number of rows that
match the conditions.
* We specify the table alias `c` for the `"SH"."CUSTOMERS"` table to make the
query more readable.
* We use the `WHERE` clause to filter the rows based on two conditions:
    + `c."CUST_CITY" = 'San Francisco'`: We filter the rows where the
customer's city is San Francisco.
    + `c."CUST_MARITAL_STATUS" = 'Married'`: We filter the rows where the
customer's marital status is Married.
* The double quotes around the table and column names are used to make the
query case-sensitive, as required in Oracle SQL.
Note: The column names and table names are assumed to be case-sensitive, so
we enclose them in double quotes. If the column names and table names are not
case-sensitive, you can remove the double quotes.
```

```
Collapse
 This snippet was truncated for display; see it in full

SQL> select ai narrate how many customers in San Francisco are married;

RESPONSE
-----------------------------------------------------
There are 28 customers in San Francisco who are married.


SQL> select ai chat what is Autonomous Database;

RESPONSE



An Autonomous Database is a type of database that uses artificial
intelligence (AI) and machine learning (ML) to automate many of the
administrative and maintenance tasks typically performed by a database
administrator (DBA). This allows the database to manage itself, without human
intervention, to a large extent.
Autonomous databases are designed to be self-driving, self-securing, and self-
repairing, which means they can:
1. **Automate administrative tasks**: Such as provisioning, patching,
upgrading, and tuning, which frees up DBAs to focus on higher-level tasks.
2. **Optimize performance**: By automatically adjusting parameters, indexing,
and caching to ensure optimal performance and efficiency.
3. **Detect and respond to security threats**: By using AI-powered security
tools to identify and respond to potential security threats in real-time.
4. **Heal itself**: By automatically detecting and repairing errors,
corruption, or other issues that may arise.
5. **Scale up or down**: To match changing workload demands, without the need
for manual intervention.
The benefits of Autonomous Databases include:
1. **Increased efficiency**: By automating routine tasks, DBAs can focus on
more strategic activities.
2. **Improved performance**: Autonomous databases can optimize performance in
real-time, leading to faster query response times and better overall system
performance.
3. **Enhanced security**: AI-powered security tools can detect and respond to
threats more quickly and effectively than human administrators.
4. **Reduced costs**: By minimizing the need for manual intervention,
Autonomous Databases can help reduce labor costs and improve resource
utilization.
5. **Improved reliability**: Autonomous Databases can detect and repair
errors more quickly, reducing downtime and improving overall system
reliability.
Oracle Autonomous Database is a popular example of an Autonomous Database,
which was introduced in 2018. Other vendors, such as Amazon, Microsoft, and
Google, also offer Autonomous Database services as part of their cloud
offerings.
In summary, Autonomous Databases are designed to be self-managing, self-
optimizing, and self-healing, which can lead to improved performance,
security, and efficiency, while reducing costs and administrative burdens.
```

```
--
--Drop the profile
--
EXEC DBMS_CLOUD_AI.DROP_PROFILE('GENAI');

PL/SQL procedure successfully completed.
```

Example: Select AI with OCI Generative AI Using LLAMA Model
This example showcases the `chat` feature from OCI Generative AI. It highlights the model's capabilities through two prompts: analyzing customer comments to gauze their sentiment and generate an introductory paragraph on rock climbing.

```
BEGIN


DBMS_CLOUD.CREATE_CREDENTIAL(
        credential_name =>
'GENAI_CRED',
        user_ocid       => 'ocid1.user.oc1..aaa',
        tenancy_ocid    => 'ocid1.tenancy.oc1..aaa',
        private_key     => '<your_api_key>',
        fingerprint     => '<your_fingerprint>'
      );


END;

    /

    PL/SQL procedure successfully completed.

BEGIN


DBMS_CLOUD_AI.CREATE_PROFILE(
          profile_name => 'GENAI',
          attributes   =>'{"provider": "oci",
            "object_list": [
                        {"owner": "SH", "name": "customers"},
                        {"owner": "SH", "name": "countries"},
                        {"owner": "SH", "name": "supplementary_demographics"},
                        {"owner": "SH", "name": "profits"},
                        {"owner": "SH", "name": "promotions"},
                        {"owner": "SH", "name": "products"}
                        ]
            "model": "meta.llama-3.3-70b-instruct",
            "oci_apiformat":"GENERIC",
            "credential_name": "GENAI_CRED",
            "oci_compartment_id": "ocid1.compartment.oc1..."}');

END;

    /

    PL/SQL procedure successfully completed.
```

```
--
--Set profile
--
EXEC DBMS_CLOUD_AI.SET_PROFILE('GENAI');

    PL/SQL procedure successfully completed.

SQL> set linesize
150
SQL> SELECT AI chat what is the sentiment of this comment I am not going to
waste my time filling up this three page form. Lousy idea;
SQL>
    RESPONSE

--------------------------------------------------------------------------------
--------------------------------------------------------------------------

    The sentiment of this comment is strongly negative. The user is
expressing frustration and annoyance with the idea of filling out a three-
page form, an
       d is explicitly stating that they consider it a "lousy idea". The use
of the phrase "waste my time" also implies that they feel the task is
unnecessary
       and unproductive. The tone is dismissive and critical.

SQL> SELECT AI chat Write an enthusiastic introductory paragraph on how to
get started with rock climbing with Athletes as the target audience;
    RESPONSE

--------------------------------------------------------------------------------
--------------------------------------------------------------------------
    Rock climbing is an exhilarating and challenging sport that's perfect for
athletes looking to push their limits and test their strength, endurance,
       and mental toughness. Whether you're a seasoned athlete or just
starting out, rock climbing offers a unique and rewarding experience that
will have
       you hooked from the very first climb. With its combination of physical
and mental challenges, rock climbing is a great way to build strength, improve
       flexibility, and develop problem-solving skills. Plus, with the
supportive community of climbers and the breathtaking views from the top of
the climb,
       you'll be hooked from the very first climb. So, if you're ready to
take on a new challenge and experience the thrill of adventure, then it's
time to
       get started with rock climbing!
```

Using OCI Generative AI with the Default Model

The following example uses the default OCI Generative AI Chat Model. If you do not specify the `model_name` parameter, OCI Generative AI uses the default model as per the table in Select your AI Provider and LLMs.

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
     profile_name => 'OCI_DEFAULT',
```

```
      attributes   => '{"provider": "oci",
                        "credential_name": "OCI_CRED",
                        "object_list": [{"owner": "ADB_USER"}]
                        }');
END;
/
```

Using OCI Generative AI with Chat Model
The following example uses `cohere.command-r-plus-08-2024` as the OCI Generative AI Chat Model.

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
      profile_name => 'OCI_COHERE_COMMAND_R_PLUS',
      attributes   => '{"provider": "oci",
                        "credential_name": "OCI_CRED",
                        "object_list": [{"owner": "ADB_USER"}],
                        "model": "cohere.command-r-plus-08-2024"
                       }');
END;
/
```

Using OCI Generative AI with Chat Model Endpoint ID
The following example demonstrates how to specify the OCI Generative AI Chat Model endpoint ID instead of `model`. If you are using Meta Llama Chat Model endpoint ID, then specify `oci_apiformat` as `GENERIC`.

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
      profile_name => 'OCI_CHAT_ENDPOINT',
      attributes => '{"provider": "oci",
                      "credential_name": "OCI_CRED",
                      "object_list": [{"owner": "ADB_USER"}],
                      "oci_endpoint_id": "<endpoint_id>",
                      "oci_apiformat": "GENERIC"
                     }');
END;
/
```

Using OCI Generative AI with Chat Model OCID
This example demonstrates how to specify the OCI Generative AI Cohere Chat Model endpoint ID instead of `model`. If you are using Meta Llama Chat Model endpoint ID, then specify `oci_apiformat` as `GENERIC`.

```
BEGIN

  DBMS_CLOUD_AI.CREATE_PROFILE(
      profile_name => 'OCI_CHAT_OCID',
      attributes   => '{"provider": "oci",
                        "credential_name": "OCI_CRED",
                        "object_list": [{"owner": "ADB_USER"}],
                        "model": "<model_ocid>",
                        "oci_apiformat": "COHERE"
                       }');
```

```
END;
/
```

Example: Select AI with OpenAI

This example shows how you can use OpenAI to generate SQL statements from natural language prompts.

> **Note:**
>
> Only a DBA can run `EXECUTE` privileges and network ACL procedure.

```
--Grants EXECUTE privilege to ADB_USER
--
SQL> grant execute on DBMS_CLOUD_AI to ADB_USER;

-- Grant Network ACL for OpenAI endpoint
--
SQL> BEGIN
     DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
         host => 'api.openai.com',
         ace  => xs$ace_type(privilege_list => xs$name_list('http'),
                             principal_name => 'ADB_USER',
                             principal_type => xs_acl.ptype_db)
     );
    END;
    /

PL/SQL procedure successfully completed.

--
-- Create Credential for AI provider
--
EXEC
DBMS_CLOUD.CREATE_CREDENTIAL(
CREDENTIAL_NAME   => 'OPENAI_CRED',
username          =>  'OPENAI',
password          =>  '<your_api_token>');

PL/SQL procedure successfully completed.

--
-- Create AI profile
--
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
  profile_name   => 'OPENAI',
  attributes     =>'{"provider":
"openai",
        "credential_name": "OPENAI_CRED",
        "object_list": [{"owner": "SH", "name": "customers"},
                        {"owner": "SH", "name": "countries"},
```

```
                        {"owner": "SH", "name": "supplementary_demographics"},
                        {"owner": "SH", "name": "profits"},
                        {"owner": "SH", "name": "promotions"},
                        {"owner": "SH", "name": "products"}],
        "conversation": "true"
       }');
END;
/

PL/SQL procedure successfully completed.

--
-- Enable AI profile in current session
--
SQL> EXEC DBMS_CLOUD_AI.SET_PROFILE('OPENAI');

PL/SQL procedure successfully completed.

--
-- Get Profile in current session
--
SQL> SELECT DBMS_CLOUD_AI.get_profile() from dual;

DBMS_CLOUD_AI.GET_PROFILE()
--------------------------------------------------------------------------------
--
"OPENAI"

--
-- Use AI
--
SQL> select ai how many customers exist;

CUSTOMER_COUNT
--------------
         55500

SQL> select ai how many customers in San Francisco are married;

MARRIED_CUSTOMERS
-----------------
               18


SQL> select ai showsql how many customers in San Francisco are married;

RESPONSE
--------------------------------------------------------------------------------
--
SELECT COUNT(*) AS married_customers_count
FROM SH.CUSTOMERS c
WHERE c.CUST_CITY = 'San Francisco'
  AND c.CUST_MARITAL_STATUS = 'Married'


SQL> select ai explainsql how many customers in San Francisco are married;
```

```
RESPONSE
--------------------------------------------------------------------------------
--
SELECT COUNT(*) AS customer_count
FROM SH.CUSTOMERS AS c
WHERE c.CUST_STATE_PROVINCE = 'San Francisco' AND c.CUST_MARITAL_STATUS =
'Married';

Explanation:
- We use the 'SH' table alias for the 'CUSTOMERS' table for better
readability.
- The query uses the 'COUNT(*)' function to count the number of rows that
match the given conditions.
- The 'WHERE' clause is used to filter the results:
  - 'c.CUST_STATE_PROVINCE = 'San Francisco'' filters customers who have 'San
Francisco' as their state or province.
  - 'c.CUST_MARITAL_STATUS = 'Married'' filters customers who have 'Married'
as their marital status.
The result of this query will give you the count of customers in San
Francisco who are married, using the column alias 'customer_count' for the
result.

Remember to adjust the table and column names based on your actual schema if
they differ from the example.

Feel free to ask if you have more questions related to SQL or database in
general.

SQL> select ai narrate what are the top 3 customers in San Francisco;

RESPONSE
--------------------------------------------------------------------------------
--
The top 3 customers in San Francisco are:

1. Hector Colven - Total amount sold: $52,025.99
2. Milburn Klemm - Total amount sold: $50,842.28
3. Gavin Xie - Total amount sold: $48,677.18


SQL> select ai chat what is Autonomous Database;

RESPONSE
--------------------------------------------------------------------------------
--
Autonomous Database is a cloud-based database service provided by Oracle. It
is designed to automate many of the routine tasks involved in managing a
database,
such as patching, tuning, and backups. Autonomous Database uses machine
learning and automation to optimize performance, security, and availability,
allowing
users to focus on their applications and data rather than database
administration tasks. It offers both Autonomous Transaction Processing (ATP)
for
transactional workloads and Autonomous Data Warehouse (ADW) for analytical
```

```
workloads. Autonomous Database provides high performance, scalability, and
reliability, making it an ideal choice for modern cloud-based applications.

SQL> EXEC DBMS_CLOUD_AI.DROP_PROFILE('OPENAI');

PL/SQL procedure successfully completed.
```

The following example shows specifying a different model in your AI profile:

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
  profile_name   => 'OPENAI',
  attributes     =>'{"provider":
"openai",
        "credential_name": "OPENAI_CRED",
        "model": "gpt-3.5-turbo",
     "object_list": [{"owner": "SH", "name": "customers"},
                     {"owner": "SH", "name": "countries"},
                     {"owner": "SH", "name": "supplementary_demographics"},
                     {"owner": "SH", "name": "profits"},
                     {"owner": "SH", "name": "promotions"},
                     {"owner": "SH", "name": "products"}],
        "conversation": "true"
       }');
END;
/
```

Example: Select AI with Cohere

This example shows how you can use Cohere to generate SQL statements from natural language prompts.

> **Note:**
>
> Only a DBA can run `EXECUTE` privileges and network ACL procedure.

```
--Grants EXECUTE privilege to ADB_USER
--
SQL>GRANT execute on DBMS_CLOUD_AI to ADB_USER;
--
-- Create Credential for AI provider
--
EXEC
DBMS_CLOUD.CREATE_CREDENTIAL(
CREDENTIAL_NAME   => 'COHERE_CRED',
username          =>  'COHERE',
password          =>  'your_api_token');

PL/SQL procedure successfully completed.

--
```

```
-- Grant Network ACL for Cohere endpoint
--
SQL> BEGIN
    DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
         host => 'api.cohere.ai',
         ace  => xs$ace_type(privilege_list => xs$name_list('http'),
                             principal_name => 'ADB_USER',
                             principal_type => xs_acl.ptype_db)
    );
     END;
     /
    /

PL/SQL procedure successfully completed.

--
-- Create AI profile
--
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
  profile_name   => 'COHERE',
  attributes     =>'{"provider": "cohere",
     "credential_name": "COHERE_CRED",
        "object_list": [{"owner": "SH", "name": "customers"},
                        {"owner": "SH", "name": "sales"},
                        {"owner": "SH", "name": "products"},
                        {"owner": "SH", "name": "countries"}]
       }');
END;
/

PL/SQL procedure successfully completed.

--
-- Enable AI profile in current session
--
SQL> EXEC DBMS_CLOUD_AI.SET_PROFILE('COHERE');

PL/SQL procedure successfully completed.

--
-- Get Profile in current session
--
SQL> SELECT DBMS_CLOUD_AI.get_profile() from dual;

DBMS_CLOUD_AI.GET_PROFILE()
--------------------------------------------------------------------------------
--
"COHERE"

--
-- Use AI
--
SQL> select ai how many customers exist;

CUSTOMER_COUNT
```

```
--------------
         55500

SQL> EXEC DBMS_CLOUD_AI.DROP_PROFILE('COHERE');

PL/SQL procedure successfully completed.
```

The following example shows specifying a different model and custom attributes in your AI profile:

```
BEGIN
DBMS_CLOUD_AI.CREATE_PROFILE(
     profile_name => 'COHERE',
     attributes =>
      '{"provider": "cohere",
        "credential_name": "COHERE_CRED",
    "model": "cohere.command-a-03-2025",
        "object_list": [{"owner": "ADB_USER"}],
        "max_tokens":512,
        "stop_tokens": [";"],
        "temperature": 0.5,
        "comments": true
       }');
END;
/
```

Example: Select AI with Azure OpenAI Service

The following examples shows how you can enable access to Azure OpenAI Service using your API key or use Azure OpenAI Service Principal, create an AI profile, and generate SQL from natural language prompts.

```
-- Create Credential for AI integration
--
EXEC
DBMS_CLOUD.CREATE_CREDENTIAL(
CREDENTIAL_NAME   => 'AZURE_CRED',
username          =>  'AZUREAI',
password          =>  'your_api_token');

PL/SQL procedure successfully completed.

--
-- Grant Network ACL for OpenAI endpoint
--
BEGIN
    DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
         host => '<azure_resource_name>.openai.azure.com',
         ace  => xs$ace_type(privilege_list => xs$name_list('http'),
                             principal_name => 'ADB_USER',
                             principal_type => xs_acl.ptype_db)
       );
END;
```

```
/

PL/SQL procedure successfully completed.

--
-- Create AI profile
--
BEGIN

DBMS_CLOUD_AI.CREATE_PROFILE(
     profile_name=>
'AZUREAI',
     attributes=> '{"provider": "azure",
        "azure_resource_name": "<azure_resource_name>",
        "azure_deployment_name":
"<azure_deployment_name>"

        "credential_name": "AZURE_CRED",
        "object_list": [{"owner": "SH", "name": "customers"},
                        {"owner": "SH", "name": "countries"},
                        {"owner": "SH", "name": "supplementary_demographics"},
                        {"owner": "SH", "name": "profits"},
                        {"owner": "SH", "name": "promotions"},
                        {"owner": "SH", "name": "products"}],
        "conversation": "true"
       }');
END;
/

PL/SQL procedure successfully completed.

--
-- Enable AI profile in current session
--
EXEC DBMS_CLOUD_AI.SET_PROFILE('AZUREAI');

PL/SQL procedure successfully completed.

--
-- Get Profile in current session
--
SELECT DBMS_CLOUD_AI.get_profile() from dual;

DBMS_CLOUD_AI.GET_PROFILE()
--------------------------------------------------------------------------------
--
"AZUREAI"


--
-- Use AI
--
select ai how many customers exist;

CUSTOMER_COUNT
--------------
```

```
         55500

select ai how many customers in San Francisco are married;

MARRIED_CUSTOMERS
-----------------
               18


select ai showsql how many customers in San Francisco are married;

RESPONSE
--------------------------------------------------------------------------------
--
SELECT COUNT(*) AS married_customers_count
FROM SH.CUSTOMERS c
WHERE c.CUST_CITY = 'San Francisco'
  AND c.CUST_MARITAL_STATUS = 'Married'


select ai explainsql how many customers in San Francisco are married;

RESPONSE
--------------------------------------------------------------------------------
--
SELECT COUNT(*) AS customer_count
FROM SH.CUSTOMERS AS c
WHERE c.CUST_STATE_PROVINCE = 'San Francisco' AND c.CUST_MARITAL_STATUS =
'Married';

Explanation:
- We use the 'SH' table alias for the 'CUSTOMERS' table for better
readability.
- The query uses the 'COUNT(*)' function to count the number of rows that
match the given conditions.
- The 'WHERE' clause is used to filter the results:
  - 'c.CUST_STATE_PROVINCE = 'San Francisco'' filters customers who have 'San
Francisco' as their state or province.
  - 'c.CUST_MARITAL_STATUS = 'Married'' filters customers who have 'Married'
as their marital status.
The result of this query will give you the count of customers in San
Francisco who are married, using the column alias 'customer_count' for the
result.

Remember to adjust the table and column names based on your actual schema if
they differ from the example.

Feel free to ask if you have more questions related to SQL or database in
general.


select ai narrate what are the top 3 customers in San Francisco;

RESPONSE
--------------------------------------------------------------------------------
--
```

```
The top 3 customers in San Francisco are:

1. Hector Colven - Total amount sold: $52,025.99
2. Milburn Klemm - Total amount sold: $50,842.28
3. Gavin Xie - Total amount sold: $48,677.18


select ai chat what is Autonomous Database;

RESPONSE
--------------------------------------------------------------------------------
--
Autonomous Database is a cloud-based database service provided by Oracle. It
is designed to automate many of the routine tasks involved in managing a
database,
such as patching, tuning, and backups. Autonomous Database uses machine
learning and automation to optimize performance, security, and availability,
allowing
users to focus on their applications and data rather than database
administration tasks. It offers both Autonomous Transaction Processing (ATP)
for transactional
workloads and Autonomous Data Warehouse (ADW) for analytical workloads.
Autonomous Database provides high performance, scalability, and reliability,
making it
an ideal choice for modern cloud-based applications.


EXEC DBMS_CLOUD_AI.DROP_PROFILE('AZUREAI');

PL/SQL procedure successfully completed.
```

The following example shows specifying a different model in your AI profile:

```
BEGIN

DBMS_CLOUD_AI.CREATE_PROFILE(

profile_name=>'AZUREAI',

      attributes=>'{"provider":
"azure",
        "credential_name": "AZURE$PA",
        "model": "gpt-3.5-turbo",
        "object_list": [{"owner": "SH", "name": "customers"},
                        {"owner": "SH", "name": "countries"},
                        {"owner": "SH", "name": "supplementary_demographics"},
                        {"owner": "SH", "name": "profits"},
                        {"owner": "SH", "name": "promotions"},
                        {"owner": "SH", "name": "products"}],
    "azure_resource_name":
"<azure_resource_name>",
        "azure_deployment_name": "<azure_deployment_name>"
       }');
END;
/
```

Example: Select AI with Azure OpenAI Service Principal
Connect as a database administrator to provide access to Azure service principal authentication and then grant the network ACL permissions to the user (`ADB_USER`) who wants to use Select AI. To provide access to Azure resources, see Use Azure Service Principal to Access Azure Resources.

> **Note:**
>
> Only a DBA user can run `EXECUTE` privileges and network ACL procedure.

```
-- Connect as ADMIN user and enable Azure service principal authentication.
BEGIN
  DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(provider  => 'AZURE',
                                         params    =>
JSON_OBJECT('azure_tenantid' value 'azure_tenantid'));
END;
/

-- Copy the consent url from cloud_integrations view and consents the ADB-S
application.
SQL> select param_value from CLOUD_INTEGRATIONS where param_name =
'azure_consent_url';
PARAM_VALUE
--------------------------------------------------------------------------------
--
https://login.microsoftonline.com/<tenant_id>/oauth2/v2.0/authorize?
client_id=<client_id>&response_type=code&scope=User.read

-- On the Azure OpenAI IAM console, search for the Azure application name and
assign the permission to the application.
-- You can get the application name in the cloud_integrations view.
SQL> select param_value from CLOUD_INTEGRATIONS where param_name =
'azure_app_name';
PARAM_VALUE
--------------------------------------------------------------------------------
--
ADBS_APP_DATABASE_OCID

--
-- Grant Network ACL for Azure OpenAI endpoint
--SQL> BEGIN
    DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
         host => 'azure_resource_name.openai.azure.com',
         ace  => xs$ace_type(privilege_list => xs$name_list('http'),
                             principal_name => 'ADB_USER',
                             principal_type => xs_acl.ptype_db)
       );
       END;
       /

PL/SQL procedure successfully completed.

--
```

```
-- Create AI profile
--
BEGIN

DBMS_CLOUD_AI.CREATE_PROFILE(

profile_name=>'AZUREAI',

      attributes=>'{"provider":
"azure",
        "credential_name": "AZURE$PA",
        "object_list": [{"owner": "SH", "name": "customers"},
                        {"owner": "SH", "name": "countries"},
                        {"owner": "SH", "name": "supplementary_demographics"},
                        {"owner": "SH", "name": "profits"},
                        {"owner": "SH", "name": "promotions"},
                        {"owner": "SH", "name": "products"}],
        "azure_resource_name":
"<azure_resource_name>",
        "azure_deployment_name": "<azure_deployment_name>"
       }');
END;
/

PL/SQL procedure successfully completed.

--
-- Enable AI profile in current session
--
EXEC DBMS_CLOUD_AI.SET_PROFILE('AZUREAI');

PL/SQL procedure successfully completed.

--
-- Get Profile in current session
--
SELECT DBMS_CLOUD_AI.get_profile() from dual;

DBMS_CLOUD_AI.GET_PROFILE()
--------------------------------------------------------------------------------
--
"AZUREAI"

--
-- Use AI
--
select ai how many customers exist;

CUSTOMER_COUNT
--------------
         55500

select ai how many customers in San Francisco are married;

MARRIED_CUSTOMERS
-----------------
```

```
               18


select ai showsql how many customers in San Francisco are married;

RESPONSE
--------------------------------------------------------------------------------
--
SELECT COUNT(*) AS married_customers_count
FROM SH.CUSTOMERS c
WHERE c.CUST_CITY = 'San Francisco'
  AND c.CUST_MARITAL_STATUS = 'Married'


select ai explainsql how many customers in San Francisco are married;

RESPONSE
--------------------------------------------------------------------------------
--
SELECT COUNT(*) AS customer_count
FROM SH.CUSTOMERS AS c
WHERE c.CUST_STATE_PROVINCE = 'San Francisco' AND c.CUST_MARITAL_STATUS =
'Married';

Explanation:
- We use the 'SH' table alias for the 'CUSTOMERS' table for better
readability.
- The query uses the 'COUNT(*)' function to count the number of rows that
match the given conditions.
- The 'WHERE' clause is used to filter the results:
  - 'c.CUST_STATE_PROVINCE = 'San Francisco'' filters customers who have 'San
Francisco' as their state or province.
  - 'c.CUST_MARITAL_STATUS = 'Married'' filters customers who have 'Married'
as their marital status.
The result of this query will give you the count of customers in San
Francisco who are married, using the column alias 'customer_count' for the
result.

Remember to adjust the table and column names based on your actual schema if
they differ from the example.

Feel free to ask if you have more questions related to SQL or database in
general.

select ai narrate what are the top 3 customers in San Francisco;

RESPONSE
--------------------------------------------------------------------------------
--
The top 3 customers in San Francisco are:

1. Hector Colven - Total amount sold: $52,025.99
2. Milburn Klemm - Total amount sold: $50,842.28
3. Gavin Xie - Total amount sold: $48,677.18
```

```
select ai chat what is Autonomous Database;

RESPONSE
--------------------------------------------------------------------------------
--
Autonomous Database is a cloud-based database service provided by Oracle. It
is designed to automate many of the routine tasks involved in managing a
database,
such as patching, tuning, and backups. Autonomous Database uses machine
learning and automation to optimize performance, security, and availability,
allowing us
ers to focus on their applications and data rather than database
administration tasks. It offers both Autonomous Transaction Processing (ATP)
for transactional
workloads and Autonomous Data Warehouse (ADW) for analytical workloads.
Autonomous Database provides high performance, scalability, and reliability,
making it
an ideal choice for modern cloud-based applications.

EXEC DBMS_CLOUD_AI.DROP_PROFILE('AZUREAI');

PL/SQL procedure successfully completed.
```

Example: Select AI with Google

This example shows how you can use Google to generate, run, and explain SQL from natural language prompts or chat using the Google Gemini LLM.

The following example demonstrates using Google as your AI provider. The example demonstrates using your Google API signing key to provide network access, creating an AI profile, using Select AI actions to generate SQL queries from natural language prompts and chat responses.

```
--Grants EXECUTE privilege to ADB_USER
--
SQL> grant EXECUTE on DBMS_CLOUD_AI to ADB_USER;

--
-- Create Credential for AI provider
--
SQL> BEGIN
      DBMS_CLOUD.CREATE_CREDENTIAL(
        credential_name => 'GOOGLE_CRED',
        username    => 'GOOGLE',
        password    => '<your_api_key>'
      );
     END;
    /

PL/SQL procedure successfully completed.

--
-- Grant Network ACL for Google endpoint
--
SQL>
SQL> BEGIN
```

```
      DBMS_NETWORK_ACL_ADB_USER.APPEND_HOST_ACE(
        host => 'generativelanguage.googleapis.com',
        ace  => xs$ace_type(privilege_list => xs$name_list('http'),
                    principal_name => 'ADB_USER',
                    principal_type => xs_acl.ptype_db)
       );
     END;
    /

PL/SQL procedure successfully completed.

--
-- Create AI profile
--
SQL> BEGIN
      DBMS_CLOUD_AI.CREATE_PROFILE(
        profile_name =>'GOOGLE',
        attributes   =>'{"provider": "google",
          "credential_name": "GOOGLE_CRED",
          "object_list": [{"owner": "SH", "name":
"customers"},
                          {"owner": "SH", "name": "countries"},
                          {"owner": "SH", "name": "supplementary_demographics"},
                          {"owner": "SH", "name": "profits"},
                          {"owner": "SH", "name": "promotions"},
                          {"owner": "SH", "name": "products"}]
         }');
     END;
   /

PL/SQL procedure successfully completed.

--
-- Enable AI profile in current session
--
SQL> EXEC DBMS_CLOUD_AI.SET_PROFILE('GOOGLE');

PL/SQL procedure successfully completed.

--
-- Use AI
--
SQL> select ai how many customers exist;

CUSTOMER_COUNT
--------------
         55500

SQL> select ai how many customers in San Francisco are married;

MARRIED_CUSTOMERS
-----------------
               18


SQL> select ai showsql how many customers in San Francisco are married;
```

```
RESPONSE
--------------------------------------------------------------------------------
--
SELECT COUNT(*) AS married_customers_count
FROM SH.CUSTOMERS c
WHERE c.CUST_CITY = 'San Francisco'
  AND c.CUST_MARITAL_STATUS = 'Married'


SQL> select ai explainsql how many customers in San Francisco are married;

RESPONSE
--------------------------------------------------------------------------------
--
SELECT COUNT(*) AS customer_count
FROM SH.CUSTOMERS AS c
WHERE c.CUST_STATE_PROVINCE = 'San Francisco' AND c.CUST_MARITAL_STATUS =
'Married';

Explanation:
- We use the 'SH' table alias for the 'CUSTOMERS' table for better
readability.
- The query uses the 'COUNT(*)' function to count the number of rows that
match the given conditions.
- The 'WHERE' clause is used to filter the results:
  - 'c.CUST_STATE_PROVINCE = 'San Francisco'' filters customers who have 'San
Francisco' as their state or province.
  - 'c.CUST_MARITAL_STATUS = 'Married'' filters customers who have 'Married'
as their marital status.
The result of this query will give you the count of customers in San
Francisco who are married, using the column alias 'customer_count' for the
result.

Remember to adjust the table and column names based on your actual schema if
they differ from the example.

Feel free to ask if you have more questions related to SQL or database in
general.

SQL> select ai narrate what are the top 3 customers in San Francisco;

RESPONSE
--------------------------------------------------------------------------------
--
The top 3 customers in San Francisco are:

1. Hector Colven - Total amount sold: $52,025.99
2. Milburn Klemm - Total amount sold: $50,842.28
3. Gavin Xie - Total amount sold: $48,677.18


SQL> select ai chat what is Autonomous Database;

RESPONSE
--------------------------------------------------------------------------------
```

```
--
Autonomous Database is a cloud-based database service provided by Oracle. It
is designed to automate many of the routine tasks involved in managing a
database,
such as patching, tuning, and backups. Autonomous Database uses machine
learning and automation to optimize performance, security, and availability,
allowing
users to focus on their applications and data rather than database
administration tasks. It offers both Autonomous Transaction Processing (ATP)
for transactional
workloads and Autonomous Data Warehouse (ADW) for analytical workloads.
Autonomous Database provides high performance, scalability, and reliability,
making it
an ideal choice for modern cloud-based applications.

--
--Drop the profile
--
EXEC DBMS_CLOUD_AI.DROP_PROFILE('GOOGLE');

PL/SQL procedure successfully
completed.
```

The following example shows specifying a different model in your AI profile:

```
BEGIN
      DBMS_CLOUD_AI.CREATE_PROFILE(
        profile_name =>'GOOGLE',
        attributes   =>'{"provider": "google",
          "credential_name": "GOOGLE_CRED",
          "model": "gemini-1.5-pro",
          "object_list": [{"owner": "SH", "name":
"customers"},
                          {"owner": "SH", "name": "countries"},
                          {"owner": "SH", "name": "supplementary_demographics"},
                          {"owner": "SH", "name": "profits"},
                          {"owner": "SH", "name": "promotions"},
                          {"owner": "SH", "name": "products"}]
                       }');
END;
/
```

Example: Select AI with Anthropic

This example shows how you can use Anthropic to generate, run, and explain SQL from natural language prompts or chat using the Anthropic Claude LLM.

The following example demonstrates using Anthropic as your AI provider. The example demonstrates using your Anthropic API signing key to provide network access, creating an AI profile, and using Select AI actions to generate SQL queries from natural language prompts and chat using the Anthropic Claude LLM.

See Profile Attributes to supply the profile attributes.

```
--Grant EXECUTE privilege to ADB_USER

SQL>GRANT EXECUTE on DBMS_CLOUD_AI to ADB_USER;

--
-- Create Credential for AI provider
--

SQL>BEGIN
      DBMS_CLOUD.CREATE_CREDENTIAL(
        credential_name => 'ANTHROPIC_CRED',
        username    => 'ANTHROPIC',
        password    => '<your api key>'
      );
    END;
     /

PL/SQL procedure successfully completed.


--
-- Grant Network ACL for Anthropic endpoint
--
SQL>BEGIN
      DBMS_NETWORK_ACL_ADB_USER.APPEND_HOST_ACE(
        host => 'api.anthropic.com',
        ace  => xs$ace_type(privilege_list => xs$name_list('http'),
                    principal_name => 'ADB_USER',
                    principal_type => xs_acl.ptype_db)
       );
     END;
    /

PL/SQL procedure successfully completed.

--
-- Create AI profile
--
SQL>BEGIN
      DBMS_CLOUD_AI.CREATE_PROFILE(
        profile_name =>'ANTHROPIC',
        attributes   =>'{"provider": "anthropic",
          "credential_name": "ANTHROPIC_CRED",
          "object_list": [{"owner": "SH", "name":
"customers"},
                          {"owner": "SH", "name": "countries"},
                          {"owner": "SH", "name": "supplementary_demographics"},
                          {"owner": "SH", "name": "profits"},
                          {"owner": "SH", "name": "promotions"},
                          {"owner": "SH", "name": "products"}]
         }');
    END;
     /
```

```
PL/SQL procedure successfully completed.

--
-- Enable AI profile in current session
--
SQL>EXEC DBMS_CLOUD_AI.SET_PROFILE('ANTHROPIC');

PL/SQL procedure successfully completed.

--
-- Use AI
--
SQL> select ai how many customers exist;

CUSTOMER_COUNT
--------------
         55500

SQL> select ai how many customers in San Francisco are married;

MARRIED_CUSTOMERS
-----------------
               18


SQL> select ai showsql how many customers in San Francisco are married;

RESPONSE
--------------------------------------------------------------------------------
--
SELECT COUNT(*) AS married_customers_count
FROM SH.CUSTOMERS c
WHERE c.CUST_CITY = 'San Francisco'
  AND c.CUST_MARITAL_STATUS = 'Married'


SQL> select ai explainsql how many customers in San Francisco are married;

RESPONSE
--------------------------------------------------------------------------------
--
SELECT COUNT(*) AS customer_count
FROM SH.CUSTOMERS AS c
WHERE c.CUST_STATE_PROVINCE = 'San Francisco' AND c.CUST_MARITAL_STATUS =
'Married';

Explanation:
- We use the 'SH' table alias for the 'CUSTOMERS' table for better
readability.
- The query uses the 'COUNT(*)' function to count the number of rows that
match the given conditions.
- The 'WHERE' clause is used to filter the results:
  - 'c.CUST_STATE_PROVINCE = 'San Francisco'' filters customers who have 'San
Francisco' as their state or province.
  - 'c.CUST_MARITAL_STATUS = 'Married'' filters customers who have 'Married'
as their marital status.
```

```
The result of this query will give you the count of customers in San
Francisco who are married, using the column alias 'customer_count' for the
result.

Remember to adjust the table and column names based on your actual schema if
they differ from the example.

Feel free to ask if you have more questions related to SQL or database in
general.

SQL> select ai narrate what are the top 3 customers in San Francisco;

RESPONSE
--------------------------------------------------------------------------------
--
The top 3 customers in San Francisco are:

1. Hector Colven - Total amount sold: $52,025.99
2. Milburn Klemm - Total amount sold: $50,842.28
3. Gavin Xie - Total amount sold: $48,677.18


SQL> select ai chat what is Autonomous Database;

RESPONSE
--------------------------------------------------------------------------------
--
Autonomous Database is a cloud-based database service provided by Oracle. It
is designed to automate many of the routine tasks involved in managing a
database,
such as patching, tuning, and backups. Autonomous Database uses machine
learning and automation to optimize performance, security, and availability,
allowing
users to focus on their applications and data rather than database
administration tasks. It offers both Autonomous Transaction Processing (ATP)
for transactional
workloads and Autonomous Data Warehouse (ADW) for analytical workloads.
Autonomous Database provides high performance, scalability, and reliability,
making it
an ideal choice for modern cloud-based applications.


--
--Drop the profile
--
EXEC DBMS_CLOUD_AI.DROP_PROFILE('ANTHROPIC');

PL/SQL procedure successfully completed.
```

The following example shows specifying a different model in your AI profile:

```
BEGIN
      DBMS_CLOUD_AI.CREATE_PROFILE(
        profile_name =>'ANTHROPIC',
        attributes   =>'{"provider": "anthropic",
```

```
          "credential_name": "ANTHROPIC_CRED",
          "model": "claude-3-opus-20240229",
          "object_list": [{"owner": "SH", "name":
"customers"},
                         {"owner": "SH", "name": "countries"},
                         {"owner": "SH", "name": "supplementary_demographics"},
                         {"owner": "SH", "name": "profits"},
                         {"owner": "SH", "name": "promotions"},
                         {"owner": "SH", "name": "products"}]
        }');
END;
/
```

Example: Select AI with Hugging Face

This example shows how you can use Hugging Face to generate, run, and explain SQL from natural language prompts or chat using the Hugging Face LLM.

The following example demonstrates using Hugging Face as your AI provider. The example demonstrates using your Hugging Face API signing key to provide network access, creating an AI profile, and using Select AI actions to generate SQL queries from natural language prompts and chat using the Hugging Face LLM.

```
--Grant EXECUTE privilege to ADB_USER

SQL>GRANT EXECUTE on DBMS_CLOUD_AI to ADB_USER;

--
-- Create Credential for AI provider
--
SQL>BEGIN
      DBMS_CLOUD.CREATE_CREDENTIAL(
        credential_name => 'HF_CRED',
        username    => 'HF',
        password    => '<your_api_key>'
      );
    END;
     /

PL/SQL procedure successfully completed.

--
-- Grant Network ACL for Hugging Face endpoint
--
SQL>BEGIN
      DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
        host => 'api-inference.huggingface.co',
        ace  => xs$ace_type(privilege_list => xs$name_list('http'),
                    principal_name => 'ADB_USER',
                    principal_type => xs_acl.ptype_db)
       );
    END;
     /

PL/SQL procedure successfully completed.
```

```
--
-- Create AI profile
--
SQL>BEGIN
      DBMS_CLOUD_AI.CREATE_PROFILE(
        profile_name =>'HF',
        attributes   =>'{"provider": "huggingface",
          "credential_name": "HF_CRED",
          "object_list": [{"owner": "SH", "name":
"customers"},
                          {"owner": "SH", "name": "countries"},
                          {"owner": "SH", "name": "supplementary_demographics"},
                          {"owner": "SH", "name": "profits"},
                          {"owner": "SH", "name": "promotions"},
                          {"owner": "SH", "name": "products"}],
           "model" : "Qwen/Qwen2.5-72B-Instruct"
         }');
    END;
     /


PL/SQL procedure successfully completed.


--
-- Enable AI profile in current session
--
SQL>EXEC DBMS_CLOUD_AI.SET_PROFILE('HF');

PL/SQL procedure successfully completed.

--
-- Use AI
--SQL> select ai how many customers exist;

Customer_Count
--------------
         55500

SQL> select ai how many customers in San Francisco are married;

Married_Customers
-----------------
               46

SQL> select ai showsql how many customers in San Francisco are married;

RESPONSE
------------------------------------------------------------
SELECT COUNT("CUST_ID") AS "Married_Customers"
FROM "SH"."CUSTOMERS" "C"
WHERE "CUST_CITY" = 'San Francisco' AND "CUST_MARITAL_STATUS
" = 'Married'
```

```
SQL> select ai explainsql how many customers in San Francisco are married;

RESPONSE
------------------------------------------------------------
To answer the question "How many customers in San Francisco are married?", we
need to query the "SH"."CUSTOMERS" table and filter the results based on the
city and marital status.
We will use table aliases to make the query more readable and ensure that the
string comparisons follow the specified rules.

Here is the Oracle SQL query:

```sql
SELECT COUNT("CUST_ID") AS "Married_Customers_in_San_Francisco"
FROM "SH"."CUSTOMERS" c
WHERE c."CUST_CITY" = 'San Francisco' -- Case insensitive comparison
  AND UPPER(c."CUST_MARITAL_STATUS") = UPPER('Married'); --
Case insensitive comparison
```

Explanation:

1. **Table Alias**:

RESPONSE

   - We use the alias `c` for the "SH"."CUSTOMERS" table to make the query
more readable.

2. **Column Selection**:
   - We select the count of "CUST_ID" to get the number of customers who meet
the criteria. We alias this count as "Married_Customers_in_San_Francisco" for
clarity.

3. **WHERE Clause**:
   - **City Filter**: `c."CUST_CITY" = 'San Francisco'`
     - Since 'San Francisco' is not in double quotes, we use a case-
insensitive comparison.
   - **Marital Status Filter**: `UPPER(c."CUST_MARITAL_STATUS") =
UPPER('Married')`
     - Since 'Married' is not in double quotes, we use a case-insensitive
comparison by applying the `UPPER` function to both sides of the comparison.

This query will return the number of customers in San Francisco who are
married.

SQL> select ai narrate what are the top 3 customers in San Francisco;

RESPONSE

The top 3 customers in San Francisco are:

- Hector Colven, with a customer ID of 8866.
- Milburn Klemm, with a customer ID of 1944.
```

```
- Gavin Xie, with a customer ID of 2579.

These customers have the highest total sales among all customers in San
Francisco.


SQL> select ai chat what is Autonomous Database;

RESPONSE
------------------------------------------------------------
An Autonomous Database is a type of database management system that uses
artificial intelligence (AI) and machine learning (ML) to automate many of
the routine tasks typically performed by database administrators (DBAs).
These tasks include provisioning, patching, backups, tuning, and security.
The goal of an autonomous database is to reduce the need for human
intervention, minimize errors, and optimize performance,thereby allowing
organizations to focus on more strategic activities.

### Key Features of Autonomous Databases:

RESPONSE
------------------------------------------------------------
1. **Automated Provisioning**: The database can be set up and configured
automatically, reducing the time and effort required to get it up and running.
2. **Self-Patching and Updates**: The database automatically applies security
patches and updates without downtime, ensuring that the system remains secure
and up-to-date.
3. **Self-Tuning**: The database continuously monitors its performance and
adjusts settings to optimize query execution and resource utilization.
4. **Self-Backup and Recovery**: Automated backup and recovery processes
ensure that data is protected and can be restored quickly in the event of a
failure.
5. **Security**: Advanced security features, including threat detection and
response, are built into the database to protect against cyber threats.
6. **Scalability**: The database can automatically scale resources up or down
based on demand, ensuring optimal performance and cost efficiency.
7. **Monitoring and Diagnostics**: Real-time monitoring and diagnostics help
identify and resolve issues before they impact performance.


RESPONSE
------------------------------------------------------------
### Benefits of Autonomous Databases:
- **Reduced Operational Costs**: By automating routine tasks, the need for
dedicated DBAs is reduced, lowering operational costs.
- **Improved Reliability**: Automated processes reduce the risk of human
error, leading to more reliable and consistent performance.
- **Enhanced Security**: Continuous monitoring and automated security
measures help protect against threats.
- **Faster Time to Market**: Automated provisioning and tuning allow
applications to be deployed more quickly.

RESPONSE
------------------------------------------------------------
- **Scalability and Flexibility**: The ability to scale resources
automatically ensures that the database can handle varying workloads
```

```
efficiently.

### Use Cases:
- **Cloud Applications**: Autonomous databases are particularly useful in
cloud environments where scalability and reliability are critical.
- **Data Warehousing**: They can handle large volumes of data and complex
queries, making them ideal for data warehousing and analytics.

RESPONSE
------------------------------------------------------------
- **IoT and Real-Time Data Processing**: They can process and analyze real-
time data from IoT devices efficiently.
- **E-commerce**: They can handle high transaction volumes and ensure fast
response times for online shopping platforms.


### Examples of Autonomous Databases:
- **Oracle Autonomous Database**: One of the first and most well-known
autonomous databases, offering both transactional
 and data warehousing capabilities.
- **Amazon Aurora**: A managed relational database service that includes
automated scaling, patching, and backups.
- **Microsoft Azure SQL Database Managed Instance**: Provides a high level of
automation and management for SQL Server databases in the cloud.
- **Google Cloud Spanner**: A globally distributed, horizontally scalable
relational database that is highly available and consistent.

Autonomous databases represent a significant advancement in database
technology, offering organizations a more efficient, secure, and cost-
effective way to manage their data.

SQL> EXEC DBMS_CLOUD_AI.DROP_PROFILE('HF');

PL/SQL procedure successfully completed.
```

Example: Enable Conversations in Select AI

This example illustrates enabling conversations in Select AI.

> **Note:**
>
> Only a DBA can run `EXECUTE` privileges and network ACL procedure.

Create your AI profile. Set the `conversation` attribute to `true` in the profile, this action includes content from prior interactions or prompts, potentially including schema metadata, and set your profile. Once the profile is enabled, you can begin having conversations with your data. Use natural language to ask questions and follow up as needed.

```
--Grants EXECUTE privilege to ADB_USER
--
SQL> grant execute on DBMS_CLOUD_AI to ADB_USER;
```

```
-- Grant Network ACL for OpenAI endpoint
--
SQL> BEGIN
     DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
         host => 'api.openai.com',
         ace  => xs$ace_type(privilege_list => xs$name_list('http'),
                             principal_name => 'ADB_USER',
                             principal_type => xs_acl.ptype_db)
     );
    END;
    /

PL/SQL procedure successfully completed.

--
-- Create Credential for AI provider
--
EXEC
DBMS_CLOUD.CREATE_CREDENTIAL(
CREDENTIAL_NAME   => 'OPENAI_CRED',
username          =>  'OPENAI',
password          =>  '<your_api_token>');

PL/SQL procedure successfully completed.

--
-- Create AI profile
--
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
  profile_name   => 'OPENAI',
  attributes     =>'{"provider":
"openai",
        "credential_name": "OPENAI_CRED",
        "object_list": [{"owner": "SH", "name": "customers"},
                        {"owner": "SH", "name": "countries"},
                        {"owner": "SH", "name": "supplementary_demographics"},
                        {"owner": "SH", "name": "profits"},
                        {"owner": "SH", "name": "promotions"},
                        {"owner": "SH", "name": "products"}],
        "conversation": "true"
       }');

END;
     /

PL/SQL procedure successfully completed.

--
-- Enable AI profile in current session
--
SQL> EXEC DBMS_CLOUD_AI.SET_PROFILE('OPENAI');

PL/SQL procedure successfully completed.

--
```

```
-- Get Profile in current session
--
SQL> SELECT DBMS_CLOUD_AI.get_profile() from dual;

DBMS_CLOUD_AI.GET_PROFILE()
--------------------------------------------------------------------------------
--
"OPENAI"

--
-- Use AI
--
what are the total number of customers;

CUSTOMER_COUNT
--------------
         55500

break out count of customers by country;

RESPONSE
-----------------
COUNTRY_NAME             CUSTOMER_COUNT
Italy                  7780
Brazil                  832
Japan                  624
United Kingdom              7557
Germany             8173
United States of America     18520
France                  3833
Canada                  2010
Spain                  2039
China                  712
Singapore              597
New Zealand              244
Poland     708
Australia              831
Argentina              403
Denmark             383
South Africa              88
Saudi Arabia              75
Turkey                  91


what age group is most common;

RESPONSE
--------------------------------------------------------------------------------
--
AGE_GROUP     CUSTOMER_COUNT
65+         28226

select ai keep the top 5 customers and their country by their purchases and
include a rank in the result;

RESPONSE
```

```
--------------------------------------------------------------------------------
--
RANK     CUSTOMER_NAME         COUNTRY     PURCHASES
1     Abigail Ruddy         Japan         276
2     Abigail Ruddy         Italy         168
3     Abigail Ruddy         Japan         74
3     Abner Robbinette     Germany     74
5     Abner Everett         France         68


SQL> EXEC DBMS_CLOUD_AI.DROP_PROFILE('OPENAI');

PL/SQL procedure successfully completed.
```

Example: Set Up and Use Select AI with RAG

This example guides you through setting up credentials, configuring network access, and creating a vector index for integrating OCI Generative AI vector store cloud services with OpenAI using Oracle Autonomous Database.

The setup concludes with creating an AI profile that uses the vector index to enhance LLM responses. Finally, this example uses the Select AI `narrate` action, which returns a response that has been enhanced using information from the specified vector database.

The following example demonstrates building and querying vector index in Oracle 23ai.

```
--Grants EXECUTE privilege to ADB_USER
GRANT EXECUTE on DBMS_CLOUD_AI to ADB_USER;

--Grants EXECUTE privilege DBMS_CLOUD_PIPELINE to ADB_USER
GRANT EXECUTE on DBMS_CLOUD_PIPELINE to ADB_USER;

-- Create the OpenAI credential
BEGIN
      DBMS_CLOUD.CREATE_CREDENTIAL(
        credential_name => 'OPENAI_CRED',
        username => 'OPENAI_CRED',
        password => '<your_api_key>'
      );
END;
/

PL/SQL procedure successfully completed.

 -- Append the OpenAI endpoint
BEGIN
        DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
             host => 'api.openai.com',
             ace  => xs$ace_type(privilege_list => xs$name_list('http'),
                     principal_name => 'ADB_USER',
                     principal_type => xs_acl.ptype_db)
       );
END;
/

PL/SQL procedure successfully completed.
```

```
-- Create the object store credential
BEGIN
      DBMS_CLOUD.CREATE_CREDENTIAL(
        credential_name => 'OCI_CRED',
        username => '<your_username>',
        password => '<OCI_profile_password>'
      );
END;
/

PL/SQL procedure successfully completed.

 -- Create the profile with the vector index.

BEGIN
      DBMS_CLOUD_AI.CREATE_PROFILE(
          profile_name =>'OPENAI_ORACLE',
          attributes   =>'{"provider": "openai",
            "credential_name": "OPENAI_CRED",
            "vector_index_name": "MY_INDEX",
            "temperature": 0.2,
            "max_tokens": 4096,
            "model": "gpt-3.5-turbo-1106"
          }');
END;
/

PL/SQL procedure successfully completed.

-- Set profile
EXEC DBMS_CLOUD_AI.SET_PROFILE('OPENAI_ORACLE');

PL/SQL procedure successfully
completed.

-- create a vector index with the vector store name, object store location and
-- object store credential
BEGIN
       DBMS_CLOUD_AI.CREATE_VECTOR_INDEX(
         index_name  => 'MY_INDEX',
         attributes  => '{"vector_db_provider": "oracle",
                          "location": "https://swiftobjectstorage.us-
phoenix-1.oraclecloud.com/v1/my_namespace/my_bucket/my_data_folder",
                          "object_storage_credential_name": "OCI_CRED",
                          "profile_name": "OPENAI_ORACLE",
                          "vector_dimension": 1536,
                          "vector_distance_metric": "cosine",
                          "chunk_overlap":128,
                          "chunk_size":1024
      }');
END;
/
PL/SQL procedure successfully completed.
```

```
-- After the vector index is populated, we can now query the index.



-- Set profile
EXEC DBMS_CLOUD_AI.SET_PROFILE('OPENAI_ORACLE');

PL/SQL procedure successfully completed.

-- Select AI answers the question with the knowledge available in the vector
database.

set pages 1000
set linesize 150
SELECT AI narrate how can I deploy an oracle machine learning model;
RESPONSE
To deploy an Oracle Machine Learning model, you would first build your model
within the Oracle database. Once your in-database models are built, they
become immediately available for use, for instance, through a SQL query using
the prediction operators built into the SQL language.

The model scoring, like model building, occurs directly in the database,
eliminating the need for a separate engine or environment within which the
model and corresponding algorithm code operate. You can also use models from
a different schema (user account) if the appropriate permissions are in place.

Sources:
  - Manage-your-models-with-Oracle-Machine-Learning-on-Autonomous-
Database.txt (https://objectstorage.../v1/my_namespace/my_bucket/
my_data_folder/Manage-your-models-with-Oracle-Machine-Learning-on-Autonomous-
Database.txt)
  - Develop-and-deploy-machine-learning-models-using-Oracle-Autonomous-
Database-Machine-Learning-and-APEX.txt (https://objectstorage.../v1/
my_namespace/my_bucket/my_data_folder/Develop-and-deploy-machine-learning-
models-using-Oracle-Autonomous-Database-Machine-Learning-and-APEX.txt)
```

Example: Select AI with In-database Transformer Models

This example demonstrates how you can import a pretrained transformer model that is stored in Oracle object storage into your Oracle Database 23ai instance and then use the imported in-database model in Select AI profile to generate vector embeddings for document chunks and user prompts.

To use in-database transformer models in your Select AI profile, be sure you have:

- your pretrained model imported in your Oracle Database 23ai instance.
- optionally, access to Oracle object storage.

Import a Pretrained Transformer Model into your Oracle Database 23ai From Oracle Object Storage
Review the steps in Import Pretrained Models in ONNX Format for Vector Generation Within the Database and the blog Pre-built Embedding Generation model for Oracle Database 23ai to import a pretrained transformer model into your database.

The following example shows how to import a pretained transformer model from Oracle object storage into your database and then view the imported model.

```
- Create a Directory object, or use an existing directory object
CREATE OR REPLACE DIRECTORY ONNX_DIR AS 'onnx_model';

-- Object storage bucket
VAR location_uri VARCHAR2(4000);
EXEC :location_uri := 'https://adwc4pm.objectstorage.us-
ashburn-1.oci.customer-oci.com/p/
eLddQappgBJ7jNi6Guz9m9LOtYe2u8LWY19GfgU8flFK4N9YgP4kTlrE9Px3pE12/n/adwc4pm/b/
OML-Resources/o/';

-- Model file name
VAR file_name VARCHAR2(512);
EXEC :file_name := 'all_MiniLM_L12_v2.onnx';

-- Download ONNX model from object storage into the directory object
BEGIN
  DBMS_CLOUD.GET_OBJECT(
        credential_name => NULL,
        directory_name  => 'ONNX_DIR',
        object_uri      => :location_uri || :file_name);
END;
/

-- Load the ONNX model into the database
BEGIN
  DBMS_VECTOR.LOAD_ONNX_MODEL(
        directory  => 'ONNX_DIR',
        file_name  => :file_name,
        model_name => 'MY_ONNX_MODEL');
END;
/

-- Verify
SELECT model_name, algorithm, mining_function
FROM user_mining_models
WHERE  model_name='MY_ONNX_MODEL';
```

Use In-database Transformer Models in Select AI Profiles

These examples illustrate how to use in-database transformer models within a Select AI profile. One profile is configured only for generating vector embeddings, while the other supports both Select AI actions and vector index creation.

Review Examples of Privileges to Run Select AI to complete the prerequisites.

The following is an example for generating vector embeddings only:

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
     profile_name => 'EMBEDDING_PROFILE',
     attributes   => '{"provider" : "database",
                       "embedding_model": "MY_ONNX_MODEL"}'
  );
```

```
END;
/
```

The following is an example for general Select AI actions and vector index generation where you can specify a supported AI provider. This example uses OCI Gen AI profile and credentials. See Select your AI Provider and LLMs for list of supported providers. However, if you want to use in-database transformer model for generating vector embeddings, then use "database: <MY_ONNX_MODEL>" in embedding_model attribute:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'GENAI_CRED',
    user_ocid       => 'ocid1.user.oc1..aaaa...',
    tenancy_ocid    => 'ocid1.tenancy.oc1..aaaa...',
    private_key     => '<your_api_key>',
    fingerprint     => '<your_fingerprint>'
  );
END;
/

BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
     profile_name => 'OCI_GENAI',
     attributes   => '{"provider": "oci",
                       "model": "meta.llama-3.3-70b-instruct",
                       "credential_name": "GENAI_CRED",
                       "vector_index_name": "MY_INDEX",
                       "embedding_model": "database: MY_ONNX_MODEL"}'
  );
END;
/
```

Use Select AI with an In-database Transformer Model from Another Schema
This example demonstrates how to use Select AI with an in-database transformer model if another schema owner owns the model. Specify schema_name.object_name as the fully qualified name of the model in embedding_model attribute. If the current user is the schema owner or owns the model, you can omit the schema name.

Be sure to have the following privileges if a different schema owner owns the model:

- CREATE ANY MINING MODEL system privilege
- SELECT ANY MINING MODEL system privilege
- SELECT MINING MODEL object privilege on the specific model

To grant a system privilege, you must either have been granted the system privilege with the ADMIN OPTION or have been granted the GRANT ANY PRIVILEGE system privilege.

See System Privileges for Oracle Machine Learning for SQL to review the privileges.

The following statements allow `ADB_USER1` to score data and view model details in any schema as long as `SELECT` access has been granted to the data. However, `ADB_USER1` can only create models in the `ADB_USER1` schema.

```
GRANT CREATE MINING MODEL TO ADB_USER1;
GRANT SELECT ANY MINING MODEL TO ADB_USER1;


BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
     profile_name => 'OCI_GENAI',
     attributes   => '{"provider": "oci",
                       "credential_name": "GENAI_CRED",
                       "vector_index_name": "MY_INDEX",
                       "embedding_model": "database:
ADB_USER1.MY_ONNX_MODEL"}'
  );
END;
/
```

The following example shows how you can specify case sensitive model object name:

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
     profile_name => 'OCI_GENAI',
     attributes   => '{"provider": "oci",
                       "credential_name": "GENAI_CRED",
                       "model": "meta.llama-3.3-70b-instruct",
                       "vector_index_name": "MY_INDEX",
                       "embedding_model": "database:
\"adb_user1\".\"my_model\""}'
  );
END;
/
```

End-to-end Examples with Different AI Providers

These examples demonstrate end-to-end steps for using in-database transformer model with Select AI RAG. One profile uses *database* as the `provider` exclusively created for generating embedding vectors while the other profile uses *oci* as the `provider` created for Select AI actions as well as vector index.

Review Perform Prerequisites for Select AI to provide the necessary privileges.

```
--Grant create any directory privilege to the user
GRANT CREATE ANY DIRECTORY to ADB_USER;

- Create a Directory object, or use an existing directory object
CREATE OR REPLACE DIRECTORY ONNX_DIR AS 'onnx_model';

-- Object storage bucket
VAR location_uri VARCHAR2(4000);
EXEC :location_uri := 'https://adwc4pm.objectstorage.us-
ashburn-1.oci.customer-oci.com/p/
eLddQappgBJ7jNi6Guz9m9LOtYe2u8LWY19GfgU8flFK4N9YgP4kTlrE9Px3pE12/n/adwc4pm/b/
OML-Resources/o/';
```

```
-- Model file name
VAR file_name VARCHAR2(512);
EXEC :file_name := 'all_MiniLM_L12_v2.onnx';

-- Download ONNX model from object storage into the directory object
BEGIN
  DBMS_CLOUD.GET_OBJECT(
        credential_name => NULL,
        directory_name  => 'ONNX_DIR',
        object_uri      => :location_uri || :file_name);
END;
/

-- Load the ONNX model into the database
BEGIN
  DBMS_VECTOR.LOAD_ONNX_MODEL(
        directory  => 'ONNX_DIR',
        file_name  => :file_name,
        model_name => 'MY_ONNX_MODEL');
END;
/

-- Verify
SELECT model_name, algorithm, mining_function
FROM user_mining_models
WHERE  model_name='MY_ONNX_MODEL';


--Administrator grants EXECUTE privilege to ADB_USER
GRANT EXECUTE on DBMS_CLOUD_AI to ADB_USER;

--Administrator grants EXECUTE privilege DBMS_CLOUD_PIPELINE to ADB_USER
GRANT EXECUTE on DBMS_CLOUD_PIPELINE to ADB_USER;

-- Create the object store credential
BEGIN
      DBMS_CLOUD.CREATE_CREDENTIAL(
        credential_name => 'OCI_CRED',
        username => '<your_username>',
        password => '<OCI_profile_password>'
      );
END;
/

PL/SQL procedure successfully completed.

 -- Create the profile with Oracle Database.

BEGIN
      DBMS_CLOUD_AI.CREATE_PROFILE(
          profile_name =>'EMBEDDING_PROFILE',
          attributes   =>'{"provider": "database",
            "embedding_model": "MY_ONNX_MODEL"
          }');
END;
```

```
/

PL/SQL procedure successfully completed.

-- Set profile
EXEC DBMS_CLOUD_AI.SET_PROFILE('EMBEDDING_PROFILE');

PL/SQL procedure successfully
completed.
```

This example uses *oci* as the `provider`.

```
--Grant create any directory privilege to the user
GRANT CREATE ANY DIRECTORY to ADB_USER;

- Create a Directory object, or use an existing directory object
CREATE OR REPLACE DIRECTORY ONNX_DIR AS 'onnx_model';

-- Object storage bucket
VAR location_uri VARCHAR2(4000);
EXEC :location_uri := 'https://adwc4pm.objectstorage.us-
ashburn-1.oci.customer-oci.com/p/
eLddQappgBJ7jNi6Guz9m9LOtYe2u8LWY19GfgU8flFK4N9YgP4kTlrE9Px3pE12/n/adwc4pm/b/
OML-Resources/o/';

-- Model file name
VAR file_name VARCHAR2(512);
EXEC :file_name := 'all_MiniLM_L12_v2.onnx';

-- Download ONNX model from object storage into the directory object
BEGIN
  DBMS_CLOUD.GET_OBJECT(
        credential_name => NULL,
        directory_name  => 'ONNX_DIR',
        object_uri      => :location_uri || :file_name);
END;
/

-- Load the ONNX model into the database
BEGIN
  DBMS_VECTOR.LOAD_ONNX_MODEL(
        directory  => 'ONNX_DIR',
        file_name  => :file_name,
        model_name => 'MY_ONNX_MODEL');
END;
/

-- Verify
SELECT model_name, algorithm, mining_function
FROM user_mining_models
WHERE  model_name='MY_ONNX_MODEL';


--Administrator Grants EXECUTE privilege to ADB_USER
```

```
GRANT EXECUTE on DBMS_CLOUD_AI to ADB_USER;

--Administrator Grants EXECUTE privilege DBMS_CLOUD_PIPELINE to ADB_USER
GRANT EXECUTE on DBMS_CLOUD_PIPELINE to ADB_USER;

-- Create the object store credential
BEGIN
      DBMS_CLOUD.CREATE_CREDENTIAL(
        credential_name => 'OCI_CRED',
        username => '<your_username>',
        password => '<OCI_profile_password>'
      );
END;
/
--Create GenAI credentials
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'GENAI_CRED',
    user_ocid       => 'ocid1.user.oc1..aaaa...',
    tenancy_ocid    => 'ocid1.tenancy.oc1..aaaa...',
    private_key     => '<your_api_key>',
    fingerprint     => '<your_fingerprint>'
  );
END;
/
--Create OCI AI profile
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
     profile_name => 'OCI_GENAI',
     attributes   => '{"provider": "oci",
                       "model": "meta.llama-3.3-70b-instruct",
                       "credential_name": "GENAI_CRED",
                       "vector_index_name": "MY_INDEX",
                       "embedding_model": "database: MY_ONNX_MODEL"}'
  );
END;
/

-- Set profile
EXEC DBMS_CLOUD_AI.SET_PROFILE('OCI_GENAI');

PL/SQL procedure successfully
completed.

-- create a vector index with the vector store name, object store location and
-- object store credential
BEGIN
       DBMS_CLOUD_AI.CREATE_VECTOR_INDEX(
         index_name  => 'MY_INDEX',
         attributes  => '{"vector_db_provider": "oracle",
                          "location": "https://swiftobjectstorage.us-
phoenix-1.oraclecloud.com/v1/my_namespace/my_bucket/my_data_folder",
                          "object_storage_credential_name": "OCI_CRED",
                          "profile_name": "OCI_GENAI",
                          "vector_dimension": 384,
                          "vector_distance_metric": "cosine",
```

```
                          "chunk_overlap":128,
                          "chunk_size":1024
      }');
END;
/
PL/SQL procedure successfully completed.



-- Set profile
EXEC DBMS_CLOUD_AI.SET_PROFILE('OCI_GENAI');

PL/SQL procedure successfully completed.

-- Select AI answers the question with the knowledge available in the vector
database.

set pages 1000
set linesize 150
SELECT AI narrate how can I deploy an oracle machine learning model;
RESPONSE
To deploy an Oracle Machine Learning model, you would first build your model
within the Oracle database. Once your in-database models are
built, they become immediately available for use, for instance, through a SQL
query using the prediction operators built into the SQL
language.

The model scoring, like model building, occurs directly in the database,
eliminating the need for a separate engine or environment within
which the model and corresponding algorithm code operate. You can also use
models from a different schema (user account) if the appropriate
permissions are in place.

Sources:
  - Manage-your-models-with-Oracle-Machine-Learning-on-Autonomous-
Database.txt (https://objectstorage.../v1/my_namespace/my_bucket/
my_data_folder/Manage-your-models-with-Oracle-Machine-Learning-on-Autonomous-
Database.txt)
  - Develop-and-deploy-machine-learning-models-using-Oracle-Autonomous-
Database-Machine-Learning-and-APEX.txt
(https://objectstorage.../v1/my_namespace/my_bucket/my_data_folder/Develop-
and-deploy-machine-learning-models-using-Oracle-Autonomous-
Database-Machine-Learning-and-APEX.txt)
```

Example: Improve SQL Query Generation

These examples demonstrate how comments, annotations, foreign key, and referential integrity constraints in database tables and columns can improve the generation of SQL queries from natural language prompts.

Example: Improve SQL Generation with Table and Column Comments

If you have table and column comments in your database tables, enable `"comments":"true"` parameter in `DBMS_CLOUD_AI.CREATE_PROFILE` function to retrieve table level and column level comments. The comments are added to the metadata of the LLM for a better SQL generation.

```
-- Adding comments to table 1, table 2, and table 3. Table 1 has 3 columns,
table 2 has 7 columns, table 3 has 2 columns.

-- TABLE1
COMMENT ON TABLE table1 IS 'Contains movies, movie titles and the year it was
released';
COMMENT ON COLUMN table1.c1 IS 'movie ids. Use this column to join to other
tables';
COMMENT ON COLUMN table1.c2 IS 'movie titles';
COMMENT ON COLUMN table1.c3 IS 'year the movie was released';
-- TABLE2
COMMENT ON TABLE table2 IS 'transactions for movie views - also known as
streams';
COMMENT ON COLUMN table2.c1 IS 'day the movie was streamed';
COMMENT ON COLUMN table2.c2 IS 'genre ids. Use this column to join to other
tables';
COMMENT ON COLUMN table2.c3 IS 'movie ids. Use this column to join to other
tables';
COMMENT ON COLUMN table2.c4 IS 'customer ids. Use this column to join to
other tables';
COMMENT ON COLUMN table2.c5 IS 'device used to stream, watch or view the
movie';
COMMENT ON COLUMN table2.c6 IS 'sales from the movie';
COMMENT ON COLUMN table2.c7 IS 'number of views, watched, streamed';

-- TABLE3
COMMENT ON TABLE table3 IS 'Contains the genres';
COMMENT ON COLUMN table3.c1 IS 'genre id. use this column to join to other
tables';
COMMENT ON COLUMN table3.c2 IS 'name of the genre';


BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
    profile_name => 'myprofile',
    attributes =>
        '{"provider": "azure",
          "azure_resource_name": "my_resource",
          "azure_deployment_name": "my_deployment",
          "credential_name": "my_credential",
          "comments":"true",
          "object_list": [
            {"owner": "moviestream", "name": "table1"},
            {"owner": "moviestream", "name": "table2"},
            {"owner": " moviestream", "name": "table3"}
          ]
          }'
    );

    DBMS_CLOUD_AI.SET_PROFILE(
        profile_name => 'myprofile'
    );
```

```
END;
/

--Prompts
select ai what are our total views;
RESPONSE
------------------------------------------------
TOTAL_VIEWS
-----------
   97890562

select ai showsql what are our total views;

RESPONSE
--------------------------------------------------------------------------------
SELECT SUM(QUANTITY_SOLD) AS total_views
FROM "moviestream"."table"

select ai what are our total views broken out by device;

DEVICE                     TOTAL_VIEWS
-------------------------- -----------
mac                           14719238
iphone                        20793516
ipad                          15890590
pc                            14715169
galaxy                        10587343
pixel                         10593551
lenovo                         5294239
fire                           5296916

8 rows selected.

select ai showsql what are our total views broken out by device;

RESPONSE

--------------------------------------------------------------------------------
---------
SELECT DEVICE, COUNT(*) AS TOTAL_VIEWS
FROM "moviestream"."table"
GROUP BY DEVICE
```

Example: Improve SQL Query Generation with Table and Column Annotations
This example demonstrates the integration of annotations in Select AI, applicable in Oracle Database 23ai. The annotations are added to the metadata that is sent to the LLM.

If you have a table with annotations in your schema, enable `"annotations":"true"` in the `DBMS_CLOUD_AI.CREATE_PROFILE` function to instruct Select AI to add annotations to the metadata.

```
--
-- Annotations
--
```

```
CREATE TABLE emp2 (
    empno NUMBER,
    ename VARCHAR2(50) ANNOTATIONS (display 'lastname'),
    salary NUMBER ANNOTATIONS ("person_salary", "column_hidden"),
    deptno NUMBER ANNOTATIONS (display 'department')
)ANNOTATIONS (requires_audit 'yes', version '1.0', ówner 'HR Organization');

Table created.

 BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
   profile_name => 'GOOGLE_ANNOTATIONS',
   attributes   => '{"provider": "google",
      "credential_name": "GOOGLE_CRED",
      "object_list": [{"owner": "ADB_USER", "name": "emp2"}],
      "annotations" : "true"
      }');
END;
/

PL/SQL procedure successfully completed.

EXEC DBMS_CLOUD_AI.SET_PROFILE('GOOGLE_ANNOTATIONS');

PL/SQL procedure successfully completed.
```

Example: Improve SQL Query Generation with Foreign Key and Referential Key Constraints

This example demonstrates the ability of the LLM to generate accurate `JOIN` conditions by retrieving the foreign key and referential key constraints into the metadata of the LLM. The foreign key and referential key constraints provide structured relationship data between the tables to the LLM.

Enable `"constraints":"true"` in the `DBMS_CLOUD_AI.CREATE_PROFILE` function for Select AI to retrieve foreign key and referential key.

```
--
-- Referential Constraints
--
CREATE TABLE dept_test (
    deptno NUMBER PRIMARY KEY,
    dname VARCHAR2(50)
);

Table created.

CREATE TABLE emp3 (
    empno NUMBER PRIMARY KEY,
    ename VARCHAR2(50),
    salary NUMBER,
    deptno NUMBER,
    CONSTRAINT emp_dept_fk FOREIGN KEY (deptno) REFERENCES dept_test(deptno)
);

Table created.
```

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
    profile_name=>'GOOGLE_CONSTRAINTS',
    attribues   =>'{"provider": "google",
      "credential_name": "GOOGLE_CRED",
      "object_list": [{"owner": "ADB_USER", "name": "dept_test"},
                      {"owner": "ADB_USER", "name": "emp3"}],
      "constraints" : "true"
      }');
END;
/

PL/SQL procedure successfully completed.

EXEC DBMS_CLOUD_AI.SET_PROFILE('GOOGLE_CONSTRAINTS');

PL/SQL procedure successfully completed.
```

Example: Automatically Detect Relevant Table Metadata
These examples shows how Select AI automatically detects relevant tables and sends metadata only for those specific tables relevant to the query in Oracle Database 23ai. To enable this feature, set `object_list_mode` to *automated*. This automatically creates a vector index named `<profile_name>_OBJECT_LIST_VECINDEX`. The vector index is initialized with default attributes and values such as `refresh_rate`, `similarity_threshold`, and `match_limit`. You can modify some of the attributes through `DBMS_CLOUD_AI.UPDATE_VECTOR_INDEX`. See UPDATE_VECTOR_INDEX Procedure for more information.

One profile is configured to use `object_list` to specify the schema or the objects in the schema while the other does not specify `object_list`. However, the same SQL construct is expected.

Review Perform Prerequisites for Select AI to provide access to the `DBMS_CLOUD_AI` package and provide network access to the AI provider.

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
    profile_name=>'OCI_AUTO',
    attributes=>'{"provider":
"oci",
      "credential_name": "GENAI_CRED",
      "object_list": [{"owner": "SH"}],
      "oci_compartment_id": "ocid1.compartment.oc1..aaaa...",
      "model" : "meta.llama-3.3-70b-instruct"
      }');
END;
/

PL/SQL procedure successfully completed.

EXEC DBMS_CLOUD_AI.SET_PROFILE('OCI_AUTO');

PL/SQL procedure successfully completed.

select ai showsql how many customers in San Francisco are married;
```

```
RESPONSE
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
-----------------------------------
SELECT COUNT(DISTINCT c."CUST_ID") AS "NUMBER_OF_CUSTOMERS"
FROM "SH"."CUSTOMERS" c
WHERE UPPER(c."CUST_CITY") = UPPER('San Francisco')
AND UPPER(c."CUST_MARITAL_STATUS") = UPPER('married')
```

The following example compares the same scenario without using `object_list`. When you don't specify `object_list`, Select AI automatically chooses all objects available to the current schema.

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
    profile_name=>'OCI_AUTO1',
    attributes=>'{"provider":
"oci",
      "credential_name": "GENAI_CRED",
      "oci_compartment_id": "ocid1.compartment.oc1..aaaa...",
      "object_list_mode": "automated",
      "model" : "meta.llama-3.3-70b-instruct"
      }');
END;
/
PL/SQL procedure successfully completed.

EXEC DBMS_CLOUD_AI.SET_PROFILE('OCI_AUTO1');

PL/SQL procedure successfully completed.

select ai showsql how many customers in San Francisco are married?;

RESPONSE
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
----------
SELECT COUNT(c."CUST_ID") AS "Number_of_Customers"
FROM "SH"."CUSTOMERS" c
WHERE UPPER(c."CUST_CITY") = UPPER('San Francisco')
AND UPPER(c."CUST_MARITAL_STATUS") = UPPER('Married')
```

Example: Generate Synthetic Data

This example explores how you can generate synthetic data mimicking the characteristics and distribution of real data.

The following example shows how to create a few tables in your schema, use OCI Generative AI as your AI provider to create an AI profile, synthesize data into those tables using the `DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA` function, and query or generate responses to natural language prompts with Select AI.

```
--Create tables or use cloned tables

CREATE TABLE ADB_USER.Director (
```

```
    director_id     INT PRIMARY KEY,
    name            VARCHAR(100)
);
CREATE TABLE ADB_USER.Movie (
    movie_id        INT PRIMARY KEY,
    title           VARCHAR(100),
    release_date    DATE,
    genre           VARCHAR(50),
    director_id     INT,
    FOREIGN KEY (director_id) REFERENCES ADB_USER.Director(director_id)
);
CREATE TABLE ADB_USER.Actor (
    actor_id        INT PRIMARY KEY,
    name            VARCHAR(100)
);
CREATE TABLE ADB_USER.Movie_Actor (
    movie_id        INT,
    actor_id        INT,
    PRIMARY KEY (movie_id, actor_id),
    FOREIGN KEY (movie_id) REFERENCES ADB_USER.Movie(movie_id),
    FOREIGN KEY (actor_id) REFERENCES ADB_USER.Actor(actor_id)
);

-- Create the GenAI credential
BEGIN
  DBMS_CLOUD.create_credential(
    credential_name => 'GENAI_CRED',
    user_ocid       => 'ocid1.user.oc1....',
    tenancy_ocid    => 'ocid1.tenancy.oc1....',
    private_key     => 'vZ6cO...',
    fingerprint     => '86:7d:...'
  );
END;
/

-- Create a profile
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
      profile_name
=>'GENAI',
      attributes  =>'{"provider":
"oci",
        "credential_name": "GENAI_CRED",
        "object_list": [{"owner": "ADB_USER",
        "oci_compartment_id": "ocid1.compartment.oc1...."}]
       }');
END;
/


EXEC DBMS_CLOUD_AI.set_profile('GENAI');

-- Run the API for single table
BEGIN
    DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA(
        profile_name => 'GENAI',
```

```
        object_name  => 'Director',
        owner_name   => 'ADB_USER',
        record_count => 5
    );
END;
/
PL/SQL procedure successfully completed.


-- Query the table to see results
SQL> SELECT * FROM ADB_USER.Director;

DIRECTOR_ID NAME
-----------
--------------------------------------------------------------------------------
----------------------
          1 John Smith
          2 Emily Chen
          3 Michael Brown
          4 Sarah Taylor
          5 David Lee


-- Or ask select ai to show the results
SQL> select ai how many directors are there;

NUMBER_OF_DIRECTORS
-------------------
                  5
```

Example: Generate Synthetic Data for Multiple Tables
After you create and set your AI provider profile, use the `DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA` to generate data for multiple tables. You can query or use Select AI to respond to the natural language prompts.

```
BEGIN
    DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA(
        profile_name => 'GENAI',
        object_list => '[{"owner": "ADB_USER", "name":
"Director","record_count":5},
                        {"owner": "ADB_USER", "name":
"Movie_Actor","record_count":5},
                        {"owner": "ADB_USER", "name":
"Actor","record_count":10},
                        {"owner": "ADB_USER", "name":
"Movie","record_count":5,"user_prompt":"all movies released in 2009"}]'
    );
END;
/
PL/SQL procedure successfully completed.


-- Query the table to see results
SQL> select * from ADB_USER.Movie;
```

```
 MOVIE_ID
TITLE                                                    RELEASE_D
                GENRE                                   DIRECTOR_ID
---------- ------------------------------------------------------- ---------
-----------------------------------------------------------
-----------
         1 The Dark Knight                                       15-
JUL-09                            Action                           8
         2 Inglourious Basterds                                  21-
AUG-09                            War                              3
         3 Up in the Air                                         04-
SEP-09                            Drama                            6
         4 The Hangover                                          05-
JUN-09                            Comedy                           1
         5 District 9                                            14-
AUG-09                            Science Fiction                 10



-- Or ask select ai to show the results
SQL> select ai how many actors are there;

Number of Actors
----------------
              10
```

Example: Guide Synthetic Data Generation with Sample Rows
To guide AI service in generating synthetic data, you can randomly select existing records from a table. For instance, by adding `{"sample_rows": 5}` to the `params` argument, you can send 5 sample rows from a table to the AI provider. This example generates 10 additional rows based on the sample rows from the `Transactions` table.

```
BEGIN
  DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA(
    profile_name => 'GENAI',
    object_name  => 'Transactions',
    owner_name   => 'ADB_USER',
    record_count => 10,
    params       => '{"sample_rows":5}'
  );
END;
/
```

Example: Customize Synthetic Data Generation with User Prompts
The `user_prompt` argument enables you to specify additional rules or requirements for data generation. This can be applied to a single table or as part of the `object_list` argument for multiple tables. For example, in the following calls to `DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA`, the prompt instructs the AI to generate synthetic data on movies released in 2009.

```
-- Definition for the Movie table CREATE TABLE Movie

CREATE TABLE Movie (
    movie_id        INT PRIMARY KEY,
```

```
    title           VARCHAR(100),
    release_date    DATE,
    genre           VARCHAR(50),
    director_id     INT,
    FOREIGN KEY (director_id) REFERENCES Director(director_id)
);



BEGIN
  DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA(
    profile_name      => 'GENAI',
    object_name       => 'Movie',
    owner_name        => 'ADB_USER',
    record_count      => 10,
    user_prompt       => 'all movies are released in 2009',
    params            => '{"sample_rows":5}'
  );
END;
/

BEGIN
    DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA(
        profile_name => 'GENAI',
        object_list => '[{"owner": "ADB_USER", "name":
"Director","record_count":5},
                        {"owner": "ADB_USER", "name":
"Movie_Actor","record_count":5},
                        {"owner": "ADB_USER", "name":
"Actor","record_count":10},
                        {"owner": "ADB_USER", "name":
"Movie","record_count":5,"user_prompt":"all movies are released in 2009"}]'
    );
END;
/
```

Example: Improve Synthetic Data Quality by Using Table Statistics
If a table has column statistics or is cloned from a database that includes metadata, Select AI can use these statistics to generate data that closely resembles or is consistent with the original data.

For `NUMBER` columns, the high and low values from the statistics guide the value range. For instance, if the `SALARY` column in the original `EMPLOYEES` table ranges from 1000 to 10000, the synthetic data for this column will also fall within this range.

For columns with distinct values, such as a `STATE` column with values *CA*, *WA*, and *TX*, the synthetic data will use these specific values. You can manage this feature using the `{"table_statistics": true/false}` parameter. By default, the table statistics are enabled.

```
BEGIN
  DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA(
    profile_name      => 'GENAI',
    object_name       => 'Movie',
    owner_name        => 'ADB_USER',
    record_count      => 10,
    user_prompt => 'all movies released in 2009',
```

```
    params           => '{"sample_rows":5,"table_statistics":true}'
  );
END;
/
```

Example: Use Column Comments to Guide Data Generation
If column comments exist, Select AI automatically includes them to provide additional information for the LLM during data generation. For example, a comment on the `Status` column in a Transaction table might list allowed values such as *successful*, *failed*, *pending*, *canceled*, and *need manual check*. You can also add comments to further explain the column, giving AI services more precise instructions or hints for generating accurate data. By default, comments are disabled. See Optional Parameters for more details.

```
-- Use comment on column
COMMENT ON COLUMN Transaction.status IS 'the value for state should either be
''successful'', ''failed'', ''pending'' or ''canceled''';
/

BEGIN
    DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA(
        profile_name  => 'GENAI',
        object_name   => 'employees',
        owner_name    => 'ADB_USER',
        record_count  => 10
        params        => '{"comments":true}'

    );
END;
/
```

Example: Set Unique Values in Synthetic Data Generation
When generating large amounts of synthetic data with LLMs, duplicate values are likely to occur. To prevent this, set up a unique constraint on the relevant column. This ensures that Select AI ignores rows with duplicate values in the LLM response. Additionally, to restrict values for certain columns, you can use the `user_prompt` or add comments to specify the allowed values, such as limiting a `STATE` column to *CA*, *WA*, and *TX*.

```
-- Use 'user_prompt'
BEGIN
    DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA(
        profile_name  => 'GENAI',
        object_name   => 'employees',
        owner_name    => 'ADB_USER',
        user_prompt   => 'the value for state should either be CA, WA, or TX',
        record_count  => 10
    );
END;
/


-- Use comment on column
COMMENT ON COLUMN EMPLOYEES.state IS 'the value for state should either be
CA, WA, or TX'
/
```

Example: Enhance Synthetic Data Generation by Parallel Processing
To reduce runtime, Select AI splits synthetic data generation tasks into smaller chunks for tables without primary keys or with numeric primary keys. These tasks run in parallel, interacting with the AI provider to generate data more efficiently. The Degree of Parallelism (DOP) in your database, influenced by your Autonomous Database service level and ECPU or OCPU settings, determines the number of records each chunk processes. Running tasks in parallel generally improves performance, especially when generating large amounts of data across many tables. To manage the parallel processing of synthetic data generation, set `priority` as an optional parameter. See Optional Parameters.

Example: Enable or Disable Data Access

This example illustrates how administrators can control data access and prevent Select AI from sending actual schema tables to the LLM.

Disabling Data Access
To restrict access to schema tables, log in as an administrator and run the following procedure.

```
EXEC DBMS_CLOUD_AI.DISABLE_DATA_ACCESS;

PL/SQL procedure successfully completed.
```

Disabling data access limits Select AI's `narrate` action and Synthetic Data Generation. The `narrate` action and synthetic data generation raise an error.

Log in as database user, create and configure your AI profile. Review Perform Prerequisites for Select AI to configure your AI profile.

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
          profile_name =>'DATA_ACCESS',
          attributes   =>'{"provider": "openai",
            "credential_name": "OPENAI_CRED",
            "object_list": [{"owner":"SH"}]
          }');
END;
/

EXEC DBMS_CLOUD_AI.SET_PROFILE('DATA_ACCESS');


select ai how many customers;

NUM_CUSTOMERS
55500

select ai narrate what are the top 3 customers in San Francisco;

ORA-20000: Data access is disabled for SELECT AI.
ORA-06512: at "C##CLOUD$SERVICE.DBMS_CLOUD", line 2228
ORA-06512: at "C##CLOUD$SERVICE.DBMS_CLOUD_AI", line 13157
ORA-06512: at line 1 https://docs.oracle.com/error-help/db/ora-20000/
The stored procedure 'raise_application_error' was called which causes this
error to be generated
Error at Line: 1 Column: 6
```

The following example shows the errors that are triggered when you try to generate synthetic data.

```
BEGIN
DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA(
profile_name => 'DATA_ACCESS_SDG',
object_name => 'CUSTOMERS_NEW',
owner_name => 'ADB_USER,
record_count => 5
);
END;
/

ERROR at line 1:

ORA-20000: Data access is disabled for SELECT AI.
ORA-06512: at "C##CLOUD$SERVICE.DBMS_CLOUD", line 2228
ORA-06512: at "C##CLOUD$SERVICE.DBMS_CLOUD_AI", line 13401

ORA-06512: at line 2
```

Enabling Data Access

The following example shows enabling data access. Log in as an administrator and run the following procedure:

```
EXEC DBMS_CLOUD_AI.ENABLE_DATA_ACCESS;

PL/SQL procedure successfully completed.
```

Log in as database user, create and configure your AI profile. Review Perform Prerequisites for Select AI to configure your AI profile. Run `narrate` action and separately generate synthetic data.

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
          profile_name =>'DATA_ACCESS_NEW',
          attributes   =>'{"provider": "openai",
            "credential_name": "OPENAI_CRED",
            "object_list": [{"owner":"SH"}]
          }');
   END;
   /

PL/SQL procedure successfully completed.

EXEC DBMS_CLOUD_AI.SET_PROFILE('DATA_ACCESS_NEW');

PL/SQL procedure successfully completed.


select ai how many customers;

NUM_CUSTOMERS
55500
```

```
select ai narrate what are the top 3 customers in San Francisco;

"RESPONSE"
"The top 3 customers in San Francisco are Cody Seto, Lauren Yaskovich, and
Ian Mc"
```

The following example shows successful synthetic data generation after enabling data access.

```
BEGIN
DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA(
profile_name => 'DATA_ACCESS_SDG',
object_name => 'CUSTOMERS_NEW',
owner_name => 'ADB_USER',
record_count => 5
);
END;
/

PL/SQL procedure successfully completed.
```

Example: Restrict Table Access in AI Profile

This example demonstrates how to restrict table access and instruct the LLM to use only the tables specified in the `object_list` of the AI profile.

Set `enforce_object_list` to *true* to restrict table access to the LLM.

As a database user, create and configure your AI profile. Review Perform Prerequisites for Select AI to configure your AI profile.

```
BEGIN
      DBMS_CLOUD_AI.CREATE_PROFILE(
        profile_name =>'GOOGLE_ENFORCED',
        attributes   =>'{"provider": "google",
          "credential_name": "GOOGLE_CRED",
          "object_list": [{"owner": "ADB_USER", "name": "GENRE"},
                  {"owner": "ADB_USER", "name": "CUSTOMER"},
                  {"owner": "ADB_USER", "name": "PIZZA_SHOP"},
                  {"owner": "ADB_USER", "name": "STREAMS"},
                  {"owner": "ADB_USER", "name": "MOVIES"},
                  {"owner": "ADB_USER", "name": "ACTIONS"}],
          "enforce_object_list" : "true"
          }');
END;
/

PL/SQL procedure successfully completed.


EXEC DBMS_CLOUD_AI.set_profile('GOOGLE_ENFORCED');

PL/SQL procedure successfully completed.

select ai showsql please list the user tables;
```

```
RESPONSE
--------------------------------------------------------------------------------
--------------
SELECT 'ADB_USER.GENRE' AS TABLE_NAME FROM DUAL UNION ALL SELECT
'ADB_USER.CUSTOMER' AS
TABLE_NAME FROM DUAL UNION ALL SELECT 'ADB_USER.PIZZA_SHOP' AS TABLE_NAME
FROM DUAL UNION
 ALL SELECT 'ADB_USER.STREAMS' AS TABLE_NAME FROM DUAL UNION ALL SELECT
'ADB_USER.MOVIES'
AS TABLE_NAME FROM DUAL
--
```

Setting `enforce_object_list` to *false* instructs the LLM to use other tables and views based on its prior knowledge.

```
BEGIN
      DBMS_CLOUD_AI.CREATE_PROFILE(
        profile_name =>'GOOGLE_ENFORCED1',
        attributes   =>'{"provider": "google",
          "credential_name": "GOOGLE_CRED",
          "object_list": [{"owner": "ADB_USER", "name": "GENRE"},
                  {"owner": "ADB_USER", "name": "CUSTOMER"},
                  {"owner": "ADB_USER", "name": "PIZZA_SHOP"},
                  {"owner": "ADB_USER", "name": "STREAMS"},
                  {"owner": "ADB_USER", "name": "MOVIES"},
                  {"owner": "ADB_USER", "name": "ACTIONS"}],
          "enforce_object_list" : "false"
          }');
END;
/

PL/SQL procedure successfully completed.


EXEC DBMS_CLOUD_AI.set_profile('GOOGLE_ENFORCED1');

PL/SQL procedure successfully completed.

select ai showsql please list the user tables;
RESPONSE
----------------------------------
SELECT TABLE_NAME FROM USER_TABLES
```

Example: Specify Case Sensitivity for Columns

This example shows how you can set case sensitivity for columns in AI profile.

Set `case_sensitive_values` to *false* to retrieve queries that are not case sensitive.

As a database user, create and configure your AI profile. Review Perform Prerequisites for Select AI to configure your AI profile.

```
BEGIN
      DBMS_CLOUD_AI.create_profile(
```

```
        profile_name   =>'GOOGLE',
        attributes     =>'{"provider": "google",
          "credential_name": "GOOGLE_CRED",
          "object_list": [{"owner": "ADB_USER", "name": "GENRE"},
                  {"owner": "ADB_USER", "name": "CUSTOMER"},
                  {"owner": "ADB_USER", "name": "PIZZA_SHOP"},
                  {"owner": "ADB_USER", "name": "STREAMS"},
                 {"owner": "ADB_USER", "name": "MOVIES"},
                 {"owner": "ADB_USER", "name": "ACTIONS"}],
         "case_sensitive_values" : "false"
         }');
END;
/

PL/SQL procedure successfully completed.

-- With "case_sensitive_values" set to "false", LLM will give back case
insensitive query.

select ai showsql how many people watch Inception;

RESPONSE
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
--------------------------------------------------------------------
SELECT COUNT(DISTINCT c.CUSTOMER_ID) AS "COUNT"
FROM "ADB_USER"."CUSTOMER" c
JOIN "ADB_USER"."STREAMS" s ON c.CUSTOMER_ID = s.CUSTOMER_ID
JOIN "ADB_USER"."MOVIES" m ON s.MOVIE_ID = m.MOVIE_ID
WHERE UPPER(m.TITLE) = UPPER('Inception')
```

You can specify case sensitive query using double quotes even though the `case_sensitive_values` is set to *false*.

```
select ai showsql how many people watch "Inception";

RESPONSE
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
-----------------------------------------------------
SELECT COUNT(DISTINCT c.CUSTOMER_ID) AS "COUNT"
FROM "ADB_USER"."CUSTOMER" c JOIN "ADB_USER"."STREAMS" s ON
c.CUSTOMER_ID = s.CUSTOMER_ID JOIN "ADB_USER"."MOVIES" m ON
s.MOVIE_ID = m.MOVIE_ID WHERE m.TITLE = 'Inception'
```

JSON Document Stores

Autonomous Database has full support for data represented as JSON documents. In Autonomous Databases, JSON documents can coexist with relational data.

For more details on JSON Documents, see JSON Developer's Guide.

Cache Messages with Singleton Pipes and Using Pipes for Persistent Messaging

On Autonomous Database the `DBMS_PIPE` package includes support for Singleton Pipes and Persistent Messaging Pipes.

- **Singleton Pipes** allow you to cache and retrieve a custom message and share the message across multiple database sessions with concurrent reads.
- **Persistent Messaging Pipes** allow you to store messages in Cloud Object Storage. Using persistent messages you can send and receive messages within a single database or between multiple databases in the same or in different regions, without message size limitations.
- Cache Messages with Singleton Pipes
  Singleton Pipe is an addition to the `DBMS_PIPE` package that allows you to cache and retrieve a custom message and share the message across multiple database sessions with concurrent reads.
- Use Persistent Messaging with Messages Stored in Cloud Object Store
  The `DBMS_PIPE` package has extended functionality on Autonomous Database to support persistent messaging, where messages are stored in Cloud Object Store.

Cache Messages with Singleton Pipes

Singleton Pipe is an addition to the `DBMS_PIPE` package that allows you to cache and retrieve a custom message and share the message across multiple database sessions with concurrent reads.

- About Caching Messages with Singleton Pipes
  The `DBMS_PIPE` package has extended functionality on Autonomous Database to support Singleton Pipes.
- Automatic Refresh of Cached Message with a Cache Function
  The `DBMS_PIPE` package allows you to automatically populate a Singleton Pipe message using a user-defined cache function.
- Create an Explicit Singleton Pipe
  Describes the steps to create a Singleton Pipe with a specified pipe name (an Explicit Singleton Pipe).
- Create an Explicit Singleton Pipe with a Cache Function
  Describes the steps to create a Singleton Pipe with a specified pipe name, an Explicit Singleton Pipe, and provide a cache function. A cache function allows you to automatically populate the message in a singleton pipe.

About Caching Messages with Singleton Pipes

The `DBMS_PIPE` package has extended functionality on Autonomous Database to support Singleton Pipes.

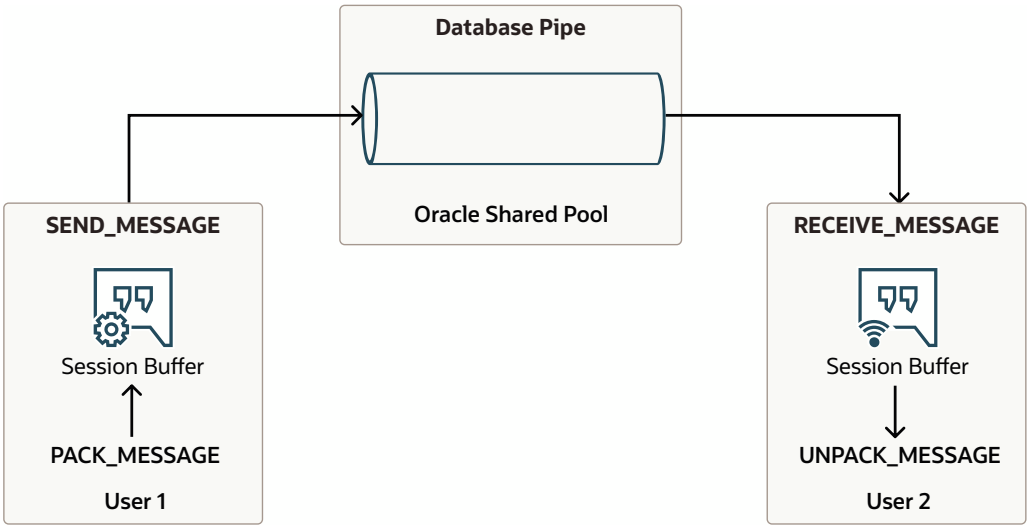
A Singleton Pipe in `DBMS_PIPE`:

- Provides in-memory caching of custom data using Singleton Pipe messages.
- Supports the ability to cache and retrieve a custom message of up to 32,767 bytes.

- Supports sharing a cached message across multiple database sessions with concurrent reads. This provides high throughput and supports concurrent reads of messages across database sessions.
- Supports Read-Only and Read-Write databases.
- Supports several cache invalidation methods:
  - Explicit cache invalidation controlled by user.
  - Cache invalidation after a user specified time interval (in seconds). This invalidation method is controlled by the message sender, using the `shelflife` parameter, instead of by message readers. This avoids the common pitfalls due to incorrect use of cache by readers.

About Standard Pipes and Singleton Pipes

The DBMS_PIPE Package allows two or more database sessions to communicate using in-memory messages. Pipe functionality has several applications such as external service interface, debugging, independent transactions, and alerts.

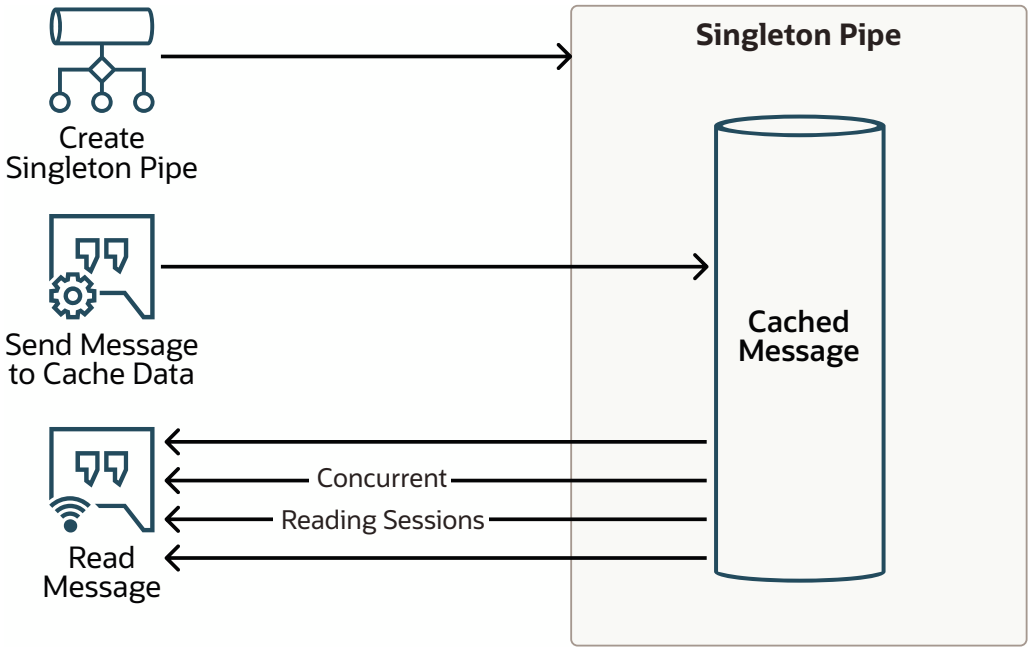


A Singleton Pipe can be any one of the supported `DBMS_PIPE` types:

- **Implicit Pipe:** Automatically created when a message is sent with an unknown pipe name using the `DBMS_PIPE.SEND_MESSAGE` function.
- **Explicit Pipe:** Created using the `DBMS_PIPE.CREATE_PIPE` function with a user specified pipe name.
- **Public Pipe:** Accessible by any user with `EXECUTE` permission on `DBMS_PIPE` package.
- **Private Pipe:** Accessible by sessions with the same user as the pipe creator.

Singleton Pipes provide the ability to cache a single message in the memory of the Autonomous Database instance.

The following shows the general workflow for using singleton pipes.



Singleton Pipe Overview and Features

- **Singleton Messages**
  - A Singleton Pipe can cache one message in the pipe, hence the name "singleton".
  - The message in a Singleton Pipe can be comprised of multiple fields, up to a total message size of 32,767 bytes.
  - `DBMS_PIPE` supports the ability to pack multiple attributes in a message using `DBMS_PIPE.PACK_MESSAGE` procedure.
  - For a Public Singleton Pipe, the message can be received by any database session with execute privilege on `DBMS_PIPE` package.
  - For Private Singleton Pipe, the message can be received by sessions with the same user as the creator of the Singleton Pipe.
- **High Message Throughput for Reads**
  - Singleton Pipes cache the message in the pipe until it is invalidated or purged. Database sessions can concurrently read a message from the Singleton Pipe.
  - Receiving a message from a Singleton Pipe is a non-blocking operation.
- **Message Caching**
  - A message is cached in a Singleton Pipe using `DBMS_PIPE.SEND_MESSAGE`.
  - If there is an existing cached message in the Singleton Pipe, then `DBMS_PIPE.SEND_MESSAGE` overwrites the previous message to maintain only one message in the Singleton Pipe.
- **Message Invalidation**
  - **Explicit Invalidation**: purges the pipe with the procedure `DBMS_PIPE.PURGE` or by overwriting the message using `DBMS_PIPE.SEND_MESSAGE`.
  - **Automatic Invalidation**: a message can be invalidated automatically after the specified `shelflife` time has elapsed.
- **No Eviction from Database Memory**

- Singleton Pipes do not get evicted from Oracle Database memory.
- An Explicit Singleton Pipe continues to reside in database memory until it is removed using `DBMS_PIPE.REMOVE_PIPE` or until the database restarts.
- An Implicit Singleton Pipe stays in database memory until there is one cached message in the pipe.

Singleton Pipe Operations

Operation	DBMS_PIPE Function or Procedure
Create an Explicit Singleton Pipe	CREATE_PIPE Function
Cache a message in Singleton Pipe	PACK_MESSAGE Procedures, SEND_MESSAGE Function
Read a cached message from Singleton Pipe	RECEIVE_MESSAGE Function, UNPACK_MESSAGE Procedures
Delete a message in Singleton Pipe	PURGE Procedure
Remove an Explicit Singleton Pipe	REMOVE_PIPE Function

Automatic Refresh of Cached Message with a Cache Function

The `DBMS_PIPE` package allows you to automatically populate a Singleton Pipe message using a user-defined cache function.

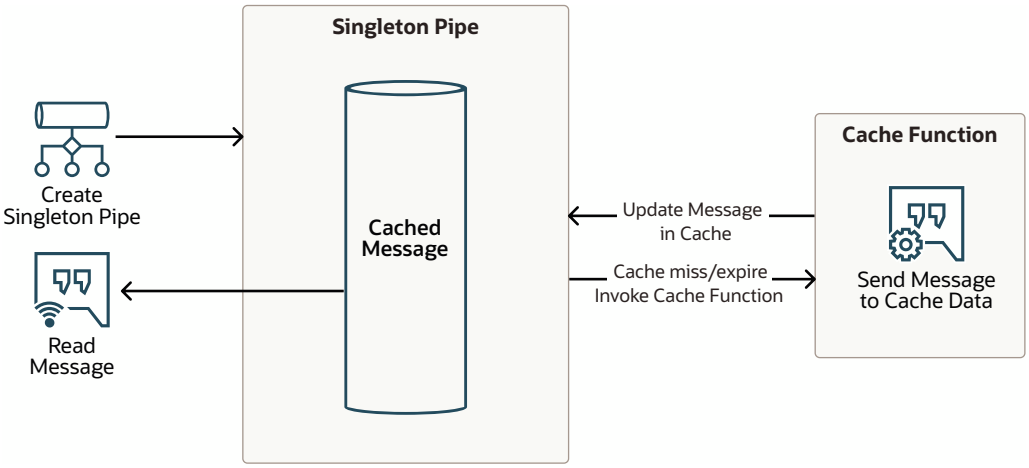
By default, after a message is invalidated with either Singleton Pipe explicit or implicit invalidation, a subsequent `DBMS_PIPE.RECEIVE_MESSAGE` results in no message being received. To add a new message to the pipe, the message must be explicitly cached by calling `DBMS_PIPE.SEND_MESSAGE`. To avoid this case, where no message is available when you read from a Singleton Pipe, you can define a cache function. With a cache function defined, the cache function is automatically invoked when you receive a message in following scenarios:

- When the Singleton Pipe is empty.
- When the message in a Singleton Pipe is invalid due to the `shelflife` time elapsed.

To use a cache function define the cache function and include the `cache_func` parameter with `DBMS_PIPE.RECEIVE_MESSAGE`. A user-defined cache function provides the following:

- The cache function can be specified when reading a message from a Singleton Pipe using `DBMS_PIPE.RECEIVE_MESSAGE`.
- When there is no message in the Singleton Pipe, `DBMS_PIPE.RECEIVE_MESSAGE` calls the cache function.
- When the message `shelflife` time has elapsed, the database automatically populates a new message in the Singleton Pipe.

Using a cache function simplifies working with Singleton Pipes. You do not need to handle failure cases for receiving a message from an empty pipe. In addition, a cache function ensures there is no cache-miss when you read messages from a Singleton Pipe, providing maximum use of the cached message.



When you define a cache function, the function name must be fully qualified with the owner schema:

- *OWNER.FUNCTION_NAME*
- *OWNER.PACKAGE.FUNCTION_NAME*

Define a cache function with the following signature:

```
CREATE OR REPLACE FUNCTION cache_function_name(
       pipename  IN VARCHAR2
) RETURN INTEGER;
```

The typical operations within a cache function are:

- Create a Singleton Pipe, for an Explicit Pipe, using `DBMS_PIPE.CREATE_PIPE`.
- Create a message to cache in the Singleton Pipe.
- Send the message to the pipe specified in the cache function, optionally specifying a `shelflife` for the implicit message.

To use a cache function, the current session user that invokes `DBMS_PIPE.RECEIVE_MESSAGE` must have required privileges to execute the cache function.

See RECEIVE_MESSAGE Function more information on defining a cache function.

Create an Explicit Singleton Pipe

Describes the steps to create a Singleton Pipe with a specified pipe name (an Explicit Singleton Pipe).

First, for this example create the `receive_message` helper function to repeatedly call `DBMS_PIPE.RECEIVE_MESSAGE`. This allows you to test singleton pipe functionality.

```
CREATE OR REPLACE FUNCTION msg_types AS
       TYPE t_rcv_row IS RECORD (c1 VARCHAR2(32767), c2 NUMBER);
       TYPE t_rcv_tab IS TABLE OF t_rcv_row;
END;


CREATE OR REPLACE FUNCTION receive_message(
      pipename    IN VARCHAR2,
      rcv_count   IN NUMBER DEFAULT 1,
```

```
     cache_func  IN VARCHAR2 DEFAULT NULL)
  RETURN msg_types.t_rcv_tab pipelined
   AS
      l_msg    VARCHAR2(32767);
      l_status NUMBER;
BEGIN
     FOR i IN 1..rcv_count LOOP
          l_status := DBMS_PIPE.RECEIVE_MESSAGE(
           pipename   => pipename,
           cache_func => cache_func,
           timeout    => 1);
        IF l_status != 0 THEN
             raise_application_error(-20000,
            'Message not received for attempt: ' || to_char(i) || ' status:
' ||
           l_status);
        END IF;

        DBMS_PIPE.UNPACK_MESSAGE(l_msg);
            pipe row(msg_types.t_rcv_row(l_msg));
    END LOOP;
RETURN;
END;
```

1. Create an explicit singleton pipe named `PIPE_TEST` with `shelflife` parameter set to 3600 (seconds).

   ```
   DECLARE
     l_status INTEGER;
   BEGIN
     l_status := DBMS_PIPE.CREATE_PIPE(
                   pipename => 'MY_PIPE1',
                   private => TRUE,
                   singleton => TRUE,
                   shelflife => 3600);
   END;
   /
   ```

   See CREATE_PIPE Function for more information.

2. Verify the singleton pipe is created.

   ```
   SELECT name, singleton, type
        FROM v$db_pipes WHERE name= '&pipename' ORDER BY 1;

   NAME                 SINGLETON  TYPE
   -------------------- ---------- -------
   PIPE_TEST            YES        PRIVATE
   ```

3. Pack and send a message on the singleton pipe.

   ```
   EXEC DBMS_PIPE.PACK_MESSAGE('This is a real message that you can get
   multiple times');
   ```

```
SELECT DBMS_PIPE.SEND_MESSAGE(pipename => '&pipename') status FROM DUAL;

STATUS
----------
0
```

See PACK_MESSAGE Procedures and SEND_MESSAGE Function for more information.

4. Receive a message from a singleton pipe.

```
SELECT * FROM receive_message(
    pipename => '&pipename',
    rcv_count => 2);

MESSAGE
--------------------------------------------------------------------------------
-----
This is a real message that you can get multiple times
This is a real message that you can get multiple times
```

The `receive_message` function is a helper function that calls `DBMS_PIPE.RECEIVE_MESSAGE`.

5. Purge the message and remove the pipe.

```
EXEC DBMS_PIPE.PURGE('&pipename');
SELECT DBMS_PIPE.REMOVE_PIPE('&pipename') status FROM DUAL;
```

Create an Explicit Singleton Pipe with a Cache Function

Describes the steps to create a Singleton Pipe with a specified pipe name, an Explicit Singleton Pipe, and provide a cache function. A cache function allows you to automatically populate the message in a singleton pipe.

1. Create a cache function, `test_cache_message` for a singleton pipe.

```
CREATE OR REPLACE FUNCTION test_cache_message(
     pipename IN VARCHAR2) return NUMBER
AS
   l_status NUMBER;
   l_data VARCHAR2(4000);
BEGIN
   l_status := DBMS_PIPE.CREATE_PIPE(
          pipename => pipename,
          private => TRUE,
          singleton => true,
          shelflife => 600);
   IF l_status != 0 THEN RETURN l_status;
   END IF;

   DBMS_PIPE.PACK_MESSAGE('This is a placeholder cache message for an
empty pipe');
   l_status := DBMS_PIPE.SEND_MESSAGE(pipename => pipename);
   RETURN l_status;
 END;
/
```

> **Note:**
>
> The current session user invoking `DBMS_PIPE.RECEIVE_MESSAGE` must have required privilege to execute the cache function.

2. Receive with a cache function and confirm the message populates in pipe. The pipe must exist as a private pipe created in the cache function.

```
SELECT * FROM receive_message(
     pipename => '&pipename',
     rcv_count => 1,
     cache_func => 'TEST_CACHE_MESSAGE');

MESSAGE
---------------
This is a placeholder cache message for an empty pipe
```

   The `receive_message` function is a helper function that calls `DBMS_PIPE.RECEIVE_MESSAGE`. See Create an Explicit Singleton Pipe for the `receive_message` definition.

   See CREATE_PIPE Function for more information.

3. Receive without the cache function to confirm the message persists in the pipe.

```
SELECT * FROM receive_message(
     pipename => '&pipename',
     rcv_count => 2);

MESSAGE
---------------
This is a placeholder cache message for an empty pipe
This is a placeholder cache message for an empty pipe
```

   The `receive_message` function is a helper function that calls `DBMS_PIPE.RECEIVE_MESSAGE`. See Create an Explicit Singleton Pipe for the `receive_message` definition.

   See CREATE_PIPE Function for more information.

Use Persistent Messaging with Messages Stored in Cloud Object Store

The `DBMS_PIPE` package has extended functionality on Autonomous Database to support persistent messaging, where messages are stored in Cloud Object Store.

- About Persistent Messaging with DBMS_PIPE
- Create an Explicit Persistent Pipe and Send a Message
  Describes the steps to create a persistent pipe with a specified pipe name (Explicit Pipe).
- Retrieve a Persistent Message on Same Database
  Describes the steps to retrieve a persistent message from an explicit pipe on the same Autonomous Database instance (the instance where the message was sent).
- Retrieve a Persistent Message by Creating a Pipe on a Different Database
  Describes the steps to retrieve a persistent message stored in Cloud Object Store with an explicit pipe on an Autonomous Database instance that is different than the instance that sent the message.

- Remove a Persistent Pipe
  Describes the steps to remove a persistent pipe.

About Persistent Messaging with DBMS_PIPE

Persistent messaging with `DBMS_PIPE` allows one or more database sessions to communicate in the same region or across regions with messages that are stored in Cloud Object Store. Persistent messages in `DBMS_PIPE`:

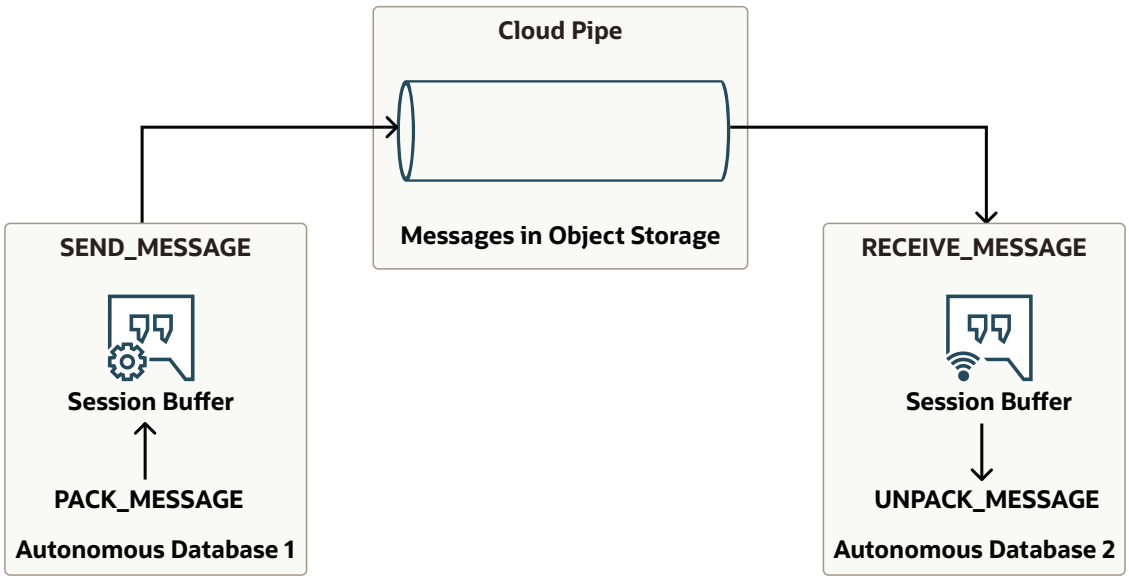
- Allow you to send and retrieve very large messages.
- Support a sending a large number of pipe messages.
- Support sending and receiving messages within a single database, across multiple databases and across databases in different regions.
- Support multiple pipes using the same Cloud Object Store location URI.

Persistent messaging pipes can be created in any of the supported `DBMS_PIPE` types:

- **Implicit Pipe:** Automatically created when a message is sent with an unknown pipe name using the `DBMS_PIPE.SEND_MESSAGE` function.
- **Explicit Pipe:** Created using the `DBMS_PIPE.CREATE_PIPE` function with a user specified pipe name.
- **Public Pipe:** Accessible by any user with `EXECUTE` permission on `DBMS_PIPE` package.
- **Private Pipe:** Accessible by sessions with the same user as the pipe creator.

> **Note:**
>
> Oracle recommends creating an explicit pipe before you send or receive messages with persistent messaging. Creating an explicit pipe with `DBMS_PIPE.CREATE_PIPE` ensures that the pipe is created with the access permissions you want, either public or private (by setting the `private` parameter).

The following shows the general workflow for `DBMS_PIPE` with persistent messaging:



Existing applications using DBMS_PIPE can continue to operate with minimal changes. You can configure existing applications that use DBMS_PIPE with a credential object and location URI using a logon trigger or using some other initialization routine. After setting the DBMS_PIPE credential and location URI, no other changes are needed to use persistent messaging. All subsequent use of the pipe stores the messages in Cloud Object Store instead of in database memory. This allows you to change the storage method for messages from in-memory to persistent Cloud Object Storage, with minimal changes.

Persistent Messaging Overview and Features

Features of DBMS_PIPE using persistent messaging:

- Messages can be sent and retrieved across multiple Autonomous Database instances in the same region or across regions.
- Persistent messages are guaranteed to either be written or read by exactly one process. This prevents message content inconsistency due to concurrent writes and reads. Using a persistent messaging pipe, DBMS_PIPE allows only one operation, sending a message or a receiving message to be active at a given time. However, if an operation is not possible due to an ongoing operation, the process retries periodically until the timeout value is reached.
- DBMS_PIPE uses DBMS_CLOUD to access Cloud Object Store. Messages can be stored in any of the supported Cloud Object Stores. See DBMS_CLOUD URI Formats for more information.
- DBMS_PIPE uses DBMS_CLOUD to access Cloud Object Store and all supported credential types are available:
  - DBMS_CLOUD.CREATE_CREDENTIAL: See CREATE_CREDENTIAL Procedure for more information.

DBMS_PIPE Privileges Authorization and Security

The DBMS_PIPE procedures run with invoker's rights. Private pipes are owned by the current user and a private pipe that is created by a user can only be used by the same user. This applies to both in-memory pipes and persistent messaging pipes where messages are stored to Cloud Object Store. Sending and receiving messages run in the invoker's schema.

Using private pipes, where messages are stored to Cloud Object Store, a credential object is required for authentication with the Cloud Object store identified by the location_uri parameter. The invoking user must have EXECUTE privilege on the credential object specified with the credential_name parameter that is used to access the Object Store.

To use a public pipe, the user, database session, must have execute privilege on DBMS_PIPE. For a public pipe using persistent messaging and storing messages to Cloud Object Store, the user, database session, must have execute privilege on DBMS_CLOUD and execute privilege on the credential object (or you can create a credential object that is allowed to access the location URI that contains the message.

DBMS_PIPE Limitation

The DBMS_PIPE package does not support sending messages between databases that use different character sets. For example, if you have one Autonomous Database instances that uses AL32UTF8 and another instance that uses WE8MSWIN1252, you cannot send messages with DBMS_PIPE between these two databases. In this case, you the system will raise the error ORA-12704 if you attempt to send messages with DBMS_PIPE between these two databases.

See Choose a Character Set for Autonomous Database for more information.

Create an Explicit Persistent Pipe and Send a Message

Describes the steps to create a persistent pipe with a specified pipe name (Explicit Pipe).

1. Store your object store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`. For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'my_persistent_pipe_cred',
    username => 'adb_user@example.com',
    password => 'password'
  );
END;
/
```

   This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. After you store the credentials you can then use the same credential name to access Cloud Object Store to send and receive messages with `DBMS_PIPE`.

   For detailed information about the parameters, see CREATE_CREDENTIAL Procedure. For Oracle Cloud Infrastructure Object Storage, it is required that the credential uses native Oracle Cloud Infrastructure authentication.

   Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

   > **Note:**
   >
   > Some tools like SQL*Plus and SQL Developer use the ampersand character (`&`) as a special character. If you have the ampersand character in your password use the `SET DEFINE OFF` command in those tools as shown in the example to disable the special character and get the credential created properly.

2. Create an explicit pipe to send and retrieve messages. For example, create a pipe named `ORDER_PIPE`.

```
DECLARE
  r_status INTEGER;
BEGIN
    r_status := DBMS_PIPE.CREATE_PIPE(pipename => 'ORDER_PIPE');
END;
/
```

   See CREATE_PIPE Function for more information.

3. Verify that the pipe is created.

```
SELECT ownerid, name, type FROM v$db_pipes
       WHERE name = 'ORDER_PIPE';
```

```
OWNERID NAME       TYPE
------- ---------- -------
     80 ORDER_PIPE PRIVATE
```

4. Use `DBMS_PIPE` procedures to set the default access credential and location URI to store persistent messages to Cloud Object Store.

```
BEGIN
    DBMS_PIPE.SET_CREDENTIAL_NAME('my_persistent_pipe_cred');
    DBMS_PIPE.SET_LOCATION_URI('https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname1/');
END;
/
```

These procedures set the default credential name and the default location URI for use with `DBMS_PIPE` procedures.

If you use Oracle Cloud Infrastructure Object Storage to store messages, you can use Oracle Cloud Infrastructure Native URIs or Swift URIs. However, the location URI and the credential must match in type as follows:

- If you use a native URI format to access Oracle Cloud Infrastructure Object Storage, you must use Native Oracle Cloud Infrastructure Signing Keys authentication in the credential object.
- If you use Swift URI format to access Oracle Cloud Infrastructure Object Storage, you must use an auth token authentication in the credential object.

See SET_CREDENTIAL_NAME Procedure and SET_LOCATION_URI Procedure for more information.

5. Pack and send a message on the pipe.

```
DECLARE
  l_result INTEGER;
  l_date   DATE;
BEGIN
    l_date := sysdate;
    DBMS_PIPE.PACK_MESSAGE(l_date);         -- date of order
    DBMS_PIPE.PACK_MESSAGE('C123');         -- order number
    DBMS_PIPE.PACK_MESSAGE(5);              -- number of items in order
    DBMS_PIPE.PACK_MESSAGE('Printers');     -- type of item in order


    l_result := DBMS_PIPE.SEND_MESSAGE(
                    pipename => 'ORDER_PIPE',
                    credential_name => DBMS_PIPE.GET_CREDENTIAL_NAME,
                    location_uri => DBMS_PIPE.GET_LOCATION_URI);

    IF l_result = 0 THEN
        DBMS_OUTPUT.put_line('DBMS_PIPE sent order successfully');
    END IF;

END;
/
```

See PACK_MESSAGE Procedures and SEND_MESSAGE Function for more information.

Retrieve a Persistent Message on Same Database

Describes the steps to retrieve a persistent message from an explicit pipe on the same Autonomous Database instance (the instance where the message was sent).

On an Autonomous Database instance you can receive messages sent to a pipe from a different session. The `DBMS_PIPE` procedures are invoker's rights procedures and run as the current invoked user.

Private pipes are owned by the current user that creates the pipe. Private pipes can only be accessed by the same user that created the pipe. This applies to pipes using in-memory messages and to pipes using persistent messaging with messages stored in Cloud Object Store.

Public pipes can be accessed by any database session having execute privilege on `DBMS_PIPE`. This applies to pipes using in-memory messages and to pipes using persistent messaging with messages stored in Cloud Object Store.

1. Verify that the pipe is created.

```
SELECT ownerid, name, type FROM v$db_pipes
       WHERE name = 'ORDER_PIPE';


OWNERID NAME       TYPE
------- ---------- -------
     80 ORDER_PIPE PRIVATE
```

   When you are on the same Autonomous Database instance and the pipe exists, you do not need to run `DBMS_PIPE.CREATE_PIPE` before you receive a message. This applies when the pipe was created on the same instance, as shown in Create an Explicit Persistent Pipe and Send a Message.

2. Receive a message from the pipe.

```
DECLARE
    message1  DATE;
    message2  VARCHAR2(100);
    message3  INTEGER;
    message4  VARCHAR2(100);
    l_result  INTEGER;

BEGIN

    DBMS_PIPE.SET_CREDENTIAL_NAME('my_persistent_pipe_cred');
    DBMS_PIPE.SET_LOCATION_URI('https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname1/');
    l_result := DBMS_PIPE.RECEIVE_MESSAGE (
                  pipename => 'ORDER_PIPE',
                  timeout  => DBMS_PIPE.MAXWAIT,
                  credential_name => DBMS_PIPE.GET_CREDENTIAL_NAME,
                  location_uri => DBMS_PIPE.GET_LOCATION_URI);

    IF l_result = 0 THEN
        DBMS_PIPE.unpack_message(message1);
        DBMS_PIPE.unpack_message(message2);
```

```
        DBMS_PIPE.unpack_message(message3);
        DBMS_PIPE.unpack_message(message4);

        DBMS_OUTPUT.put_line('Order Received Successfully On: ' ||
TO_CHAR(sysdate, 'dd-mm-yyyy hh24:mi:ss'));
        DBMS_OUTPUT.put_line('Date of Order: ' || message1);
        DBMS_OUTPUT.put_line('Order Number: ' || message2);
        DBMS_OUTPUT.put_line('Number of Items In Order: ' || message3);
        DBMS_OUTPUT.put_line('Item Type in Order: ' || message4);
    END IF;

END;
/
```

When you are on the same Autonomous Database instance, the credential already exists and you do not need to run `DBMS_CLOUD.CREATE_CREDENTIAL` to receive a message. This applies when the pipe was created on the same instance, as shown in Create an Explicit Persistent Pipe and Send a Message.

See SET_CREDENTIAL_NAME Procedure and SET_LOCATION_URI Procedure for more information.

See RECEIVE_MESSAGE Function for more information.

Retrieve a Persistent Message by Creating a Pipe on a Different Database

Describes the steps to retrieve a persistent message stored in Cloud Object Store with an explicit pipe on an Autonomous Database instance that is different than the instance that sent the message.

1. Store your object store credentials using the procedure `DBMS_CLOUD.CREATE_CREDENTIAL`. For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'my_persistent_pipe_cred',
    username => 'adb_user@example.com',
    password => 'password'
  );
END;
/
```

This operation stores the credentials in the database in an encrypted format. You can use any name for the credential name. Note that this step is required only once unless your object store credentials change. Once you store the credentials you can then use the same credential name to access the Cloud Object Store to send and receive messages with `DBMS_PIPE`.

For detailed information about the parameters, see CREATE_CREDENTIAL Procedure.

Creating a credential to access Oracle Cloud Infrastructure Object Store is not required if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

> **Note:**
>
> Some tools like SQL*Plus and SQL Developer use the ampersand character (&) as a special character. If you have the ampersand character in your password use the `SET DEFINE OFF` command in those tools as shown in the example to disable the special character and get the credential created properly.

2. Create an explicit pipe with the same name as the pipe that sent the message. For example, create a pipe named `ORDER_PIPE`.

```
DECLARE
  r_status INTEGER;
BEGIN
    r_status := DBMS_PIPE.CREATE_PIPE(pipename => 'ORDER_PIPE');
END;
/
```

See CREATE_PIPE Function.

3. Verify that the pipe is created.

```
SELECT ownerid, name, type FROM v$db_pipes
       WHERE name = 'ORDER_PIPE';


OWNERID NAME       TYPE
------- ---------- -------
     80 ORDER_PIPE PRIVATE
```

4. Use `DBMS_PIPE` procedures to set the default access credential and location URI for Object Store so that `DBMS_PIPE` can access the persistent message.

```
BEGIN
    DBMS_PIPE.SET_CREDENTIAL_NAME('my_persistent_pipe_cred');
    DBMS_PIPE.SET_LOCATION_URI('https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname1/');
END;
/
```

These procedures set the default credential name and the default location URI for use with `DBMS_PIPE` procedures.

If you use Oracle Cloud Infrastructure Object Storage to store messages, you can use Oracle Cloud Infrastructure Native URIs or Swift URIs. However, the location URI and the credential must match in type as follows:

- If you use a native URI format to access Oracle Cloud Infrastructure Object Storage, you must use Native Oracle Cloud Infrastructure Signing Keys authentication in the credential object.
- If you use Swift URI format to access Oracle Cloud Infrastructure Object Storage, you must use an auth token authentication in the credential object.

See SET_CREDENTIAL_NAME Procedure and SET_LOCATION_URI Procedure for more information.

5. Receive a message from the persistent pipe.

```
DECLARE
    message1  DATE;
    message2  VARCHAR2(100);
    message3  INTEGER;
    message4  VARCHAR2(100);
    l_result  INTEGER;

BEGIN

    DBMS_PIPE.SET_CREDENTIAL_NAME('my_persistent_pipe_cred');
    DBMS_PIPE.SET_LOCATION_URI('https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname1/');
    l_result := DBMS_PIPE.RECEIVE_MESSAGE (
                  pipename => 'ORDER_PIPE',
                  timeout  => DBMS_PIPE.MAXWAIT,
                  credential_name => DBMS_PIPE.GET_CREDENTIAL_NAME,
                  location_uri => DBMS_PIPE.GET_LOCATION_URI);

    IF l_result = 0 THEN
        DBMS_PIPE.unpack_message(message1);
        DBMS_PIPE.unpack_message(message2);
        DBMS_PIPE.unpack_message(message3);
        DBMS_PIPE.unpack_message(message4);

        DBMS_OUTPUT.put_line('Order Received Successfully On: ' ||
TO_CHAR(sysdate, 'dd-mm-yyyy hh24:mi:ss'));
        DBMS_OUTPUT.put_line('Date of Order: ' || message1);
        DBMS_OUTPUT.put_line('Order Number: ' || message2);
        DBMS_OUTPUT.put_line('Number of Items In Order: ' || message3);
        DBMS_OUTPUT.put_line('Item Type in Order: ' || message4);
    END IF;

END;
/
```

See RECEIVE_MESSAGE Function for more information.

Remove a Persistent Pipe

Describes the steps to remove a persistent pipe.

Persistent pipes send and receive messages by storing messages in Cloud Object Store. Use `DBMS_PIPE.REMOVE_PIPE` to remove a persistent pipe on an Autonomous Database instance.

1. Call the `DBMS_PIPE.REMOVE_PIPE` function to remove a pipe.

```
DECLARE
   l_result  INTEGER;
BEGIN
     l_result := DBMS_PIPE.REMOVE_PIPE('ORDER_PIPE');
END;
/
```

The `REMOVE_PIPE` function removes the pipe from the Autonomous Database instance where it runs, however `REMOVE_PIPE` does not affect other Autonomous Database instances with a pipe with the same name that uses the same location URI.

2. On the Autonomous Database instance were you run `DBMS_PIPE.REMOVE_PIPE`, verify that the pipe is removed.

```
SELECT ownerid, name, type FROM v$db_pipes
       WHERE name = 'ORDER_PIPE';


No rows selected
```

Full-Text Search Queries

The PL/SQL package `DBMS_CLOUD` enables you to build a text index on the object store files, which allows you to search the text and use wildcards with your search. You can use `JSON_TEXTCONTAINS` to do full-text JSON search on files stored in Object Storage.

- Use Full-Text Search on Files in Object Storage
  The PL/SQL package DBMS_CLOUD enables you to build a text index on the object store files, which allows you to search the text and use wildcards with your search.
- Use JSON Search on Documents in Object Storage
  You can build a JSON search index on documents residing in files stored in Object Storage, which allows you to search fields in your JSON documents using `JSON_TEXTCONTAINS`, for example, searching with wildcards.

Use Full-Text Search on Files in Object Storage

The PL/SQL package DBMS_CLOUD enables you to build a text index on the object store files, which allows you to search the text and use wildcards with your search.

- About Full-Text Search on Files in Object Storage
- Create a Text Index on Object Storage Files
  Use `DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX` to create a text index on files in object storage.
- Drop an Index on the Cloud Storage Files
  Use the `DBMS_CLOUD.DROP_EXTERNAL_TEXT_INDEX` procedure to drop a text index on object storage files.
- Text Index Reference Table
  A local table is created within your database with a standard suffix *`INDEX_NAME`*`$TXTIDX`. This table is created internally when you run `DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX`.
- Monitor Text Index Creation
  When you run `DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX` the text index creation operation is logged in the `ALL_SCHEDULER_JOB_RUN_DETAILS` view.

About Full-Text Search on Files in Object Storage

You can create a text index on files in object storage. A text index allows you to perform a word-based search across very large data sets in object store.

DBMS_CLOUD provides fast and efficient ways to manage data in object store. The DBMS_CLOUD APIs let you create, copy, download, delete, and traverse files present in object store. When you define external tables you can run SQL queries on data stored in your object store (or with hybrid partitioned external tables, across data in your database and in object store). When you use DBMS_CLOUD to define a text index, this allows you to search your data for text and use wildcards.

Autonomous Database support for word-based search works for commonly used data formats, for example CSV or JSON and with formatted documents (binary), for example PDF and DOC (MS Word) formats. You can configure a refresh rate that indicates the frequency in minutes at which the index is refreshed for any new uploads or deletes.

A local table with the standard suffix *INDEX_NAME*$TXTIDX is created when you create an index on the object storage, and you can utilize the table to perform a search using the CONTAINS keyword.

See Indexing with Oracle Text for more information.

Create a Text Index on Object Storage Files

Use DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX to create a text index on files in object storage.

Formatted documents (binary) are supported when you specify the binary_files format option with DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX.

You can include a stop word list when you specify the stop_words format option with DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX.

See Indexing with Oracle Text for more information on Oracle Text stop words and working with binary files.

1. Create a credential object to access the source location.

   See CREATE_CREDENTIAL Procedure for more information.

2. Run the DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX procedure to create a text index on the object storage files.

```
BEGIN
DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX (
        credential_name => 'DEFAULT_CREDENTIAL',
        location_uri    => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/ts_data/'
        index_name      => 'EMP',
        format          => JSON_OBJECT ('refresh_rate' value 10)
);
END;
/
```

   This example creates a text index EMP on the object storage files located at the URI specified in the location_uri parameter. The refresh_rate option in the format parameter specifies that the index EMP is refreshed at an interval of 10 minutes.

   This creates a local table *INDEX_NAME*$TXTIDX. You can utilize the table *INDEX_NAME*$TXTIDX to perform a search using CONTAINS.

For example:

```
SELECT object_name FROM EMP$TXTIDX
     WHERE CONTAINS(object_name,'king') > 0;
```

This query returns the object or file names that contain the string `king`.

See Text Index Reference Table for more information.

You can query an external table using the `EXTERNAL MODIFY` clause to retrieve the actual records.

```
SELECT * FROM EMPEXTTAB EXTERNAL MODIFY ((location_url object_name));
```

> **Note:**
>
> The external table `EMPEXTTAB` is a sample external table that is created on the same `location_url`.

See Query External Data with Autonomous Database for more information.

See CREATE_EXTERNAL_TEXT_INDEX Procedure for more information.

See Configure Policies and Roles to Access Resources for more information.

Drop an Index on the Cloud Storage Files

Use the `DBMS_CLOUD.DROP_EXTERNAL_TEXT_INDEX` procedure to drop a text index on object storage files.

Run the `DBMS_CLOUD.DROP_EXTERNAL_TEXT_INDEX` procedure to drop a text index on files in object storage.

```
BEGIN
DBMS_CLOUD.DROP_EXTERNAL_TEXT_INDEX (
        index_name => 'EMP',
);
END;
/
```

This example drops the `EMP` text index.

See DROP_EXTERNAL_TEXT_INDEX Procedure for more information.

Text Index Reference Table

A local table is created within your database with a standard suffix *`INDEX_NAME`*`$TXTIDX`. This table is created internally when you run `DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX`.

You can query the *`INDEX_NAME`*`$TXTIDX` table to search for a string using the `CONTAINS` keyword. For example, when you call `DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX` procedure the *`INDEX_NAME`* value as `EMP`, this creates the `EMP$TXTIDX` the text reference table.

The text reference table has the following columns:

- `object_name`: is the file name on the object storage that contains the searched text string.
- `object_path`: is the object storage bucket or folder URI that contains the object storage file.
- `mtime`: is the last modified timestamp of the object storage file. This is the time when the file was last accessed by `DBMS_CLOUD`.

For example:

```
SELECT object_path, object_name FROM EMP$TXTIDX WHERE CONTAINS(OBJECT_NAME, 'king') > 0;

OBJECT_PATH
    OBJECT_NAME
--------------------------------------------------------------------------------------------
   -----------------------------------
https://objectstorage.us-phoenix-1.oraclecloud.com/n/example1/b/adbs_data_share/o/
ts_data/       data_2_20221026T195313585601Z.json
```

This query returns the file names and location URI on the object storage which contains the text string `king`, in either upper or lowercase.

```
SELECT object_name, mtime FROM EMP$TXTIDX;

OBJECT_NAME                    MTIME
------------------------------ -----------------------------------
data_1_20220531T165402Z.json    31-MAY-22 04.54.02.979000 PM +00:00
data_1_20220531T165427Z.json    31-MAY-22 04.54.27.997000 PM +00:00
```

This query returns file name and last modified timestamp of the object files on which the index `EMP` is created.

Monitor Text Index Creation

When you run `DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX` the text index creation operation is logged in the `ALL_SCHEDULER_JOB_RUN_DETAILS` view.

You can query the `ALL_SCHEDULER_JOB_RUN_DETAILS` view to obtain the status and any error reported by the index creation job.

The name of the `DBMS_SCHEDULER` job is derived from the `INDEX_NAME` parameter specified when you call `DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX`.

To query the `ALL_SCHEDULER_JOB_RUN_DETAILS` view, you must be logged in as the `ADMIN` user or have `READ` privilege on the `ALL_SCHEDULER_JOB_RUN_DETAILS` view.

For example, the following `SELECT` statement with a `WHERE` clause on `job_name` shows the run details for the job:

```
SELECT status, additional_info
   FROM all_scheduler_job_run_details WHERE LOWER(job_name) = LOWER('index_name$JOB');
```

You can also query for the existence of an index creation scheduler job.

For example:

```
SELECT status
    FROM all_scheduler_jobs where LOWER(job_name) = LOWER('index_name$JOB');
```

See CREATE_EXTERNAL_TEXT_INDEX Procedure for more information.

Use JSON Search on Documents in Object Storage

You can build a JSON search index on documents residing in files stored in Object Storage, which allows you to search fields in your JSON documents using `JSON_TEXTCONTAINS`, for example, searching with wildcards.

- Use JSON Search on Documents in Object Storage
- Create a JSON Search Index on JSON Documents in Files
  Use `DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX` to create a JSON search index on the JSON files stored in Object Storage.
- Drop a JSON Search Index
  Use the `DBMS_CLOUD.DROP_EXTERNAL_TEXT_INDEX` procedure to drop a JSON search index on the JSON files stored in Object Storage.
- Monitor JSON Search Index Creation
  When you run `DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX` the JSON search index creation operation is logged in the `ALL_SCHEDULER_JOB_RUN_DETAILS` view.
- JSON Search Index Reference Table
  A local table is created within your database with a standard suffix *`INDEX_NAME`*`$TXTIDX`. This table is created internally when you run `DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX`.
- JSON Search Index Reference View
  A view with the name *`INDEX_NAME`* is created when you run `DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX` to create a JSON index.
- Error Log Table

Use JSON Search on Documents in Object Storage

You can create a JSON search index on documents in files stored in Object Storage. A JSON search index is specifically designed for JSON documents and enables you to do regular as well as full text search on your documents or fragments of your documents using `JSON_TEXTCONTAINS`.

See JSON_TEXTCONTAINS Condition for more information.

You use `DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX` to create a JSON search index on files in Object Storage. You can configure a refresh rate that indicates the frequency in minutes at which the index is refreshed for any new uploads or deletes. See Create a JSON Search Index on JSON Documents in Files for more information.

The following objects are created when you create a JSON search index:

- A JSON search index *`index_name`* on the Object Storage JSON files.
- A local table *`INDEX_NAME`*`$TXTIDX`. This table is the mapping table between your documents in the files in Object Storage and the JSON search index that is created in the database. See JSON Search Index Reference Table for more information.

- A view with the standard name *INDEX_NAME*. You can utilize the view to perform a search using JSON_TEXTCONTAINS. This view is created on top of the *INDEX_NAME*$TXTIDX table. See JSON Search Index Reference View for more information.
- An error log table with the standard suffix *index_name*$txtidx_err. See Error Log Table for more information.

The JSON search index creation operation is logged in the ALL_SCHEDULER_JOB_RUN_DETAILS view. See Monitor Text Index Creation for more information.

Autonomous Database supports JSON search index creation for a variety of text files containing JSON documents. For example, you can index JSON files in both compressed and uncompressed format, or choose a different JSON document delimiter. See DBMS_CLOUD Package for more information about the file handling options.

See Indexing with Oracle Text for more information.

Create a JSON Search Index on JSON Documents in Files

Use DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX to create a JSON search index on the JSON files stored in Object Storage.

See Indexing with Oracle Text for more information on Oracle Text stop words.

1. Create a credential object to access the source location.

   See CREATE_CREDENTIAL Procedure for more information.

2. Run the DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX procedure to create a JSON search index on the object storage files.

```
BEGIN
 DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX (
        credential_name => 'OBJ_STORE_CRED',
        location_uri    => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/vcnlogs/',
        index_name      => 'VNCLOGS',
        format          =>  JSON_OBJECT ('json_index' value 'true',
'json_index_doc_len' value 12000, 'file_compression' value 'gzip')
);
END;
/
```

   This example creates the following:

   - A JSON search index VNCLOGS on the Object Storage files located at the URI specified in the location_uri parameter.
   - A local table VNCLOGS$TXTIDX. See JSON Search Index Reference Table for more information.
   - A view with the name VNCLOGS. This view is created on top of the VNCLOGS$TXTIDX table. You can utilize the view to perform a search using JSON_TEXTCONTAINS. See JSON Search Index Reference View for more information.
   - An error log table with the name VNC$TXTIDX_ERR. The table contains a list of invalid JSON documents in your files. See Error Log Table for more information.

After creating the JSON search index, you can query the VNCLOGS view. For example:

```
SELECT JSON_QUERY (data, '$' returning CLOB pretty) AS RECORD
 FROM vnclogs WHERE JSON_TEXTCONTAINS(file_line_json, '$.tuples',
'{XXX.XXX.0.1.10.10.10.10.null.XXX_345}');
```

This query returns the log details for the specified IP addresses in the JSON_TEXTCONTAINS condition. Note that the JSON Search index criteria is specified on the file_line_json column, the indexed column in your JSON search index table in the database, while the JSON documents are streamed directly from Object Storage, represented through the data column. The data column contains the content of the Object Storage file.

See CREATE_EXTERNAL_TEXT_INDEX Procedure for more information.

Drop a JSON Search Index

Use the DBMS_CLOUD.DROP_EXTERNAL_TEXT_INDEX procedure to drop a JSON search index on the JSON files stored in Object Storage.

For example:

```
BEGIN
 DBMS_CLOUD.DROP_EXTERNAL_TEXT_INDEX (
  index_name => 'VCNLOGS',
);
END;
/
```

This example drops the VCNLOGS JSON search index.

See DROP_EXTERNAL_TEXT_INDEX Procedure for more information.

Monitor JSON Search Index Creation

When you run DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX the JSON search index creation operation is logged in the ALL_SCHEDULER_JOB_RUN_DETAILS view.

You can query the ALL_SCHEDULER_JOB_RUN_DETAILS view to obtain the status and any error reported by the index creation job.

The name of the DBMS_SCHEDULER job is derived from the INDEX_NAME parameter specified when you call DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX.

To query the ALL_SCHEDULER_JOB_RUN_DETAILS view, you must be logged in as the ADMIN user or have READ privilege on the ALL_SCHEDULER_JOB_RUN_DETAILS view.

For example, the following SELECT statement with a WHERE clause on job_name shows the run details for the job:

```
SELECT status, additional_info
   FROM all_scheduler_job_run_details WHERE LOWER(job_name) = LOWER('index_name$JOB');
```

You can also query for the existence of an index creation scheduler job.

For example:

```
SELECT status
    FROM all_scheduler_jobs where LOWER(job_name) = LOWER('index_name$JOB');
```

See CREATE_EXTERNAL_TEXT_INDEX Procedure for more information.

JSON Search Index Reference Table

A local table is created within your database with a standard suffix *INDEX_NAME*$TXTIDX. This table is created internally when you run DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX.

You can query the *INDEX_NAME*$TXTIDX table to search for a string using the JSON_TEXTCONTAINS keyword. For example, when you call DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX procedure the *INDEX_NAME* value as VNCLOGS, this creates the VNCLOGS$TXTIDX JSON search index reference table.

The JSON search index reference table has the following columns:

- object_name: is the file name on the Object Storage that contains the searched text string.
- object_path: is the Object Storage bucket or folder URI that contains the object storage file.
- length: is the length of the JSON document.
- offset: is the byte offset of the beginning of the file.
- mtime: is the last modified timestamp of the object storage file. This is the time when the file was last accessed by DBMS_CLOUD.

Run the following query to retrieve the object name, object path, offset and length from the JSON search index reference table:

```
SELECT * FROM (SELECT object_name, object_path, length, offset
 FROM vnclogs$txtidx
 WHERE JSON_TEXTCONTAINS (file_line_json, '$.tuples',
'{XXX.XXX.0.1.10.10.10.10.null.XXX_345}')
 ORDER BY mtime DESC)
 WHERE rownum < 3;
```

This query provides internal information about where in your files in Object Storage the relevant documents are located. The query does not provide the actual documents.

To retrieve the documents directly, use the following query:

```
SELECT * FROM (SELECT object_name, object_path, data
 FROM vnclogs
 WHERE JSON_TEXTCONTAINS (file_line_json, '$.tuples',
'{XXX.XXX.0.1.10.10.10.10.null.XXX_345}')
 ORDER BY mtime DESC)
 WHERE rownum < 3;
```

JSON Search Index Reference View

A view with the name *INDEX_NAME* is created when you run DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX to create a JSON index.

The view itself contains no data. The data is fetched in the view when you query the *INDEX_NAME* view using SQL condition JSON_TEXTCONTAINS.

Column	Description
OBJECT_NAME	The file name on the Object Storage that contains the searched text string.
OBJECT_PATH	The Object Storage bucket or folder URI that contains the object storage file.
LENGTH	The length of the JSON document.
OFFSET	The byte offset of the beginning of the file.
DATA	The content of the Object Storage file.
FILE_LINE_JSON	The column on which the JSON_TEXTCONTAINS is performed.
MTIME	The last modified timestamp of the object storage file. This is the time when the file was last accessed by DBMS_CLOUD.

Error Log Table

An error log table *index_name*$txtidx_err is created when you create a JSON search index on the JSON files stored in Object Storage.

The table contains a list of the invalid JSON documents in your files. A JSON document is marked as invalid if the length of the source file that is indexed exceeds the default value or the specified value. The default value is 32767 bytes, and the maximum value you can specify using the json_index_doc_len format option is 200000 bytes. See Create a JSON Search Index on JSON Documents in Files for more information.

You can query the *index_name*$txtidx_err to get the list of invalid JSON files reported during the index creation.

Column	Description
OBJECT_NAME	The file name on the Object Storage that contains the searched text string.
OBJECT_PATH	The Object Storage bucket or folder URI that contains the object storage file.
OFFSET	Byte offset of beginning of the line.
LENGTH	Length of the JSON document.
LINE_NUMBER	Line number in the file at object store.

Use Oracle Workspace Manager on Autonomous Database

Workspace Manager provides an infrastructure that enables applications to create workspaces and group different versions of table row values in different workspaces.

- About Using Oracle Workspace Manager on Autonomous Database
  Use Oracle Workspace Manager to version-enable one or more user tables in the database. When a table is version-enabled, all rows in the table can support multiple versions of the data.
- Enable Oracle Workspace Manager on Autonomous Database
  Oracle Workspace Manager must be enabled to be used on Autonomous Database. You can migrate existing data that has enabled Oracle Workspace Manager.
- Disable Oracle Workspace Manager on Autonomous Database
  Use the following steps to disable Oracle Workspace Manager on Autonomous Database.

About Using Oracle Workspace Manager on Autonomous Database

Use Oracle Workspace Manager to version-enable one or more user tables in the database. When a table is version-enabled, all rows in the table can support multiple versions of the data.

Applications that can benefit from Workspace Manager typically do one or more of the following operations:

- Manage a collection of updates and insertions as a unit before incorporating them into production data
- Support a collaborative development effort
- Use a common data set to create multiple scenarios for what-if analyses or multiple editions of data for publication
- Keep a history of changes to data

Introduction to Workspace Manager provides relevant usage and reference information.

Workspace manager is enabled using the procedure `DBMS_CLOUD_ADMIN.ENABLE_FEATURE`.

Enable Oracle Workspace Manager on Autonomous Database

Oracle Workspace Manager must be enabled to be used on Autonomous Database. You can migrate existing data that has enabled Oracle Workspace Manager.

As the `ADMIN` user, run `DBMS_CLOUD_ADMIN.ENABLE_FEATURE` to enable Oracle Workspace Manager.

1. Enable Oracle Workspace Manager.

```
BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_FEATURE(
        feature_name => 'OWM');
END;
/
```

2. Restart the Autonomous Database instance.

   See Restart Autonomous Database for more information.

3. Query `dba_cloud_config` to verify that Oracle Workspace Manager is enabled.

```
SELECT param_name, param_value FROM dba_cloud_config WHERE
       UPPER(param_name) = 'OWM';

PARAM_NAME PARAM_VALUE
```

```
---------- -----------
owm        enabled
```

4. (Optional) If you have existing data using Oracle Workspace Manager on another database that you want to migrate to Autonomous Database, Oracle Workspace Manager provides import and export procedures that enable you to migrate the data. See Export_Schemas and Import_Schemas for additional information.
5. Continue with the imported schemas or the new schemas.

See ENABLE_FEATURE Procedure for more information.

Disable Oracle Workspace Manager on Autonomous Database

Use the following steps to disable Oracle Workspace Manager on Autonomous Database.

As the `ADMIN` user, run `DBMS_CLOUD_ADMIN.DISABLE_FEATURE` to disable Oracle Workspace Manager.

1. Disable Oracle Workspace Manager.

```
BEGIN
   DBMS_CLOUD_ADMIN.DISABLE_FEATURE(
        feature_name => 'OWM');
END;
/
```

2. Query `dba_cloud_config` to verify that Oracle Workspace Manager is disabled.

```
SELECT param_name, param_value FROM dba_cloud_config
     WHERE UPPER(param_name) = 'OWM';

0 rows selected.
```

See DISABLE_FEATURE Procedure for more information.

Use and Manage Elastic Pools on Autonomous Database

Use an elastic pool to consolidate your Autonomous Database instances, in terms of their allocation of compute resources, and to provide up to 87% cost savings.

> **Note:**
>
> Elastic pools are only available for Autonomous Database instances that use the ECPU compute model.

- About Elastic Pools
  Elastic pools help you improve operating efficiency and reduce costs by bringing all of your databases to the Cloud. This also supports consolidating resources and simplifying administration and operations by using Autonomous Database.
- About Elastic Pool Billing

- About Elastic Pool Billing with Autonomous Data Guard Enabled
  An Autonomous Data Guard primary database can use a local or a cross-region standby that is part of an elastic pool, either as a leader or a member.
- Create, Join or Manage an Elastic Pool
  Shows the steps to create, join, or manage an elastic pool.
- List Elastic Pool Members
  Shows the steps for the pool leader to list elastic pool members.
- Offload Queries to Member Refreshable Clones
  When you have heavy read workloads, where a relatively small amount of data is scanned by multiple queries, you can offload queries (reads) either from an elastic pool leader or from an elastic pool member to Refreshable Clones.
- Remove Pool Members from an Elastic Pool
  As an elastic pool member you can remove yourself from an elastic pool. As an elastic pool leader you can remove pool members from an elastic pool.
- Terminate an Elastic Pool
  Shows the steps to terminate an elastic pool. Only the pool leader can terminate an elastic pool.
- Elastic Pool Notes
  Provides information about elastic pool usage.

About Elastic Pools

Elastic pools help you improve operating efficiency and reduce costs by bringing all of your databases to the Cloud. This also supports consolidating resources and simplifying administration and operations by using Autonomous Database.

Elastic pools are only available for Autonomous Database instances that use the ECPU compute model.

When you need a large number of databases that can scale up and down elastically without downtime, you can benefit by creating and using elastic pools. Elastic pools have the following advantages:

- Enable operating with a fixed budget for a group of databases, while delivering performance elasticity for each individual database.
- Allow for easy migration from on-prem Oracle environments that include oversubscription, to provide a cost effective way to move to Autonomous Database.
- Support SaaS vendors with a large number of individual customer databases.
- Provide resources for using a microservices architecture, where the ability to supply of large number of databases is required.
- The pool members in an elastic pool are not billed individually (the pool leader is billed based on the pool shape). You can allocate additional ECPUs per instance for pool members, without worrying about the cost associated with the ECPU usage for the individual members. Autonomous Database IO capacity and memory allocation is directly correlated with the **ECPU count**, so by selecting a greater number of ECPUs for an instance, this allows you to run with greater IO capacity and more memory without having to pay for the additional resources. This means, using a larger number of ECPUs per instance allows you to use more IO capacity and more memory per instance, where the cost is based on the pool shape and is not based on an individual instance's **ECPU count**.

When you create an elastic pool you select a pool size from a predefined set of pool sizes. Pool size determines how much you pay for compute as well as how many ECPUs you can provision in a given pool.

There are several terms to use when you work with elastic pools:

- **Pool Leader**: Is the Autonomous Database instance that creates an elastic pool.
- **Pool Member**: Is an Autonomous Database instance that is added to an elastic pool.
- **Pool Size**: Is a value that you set when you create an elastic pool. The pool size must be one of the available elastic pool shapes.
- **Pool Capacity**: The pool capacity is the maximum number of ECPUs that an elastic pool can use, and is four times (x4) the pool size.
- **Pool Shape**: A pool shape is one of the valid pool sizes that you select when you create an elastic pool. The pool shape must be one of: 128, 256, 512, 1024, 2048, or 4096 ECPUs.

> **Note:**
>
> By default each instance in an elastic pool is automatically assigned a maintenance window. By selecting a pool shape that is 1024 ECPUs or greater, you have the option of assigning a custom 2-hour maintenance window during which the leader and all elastic pool members are patched together. To select a custom maintenance window for your elastic pool, file a Service Request at Oracle Cloud Support.

Requirements to Create an Elastic Pool

The following are the requirements for an Autonomous Database instance to create an elastic pool and become a pool leader:

- The instance must use the ECPU compute model.
- Auto scaling must be disabled.
- The instance must not be a member of an existing elastic pool.
- The maximum allowed individual **ECPU count** for an Autonomous Database instance that creates an elastic pool is 4 times the pool size specified when you create the pool.
- The instance that creates an elastic pool is subject to tenancy limits. To create an elastic pool you must have a sufficient number of ECPUs available, below the tenancy limit, to accommodate the size of the elastic pool.

> **Note:**
>
> An elastic pool leader can be of any workload type. Elastic pool billing for compute usage is always billed to the pool leader at the **Transaction Processing** compute usage rate. This includes when the elastic pool leader has the **Transaction Processing**, **Data Warehouse**, **JSON Database**, or **APEX** workload type. See About Elastic Pool Billing and ECPU Compute Model Billing Information for more information.

Requirements to Join an Elastic Pool

The following are the requirements for an Autonomous Database instance to join an elastic pool:

- The instance must use the ECPU compute model.
- An elastic pool can contain Autonomous Database instances with **Transaction Processing**, **Data Warehouse**, **JSON Database**, or **APEX** workload types.
- Auto scaling must be disabled.
- The instance must not be a member of an elastic pool.
- The maximum allowed individual **ECPU count** for an Autonomous Database instance is the available pool capacity. When an instance has an ECPU count greater than the available pool capacity, it is not allowed to join that elastic pool.

Pool Leader and Member Minimum ECPU Allocation

When an Autonomous Database instance is part of an elastic pool, the minimum allowed individual ECPU allocation is 1 ECPU.

When an Autonomous Database instance is part of an elastic pool, increments of 1 ECPU are allowed for individual Autonomous Database instance ECPU allocation.

Pool Capacity for an Elastic Pool

An elastic pool has a pool capacity of 4 times the pool size. For example, a pool with pool size of 128 ECPUs can hold up to 512 ECPUs for its leader and the members.

> **Note:**
>
> In these examples Autonomous Data Guard is not enabled. See About Elastic Pool Billing with Autonomous Data Guard Enabled for information on using elastic pools with Autonomous Data Guard.

The following are examples of Autonomous Database instances that could be in an elastic pool with a pool size of 128 and a pool capacity of 512:

- Each of these are valid for pool members in an elastic pool with a pool size of 128:
  - 1 instance with 512 ECPUs, for a total of 512 ECPUs
  - 128 instances with 4 ECPUs, for a total of 512 ECPUs
  - 256 instances with 2 ECPUs, for a total of 512 ECPUs
  - 50 instances with 10 ECPUs and 3 instances with 4 ECPUs, for a total of 512 ECPUs
- Similarly, each of the following are valid for pool members in an elastic pool with a pool size of 128:
  - 1 instance with 128 ECPUs, 2 instances with 64 ECPUs, 32 instances with 4 ECPUs, and 64 instances with 2 ECPUs, for a total of 512 ECPUs
  - 256 instances with 1 ECPU, 64 instances with 2 ECPUs, for a total of 384 ECPUs, which is less than the pool capacity of 512 ECPUs.
  - 100 instances with 4 ECPUs and 50 instances with 2 ECPUs, which is less than the pool capacity of 512 ECPUs.

These are examples, you can add pool members to a pool to match the number of instances and the number ECPUs per instance to meet your needs, based on the pool size you select.

About Elastic Pool Leader Operations

The Autonomous Database instance that creates an elastic pool is the pool leader. Autonomous Database instances that are added to an existing pool are pool members. Depending on your role, either leader or member, you can perform operations on an elastic pool.
The following operations are valid only for the pool leader:

Operation	Description
Create an elastic pool	The Autonomous Database instance that creates an elastic pool is the pool leader. See Create an Elastic Pool for more information.
Remove an elastic pool member	An elastic pool leader can remove a member from the elastic pool. See As Pool Leader Remove Members from an Elastic Pool for more information.
Terminate an elastic pool	When an elastic pool has no pool members, the pool leader can terminate the elastic pool. See Terminate an Elastic Pool for more information.
Modify elastic pool size	An elastic pool leader can modify the pool size. See Change the Elastic Pool Shape for more information.
List pool members	A pool leader can list pool members. See List Elastic Pool Members for more information.

About Elastic Pool Member Operations

The Autonomous Database instance that creates an elastic pool is the pool leader. Autonomous Database instances that are added to an existing pool are pool members. Depending on your role, either leader or member, you can perform operations on an elastic pool.
The following operations are valid for a pool member or for the pool leader:

Operation	Description
Add instance to elastic pool	An Autonomous Database instance can be added as a pool member as long as the instance is one of the supported workload types, the instance uses the ECPU compute model, and the instance is not a pool member of a different pool. The supported workload types are: **Transaction Processing**, **Data Warehouse**, **JSON Database**, or **APEX**. See Join an Existing Elastic Pool for more information.
Remove an elastic pool member	An elastic pool member can remove themselves from the elastic pool. See Remove Pool Members from an Elastic Pool for more information.

About Elastic Pool Billing

Elastic pool usage is billed to the pool leader and billing is based on the elastic pool size and the actual hourly ECPU usage of the pool leader and the members. Elastic pool usage can exceed the pool size (pool capacity can be up to four times greater than the pool size).
The billing for an elastic pool consists of only compute resources, that is ECPU usage, and all compute usage is charged to the Autonomous Database instance that is the pool leader. Any

billing for storage usage is charged separately to individual Autonomous Database instances, independent of whether the instance is in an elastic pool.

An elastic pool allows you to consolidate your Autonomous Database instances in terms of their compute resource billing. You can think of an elastic pool like a mobile phone service “family plan”, except this applies to your Autonomous Database instances. Instead of paying individually for each database, the databases are grouped into a pool in which one instance, the leader, is charged for the compute usage associated with the entire pool.

Using an elastic pool you are able to provision up to four times the number of ECPUs, over your selected pool size, and you can provision database instances that are in the elastic pool with as little as 1 ECPU per database instance. Outside of an elastic pool the minimum number of ECPUs per database instance is 2 ECPUs. For example, with a pool size of 128 you can provision 512 Autonomous Database instances (when each instance has 1 ECPU). In this example you are billed for the pool size compute resources, based on the pool size of 128 ECPUs, while you have access to 512 Autonomous Database instances. In contrast, when you individually provision 512 Autonomous Database instances without using an elastic pool you are required to allocate a minimum of 2 ECPUs for each Autonomous Database instance, and in this example you would pay for 1024 ECPUs. Using an elastic pool provides **up to 87% compute cost savings**.

After you create an elastic pool, the total ECPU usage for a given hour is charged to the Autonomous Database instance that is the pool leader. With the exception of the pool leader, individual Autonomous Database instances that are pool members are not charged for ECPU usage while they are members of an elastic pool.

Elastic pool billing is as follows:

- If the total aggregated peak ECPU utilization is equal to or below the pool size for a given hour, you are charged for the pool size number of ECPUs (one times the pool size).

  After an elastic pool is created ECPU billing continues at a minimum of one times the pool size rate, even when databases that are part of the pool are stopped. This applies to pool member databases and to the pool leader.

  In other words, if the aggregated peak ECPU utilization of the pool is less than or equal to the pool size for a given hour, you are charged for the pool size number of ECPUs (one times the pool size). This represents **up to 87% compute cost savings** over the case in which these databases are billed separately without using elastic pools.

- If the aggregated peak ECPU utilization of the pool leader and the members exceeds the pool size at any point in time in a given billing hour:

  – **Aggregated peak ECPU utilization of the pool is equal to or less than two times of the pool size number of ECPUs**: For usage that is greater than one times the pool size number of ECPUs and up to and including two times the number of ECPUs in a given billing hour: Hourly billing is two times the pool size number of ECPUs.

    In other words, if the aggregated peak ECPU utilization of the pool exceeds the pool size, but is less than or equal to two times the pool size for a given hour, you are charged for twice the pool size number of ECPUs (two times the pool size). This represents **up to 75% compute cost savings** over the case in which these databases are billed separately without using elastic pools.

  – **Aggregated peak ECPU utilization of the pool is equal to or less than four times the pool size number of ECPUs**: For usage that is greater than two times the pool size number of ECPUs and up and including to four times the pool size number of ECPUs in a given billing hour: Hourly billing is four times the pool size number of ECPUs.

In other words, if the aggregated peak ECPU utilization of the pool exceeds twice the pool size for a given hour, you are charged for four times the pool size number of ECPUs (four times the pool size). This represents **up to 50% compute cost savings** over the case in which these databases are billed separately without using elastic pools.

For example, consider an elastic pool with a pool size of 128 ECPUs and a pool capacity of 512 ECPUs:

- **Case-1**: The aggregated peak ECPU utilization of the pool leader and the members is 40 ECPUs between 2:00pm and 2:30pm, and 128 ECPUs between 2:30pm and 3:00pm.

  The elastic pool is billed 128 ECPUs, one times the pool size, for this billing hour (2-3pm). This case applies when the peak aggregated ECPU usage of the elastic pool for the billing hour is less than or equal to 128 ECPUs.

- **Case-2**: The aggregated peak ECPU utilization of the pool leader and the members is 40 ECPUs between 2:00pm and 2:30pm, and 250 ECPUs between 2:30pm and 3:00pm.

  The elastic pool is billed 256 ECPUs, two times the pool size, for this billing hour (2-3pm). This case applies when the peak aggregated ECPU usage of the elastic pool for the billing hour is less than or equal to 256 ECPUs and greater than 128 ECPUs.

- **Case-3**: The aggregated peak ECPU utilization of the pool leader and the members is 80 ECPUs between 2:00pm and 2:30pm, and 509 ECPUs between 2:30pm and 3:00pm.

  The elastic pool is billed 512 ECPUs, four times the pool size, for this billing hour (2-3pm). This case applies when the peak aggregated ECPU usage of the elastic pool for the billing hour is less than or equal to 512 ECPUs and greater than 256 ECPUs.

See How to Achieve up to 87% Compute Cost Savings with Elastic Resource Pools on Autonomous Database for more details.

Elastic Pool Billing when a Pool is Created or Terminated

When an elastic pool is created or terminated, the leader is billed for the full hour for the elastic pool. In addition, individual instances that are either added or removed from the pool are billed for any compute usage that occurs while the instance is not in the elastic pool (in this case the billing applies to the individual Autonomous Database instance).

- **Pool Creation Example**: Assume there is an Autonomous Database instance with 4 ECPUs that is not part of any elastic pool. At 2:15pm, if you create an elastic pool with this instance with a pool size of 128 ECPUs, the instance becomes a pool leader. Assuming the Autonomous Database idles between 2-3pm, and there are no other Autonomous Database instances in the pool, billing for the hour between 2-3pm is as follows:

  The bill for the period 2-3pm is: (4 * 0.25) + 128 = 129 ECPUs

  Where the (4 * 0.25) is the billing for compute for the fifteen minutes before the Autonomous Database instance created the elastic pool, and 128 ECPUs is the billing for the elastic pool for the hour when the elastic pool is created.

- **Pool Termination Example**: Assume an Autonomous Database instance with 4 ECPUs is the leader of an elastic pool and the pool size is 128 ECPUs. At 4:30pm, if you terminate the elastic pool, the database becomes a standalone Autonomous Database instance that is not part of any elastic pool. Assuming the Autonomous Database idles between 4-5pm, and there are no other Autonomous Database instances in the pool, billing for the hour between 4-5pm is as follows:

  The bill for 4-5pm is: (4 * 0.5) + 128 = 130 ECPUs

Where the (4 * 0.5) is the billing for compute for the thirty minutes after the Autonomous Database instance terminates the elastic pool, and 128 ECPUs is the billing for the elastic pool for the hour when the elastic pool was terminated.

Elastic Pool Billing with Built-in Tools

For either the pool leader or the members, compute resources that are allocated to the built-in tools, OML, Graph, or Data Transforms, are separate and do not count towards the elastic pool total allocation. For billing purposes, the elastic pool leader is billed for any built-in tool ECPU usage by either the leader or elastic pool members, in addition to the elastic pool ECPU usage.

For example, assume there is an elastic pool with a pool size of 128 ECPUs; if in a given billing hour the aggregated peak ECPU utilization of the pool leader and the members is 80 ECPUs for the billing hour, and during this hour the combined total ECPU utilization for instances using built-in tools is 30 ECPUs, the leader is charged for the pool size (128 ECPUs), plus the built-in tool ECPU usage (30 ECPUs), for a total of 158 ECPUs for that hour.

Elastic Pool Billing with Autonomous Data Guard

An elastic pool leader or a member can enable either local or cross-region Autonomous Data Guard, or both local and cross-region Autonomous Data Guard. Depending on whether the standby is local or for a cross-region standby, depending on whether the standby is the pool leader, a pool member, or not an elastic pool leader or a pool member, there are differences in billing. For an Autonomous Data Guard standby database, either the elastic pool leader is charged for the standby's compute usage or the individual standby instance is charged independently.

Elastic Pool Usage Details in OCI Usage Reports and `OCI_USAGE_DATA` View (for records until January 31, 2025)

You can obtain a detailed breakdown of elastic pool usage in the Oracle Cloud Infrastructure (OCI) usage reports and this information is also shown in the `OCI_USAGE_DATA` view. See Cost and Usage Reports and OCI_USAGE_DATA View for more information.

> **Note:**
>
> You can use OCI usage reports and `OCI_USAGE_DATA` view for records until January 31, 2025. The OCI usage reports are deprecated. You can continue to access your existing usage report CSV files until July 31, 2025. See Cost and Usage Reports for more information.

The following table shows the `product/resource` column values in an OCI usage report. The OCI usage report provides details on elastic pool usage for the pool leader and for pool members a given billing hour (similar information is available in `OCI_USAGE_DATA` view):

Elastic Pool Usage Type	Billing Hour Values Shown

Member Compute Usage	For a given pool member where the `product/resourceId` column of the OCI usage report is equal to pool member's OCID and the `product/resource` column of the OCI usage report is equal to `PIC_ADBS_DB_ECPU_PEAK`, the `usage/billedQuantity` of the OCI usage report shows the peak ECPU usage of the member in a specified billing hour.
	You can use the following query to view the same usage details in the `OCI_USAGE_DATA` view:
	`SELECT billed_quantity FROM OCI_USAGE_DATA`
	`WHERE resource_name='PIC_ADBS_DB_ECPU_PEAK' and`
	`resource_id=`*`OCID_of_the_pool_member`*
	`and`
	`interval_usage_start=`*`start_time`* `and`
	`interval_usage_end=`*`end_time`*
	Note:
	If a member has a local a local Autonomous Data Guard standby, its peak usage will be reported as two times (2 `x`) the peak usage. See About Elastic Pool Billing with Autonomous Data Guard Enabled for more information.

Elastic Pool Usage Type	Billing Hour Values Shown
Leader Compute Usage	For a given pool leader where the `product/resourceId` column of the OCI usage report is equal to pool leader's OCID and the `product/resource` column of the OCI usage report is equal to `PIC_ADBS_DB_ECPU_PEAK`, the `usage/billedQuantity` of the OCI usage report shows the peak ECPU usage of the leader in a specified billing hour. You can use the following query to view the same usage details in the `OCI_USAGE_DATA` view: `SELECT billed_quantity FROM OCI_USAGE_DATA WHERE resource_name='PIC_ADBS_DB_ECPU_PEAK' and resource_id=`*`OCID_of_the_pool_leader`* `and interval_usage_start=`*`start_time`*` and interval_usage_end=`*`end_time`* **Note:** If the leader has a local a local Autonomous Data Guard standby, its peak usage will be reported as two times (2 x) the peak usage. See About Elastic Pool Billing with Autonomous Data Guard Enabled for more information.
Aggregated Pool Compute Usage	For a given pool leader where the `product/resourceId` column of the OCI usage report is equal to pool leader's OCID and the `product/resource` column of the OCI usage report is equal to `PIC_ADBS_ELASTIC_POOL_DB_ECPU`, the `usage/billedQuantity` of the OCI usage report shows the aggregated peak ECPU usage of the leader and all members in the specified billing hour. You can use the following query to view the same usage details in the `OCI_USAGE_DATA` view: `SELECT billed_quantity FROM OCI_USAGE_DATA WHERE resource_name='PIC_ADBS_ELASTIC_POOL_DB_ECPU' and resource_id=`*`OCID_of_the_pool_leader`* `and interval_usage_start=`*`start_time`*` and interval_usage_end=`*`end_time`*

Notes for elastic pool billing information in the OCI usage report and the `OCI_USAGE_DATA` view:

- Elastic pool aggregate peak ECPU usage is shown for terminated databases.
- Elastic pool aggregate peak ECPU usage is shown for non-pool databases that were part of an elastic pool during a billing hour when the database was a member of an elastic pool.

See Billing Information: Autonomous Transaction Processing ECPU Compute Model and Track Oracle Cloud Infrastructure Resources, Cost and Usage Reports with Autonomous Database Views for more information.

Elastic Pool Usage Details in OCI Cost Reports and `OCI_COST_DATA` View for Records After January 31, 2025

> **Note:**
>
> Usage details are in the cost reports and the `OCI_COST_DATA` view for usage after January 31, 2025. See Cost and Usage Reports for more information.

You can obtain a detailed breakdown of elastic pool usage in the Oracle Cloud Infrastructure (OCI) cost reports and this information is also shown in the `OCI_COST_DATA` view. See Cost and Usage Reports and OCI_COST_DATA View for more information.

The following table shows the `product/resource` column values in an OCI cost report. The OCI cost report provides details on elastic pool usage for the pool leader and for pool members for a given billing hour and similar information is available in `OCI_COST_DATA` view.

Elastic Pool Usage Type	Billing Hour Values Shown
Member Compute Usage	For a given pool member where the `product/resourceId` column of the OCI cost report is equal to a pool member's OCID and the `product/Description` column is equal to '`Autonomous Database - Elastic Pool Individual DB Peak ECPU`', the `usage/billedQuantity` column of the OCI cost report shows the peak ECPU usage of the member in a specified billing hour. You can use the following query to view the same usage details in the `OCI_COST_DATA` view: `SELECT billed_quantity FROM OCI_COST_DATA WHERE product_description='Autonomous Database - Elastic Pool Individual DB Peak ECPU' and resource_id=`*`OCID_of_the_pool_member`*` and interval_usage_start=`*`start_time`*` and interval_usage_end=`*`end_time`* **Note:** If a member has a local a local Autonomous Data Guard standby, its peak usage will be reported as two times (2 `x`) the peak usage. See About Elastic Pool Billing with Autonomous Data Guard Enabled for more information.

Elastic Pool Usage Type	Billing Hour Values Shown
Leader Compute Usage	For a given pool leader where the `product/resourceId` column of the OCI cost report is equal to pool leader's OCID, and the `product/Description` column is equal to '`Autonomous Database - Elastic Pool Individual DB Peak ECPU`', the `usage/billedQuantity` column of the cost report shows the peak ECPU usage of the leader in a specified billing hour. You can use the following query to view the same usage details in the `OCI_COST_DATA` view: `SELECT billed_quantity FROM OCI_COST_DATA WHERE product_description='Autonomous Database - Elastic Pool Individual DB Peak ECPU' and resource_id=`*`OCID_of_the_pool_leader`*` and interval_usage_start=`*`start_time`*` and interval_usage_end=`*`end_time`* **Note:** If the leader has a local a local Autonomous Data Guard standby, its peak usage will be reported as two times (2 x) the peak usage. See About Elastic Pool Billing with Autonomous Data Guard Enabled for more information.
Aggregated Pool Compute Usage	For a given pool leader where the `product/resourceId` column of the cost report is equal to pool leader's OCID and the `product/Description` column of the cost report is equal to `product/Description` column is equal to '`Autonomous Database - Elastic Pool ECPU`', the `usage/billedQuantity` column of the cost report shows the aggregated peak ECPU usage of the leader and all members in the specified billing hour. You can use the following query to view the same usage details in the `OCI_COST_DATA` view: `SELECT billed_quantity FROM OCI_COST_DATA WHERE product_description='Autonomous Database - Elastic Pool ECPU' and resource_id=`*`OCID_of_the_pool_leader`*` and interval_usage_start=`*`start_time`*` and interval_usage_end=`*`end_time`*

Notes for elastic pool billing information in the OCI cost report and the `OCI_COST_DATA` view:

- Elastic pool aggregate peak ECPU usage is shown for terminated databases.
- Elastic pool aggregate peak ECPU usage is shown for non-pool databases that were part of an elastic pool during a billing hour when the database was a member of an elastic pool.

See Billing Information: Autonomous Transaction Processing ECPU Compute Model and Track Oracle Cloud Infrastructure Resources, Cost and Usage Reports with Autonomous Database Views for more information.

About Elastic Pool Billing with Autonomous Data Guard Enabled

An Autonomous Data Guard primary database can use a local or a cross-region standby that is part of an elastic pool, either as a leader or a member.

- About Elastic Pool Billing with Autonomous Data Guard Enabled with a Local Standby
  When an elastic pool leader or an elastic pool member enables a local Autonomous Data Guard standby, the standby database is part of the elastic pool and you are billed accordingly.
- About Elastic Pool Billing with a Cross-Region Standby
  Describes billing and elastic pool capacity details for a cross-region Autonomous Data Guard standby when the cross-region standby is added to an elastic pool.

About Elastic Pool Billing with Autonomous Data Guard Enabled with a Local Standby

When an elastic pool leader or an elastic pool member enables a local Autonomous Data Guard standby, the standby database is part of the elastic pool and you are billed accordingly.

When you add a local standby, a total of two times (2 x) the primary's ECPU allocation is counted towards the pool capacity (1 x for the primary and 1 x for the standby). Meaning a local standby multiplies the primary's peak usage by 2.

For example, if you create an elastic pool with a pool size of 128 ECPUs, with a pool capacity of 512 ECPUs, adding the following Autonomous Database instance uses the elastic pool capacity:

- 1 instance with 256 ECPUs with local Autonomous Data Guard enabled, for a total of 512 ECPUs allocation from the pool.

  When using this instance, if the peak ECPU utilization is 256 ECPUs for a given billing hour, the overall peak ECPU utilization will be reported as 512 ECPUs because of the local standby database's 2x multiplication factor. And the billing for that hour will be 512 ECPUs (4x pool size).

Similarly, if you create an elastic pool with a pool size of 128 ECPUs, with a pool capacity of 512 ECPUs, adding the following Autonomous Database instances uses the elastic pool capacity as follows:

- 128 instances with 2 ECPUs each, with local Autonomous Data Guard enabled, for a total of 512 ECPUs allocation from the pool.

  When using all of these instances, if the peak ECPU utilization is 256 ECPUs, 128 *2 ECPUs per instance, for a given billing hour, the overall peak ECPU utilization will be reported as 512 ECPUs because of the local standby database's 2x multiplication factor. And the billing for that hour will be 512 ECPUs (4x pool size)

See Enable Autonomous Data Guard for more information.

About Elastic Pool Billing with a Cross-Region Standby

Describes billing and elastic pool capacity details for a cross-region Autonomous Data Guard standby when the cross-region standby is added to an elastic pool.

The databases in an elastic pool must be located in the same region. If you have a cross-region Autonomous Data Guard standby, you can place it in an elastic pool in the standby's region.

- If you increase the primary database's compute resources, the remote elastic pool where the cross-region standby runs must have enough available capacity to accommodate the increase.

  See About Elastic Pools for more information on elastic pool capacity.
- It is not possible to terminate a primary database when its cross-region Autonomous Data Guard standby is a pool leader. In this case you must terminate the elastic pool before you terminate the primary database."

  See Terminate an Elastic Pool for more information.
- A cross-region Autonomous Data Guard standby that is in an elastic pool is billed based on the peak usage of the primary. This applies independent of whether the primary is in an elastic pool.

  For example, in a given billing hour from 1pm to 2pm where the primary's peak usage is 30 ECPUs, the cross-region Autonomous Data Guard standby also shows its peak ECPU usage as 30 ECPUs and this usage is reported to the remote region elastic pool leader.
- A cross-region Autonomous Data Guard standby that is not in an elastic pool, neither as a pool leader nor as a pool member, is billed just like a regular cross-region standby.

  See Oracle Autonomous Database Serverless Features Billing for more information.

Create, Join or Manage an Elastic Pool

Shows the steps to create, join, or manage an elastic pool.

- Create an Elastic Pool
  Shows the steps to create an elastic pool in an existing Autonomous Database instance.
- Create an Elastic Pool While Provisioning or Cloning an Instance
  You can create or join an elastic pool when you provision or clone an Autonomous Database instance.
- Join an Existing Elastic Pool
  Shows the steps to join an existing elastic pool.
- Join an Elastic Pool While Provisioning or Cloning an Instance
  You can create or join an elastic pool when you provision or clone an Autonomous Database instance.
- Change the Elastic Pool Shape
  Shows the steps for the pool leader to change the elastic pool shape for an existing elastic pool.

Create an Elastic Pool

Shows the steps to create an elastic pool in an existing Autonomous Database instance.

> **Note:**
>
> To create an elastic pool the instance must use the ECPU compute model.

- Open the Oracle Cloud Infrastructure Console by clicking the  next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then click **Autonomous Database**.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

To create an elastic pool:

1. On the Autonomous Database Details page select **Manage resource allocation**.

   If **Compute auto scaling** is selected:

   a. Deselect **Compute auto scaling**.

   b. Click **Apply**.

      The **Lifecycle state** changes to **Scaling in Progress**. After some time the **Lifecycle state** changes to **Available**.

   c. On the Autonomous Database Details page select **Manage resource allocation** to display the **Manage resource allocation** area.

2. In the **Manage resource allocation** area, click **Show advanced options**.
3. Select **Enable elastic pool**.
4. Select **Create an elastic pool**.
5. Enter or choose a pool size in the **Pool ECPU count** field.

   Select a pool size from the list of pool shapes: 128, 256, 512, 1024, 2048, or 4096.

   By default each instance in an elastic pool is automatically assigned a maintenance window. By selecting a **Pool ECPU count** that is 1024 greater, you have the option of assigning a custom 2-hour maintenance window during which the leader and all elastic pool members are patched together. To select a custom maintenance window for your elastic pool, file a Service Request at Oracle Cloud Support.

   See About Scheduled Maintenance and Patching for more information.

6. Click **Apply** to create the elastic pool.

When you click **Apply** the **Lifecycle state** changes to **Updating**. After the **Lifecycle state** changes to **Available** the changes apply immediately.

After you create an elastic pool, click **Manage resource allocation** to display the elastic pool information. In the **Manage resource allocation** area, the **Elastic pool** field shows **Enabled**, the **Pool role** field shows **Leader**, and the **Pool ECPU count** field shows the pool size you selected.



Create an Elastic Pool While Provisioning or Cloning an Instance

You can create or join an elastic pool when you provision or clone an Autonomous Database instance.

See Provision an Autonomous Database Instance for details on how to create an Autonomous Database for your workload type using the Create Autonomous Database dialog.

See Clone an Autonomous Database Instance or Clone an Autonomous Database from a Backup for details on cloning.

See Join an Elastic Pool While Provisioning or Cloning an Instance for details on joining an elastic pool while provisioning or cloning.

> **Note:**
>
> To create an elastic pool the instance must use the ECPU compute model.

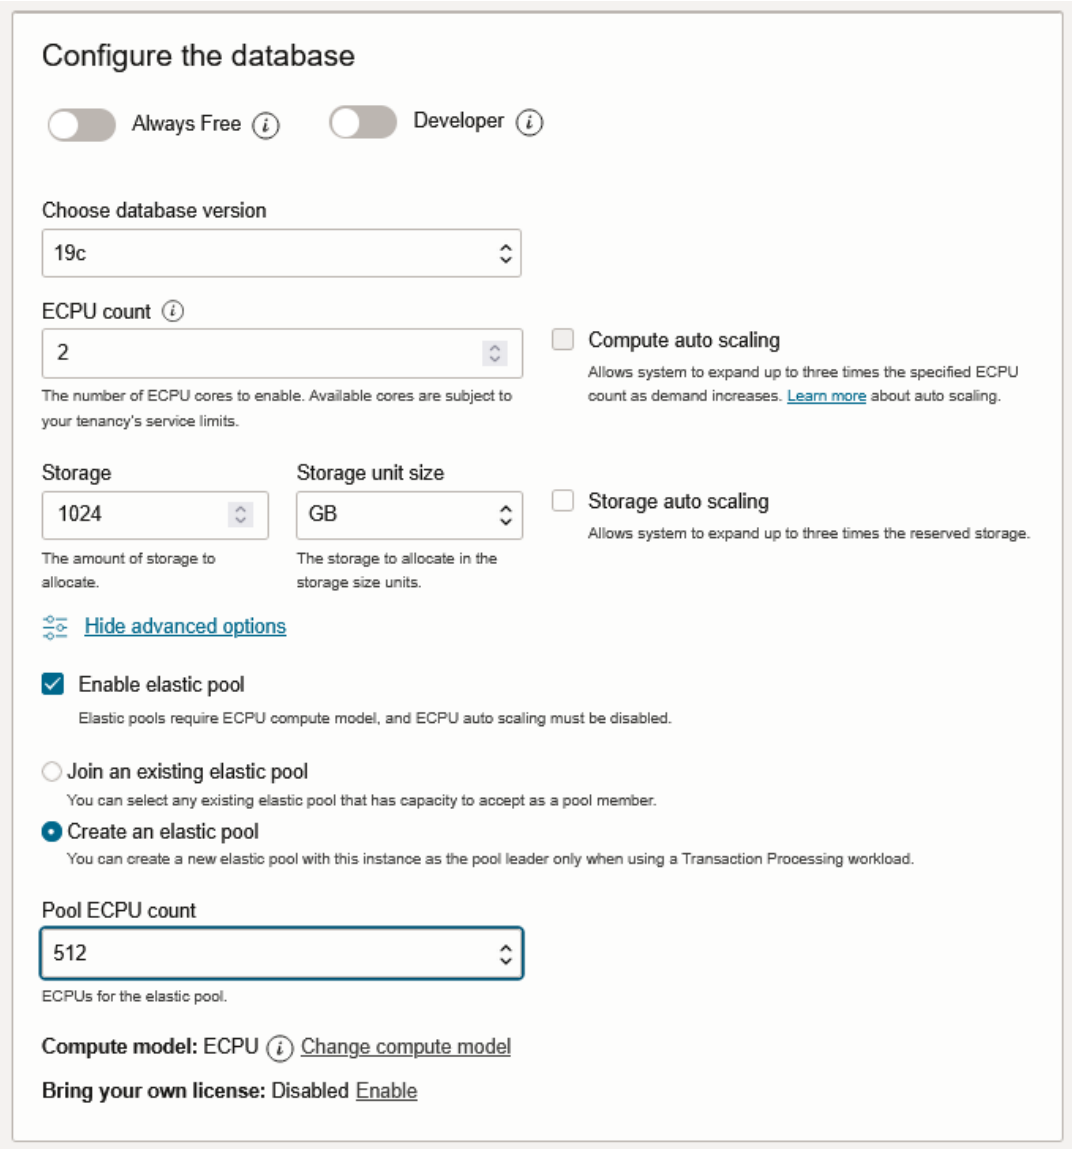
To create an elastic pool while provisioning or cloning:

1. In the **Configure the database** area, deselect **Compute auto scaling**.
2. In the **Configure the database** area, click **Show advanced options**.
3. Select **Enable elastic pool**.
4. Select **Create a elastic pool**.
5. In the **Pool ECPU count** field, select a pool size from the list of pool shapes.

   The valid values that you can select are: 128, 256, 512, 1024, 2048, or 4096.

   By default each instance in an elastic pool is automatically assigned a maintenance window. By selecting a **Pool ECPU count** that is 1024 greater, you have the option of assigning a custom 2-hour maintenance window during which the leader and all elastic pool members are patched together. To select a custom maintenance window for your elastic pool, file a Service Request at Oracle Cloud Support.

   See About Scheduled Maintenance and Patching for more information.
6. If applicable, enable the **Bring your own license** option.



See Choose Bring Your Own License Option While Provisioning or Cloning for more information.

7. Complete the remaining provisioning or cloning steps, as specified in Provision an Autonomous Database Instance, Clone an Autonomous Database Instance, or Clone an Autonomous Database from a Backup.

Join an Existing Elastic Pool

Shows the steps to join an existing elastic pool.

To join an elastic pool, the following is required for the Autonomous Database instance:

- The instance must use the ECPU compute model.

  See Compute Models in Autonomous Database for more information.
- Auto scaling must be disabled.
- The instance must not be a member of an elastic pool.

To join an elastic pool:

1. On the Autonomous Database Details page select **Manage resource allocation**.

   If **Compute auto scaling** is selected:

   a. Deselect **Compute auto scaling**.

   b. Click **Apply**.

      The **Lifecycle state** changes to **Scaling in Progress**. After some time the **Lifecycle state** changes to **Available**.

   c. On the Autonomous Database Details page select **Manage resource allocation** to display the **Manage resource allocation** area.

2. In the **Manage resource allocation** area, click **Show advanced options**.
3. Select **Enable elastic pool**.
4. Select **Join an existing elastic pool**.
5. In the **Select pool leader in *compartment*** field choose a pool leader in a compartment.

   a. Use the compartment shown or click **Change Compartment** to select a different compartment.

   b. Select a pool leader from the list of available pool leaders in the selected compartment.

6. Click **Apply** to add the instance to the elastic pool.

When you click **Apply** the **Lifecycle state** changes to **Updating**. After the **Lifecycle state** changes to **Available** the changes apply immediately.

After you create an elastic pool, click **Manage resource allocation** to see the elastic pool details. In the Manage resource allocations area, under **Elastic pool**, the **Elastic pool** field shows **Enabled**, the **Pool role** field shows **Member**, and the **Elastic pool leader** field shows a link to the pool leader.

Join an Elastic Pool While Provisioning or Cloning an Instance

You can create or join an elastic pool when you provision or clone an Autonomous Database instance.

See Provision an Autonomous Database Instance for details on how to create an Autonomous Database for your workload type using the Create Autonomous Database dialog.

See Clone an Autonomous Database Instance or Clone an Autonomous Database from a Backup for details on cloning.

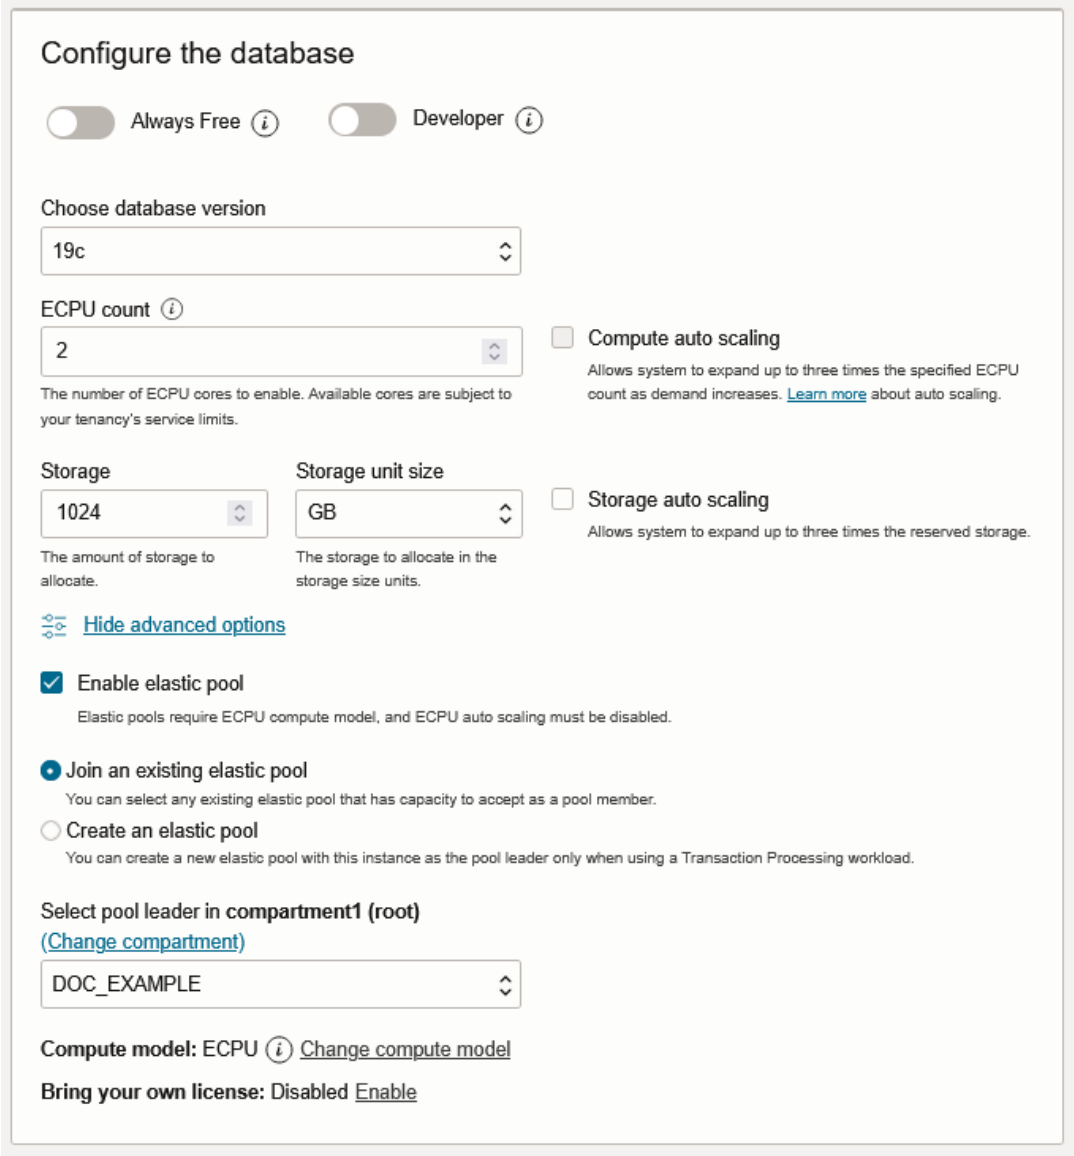
See Create an Elastic Pool While Provisioning or Cloning an Instance for details on creating an elastic pool while provisioning or cloning.

> **Note:**
>
> To join an elastic pool the instance must use the ECPU compute model.

To join an existing elastic pool while provisioning or cloning:

1. In the **Configure the database** area, deselect **Compute auto scaling**.
2. In the **Configure the database** area, click **Show advanced options**.
3. Select **Enable elastic pool**.

4. Select **Join an existing elastic pool**.
5. In the **Select pool leader in *compartment*** field choose a pool leader in a compartment.
   a. Use the compartment shown or click **Change Compartment** to select a different compartment.
   b. Select a pool leader from the list of available pool leaders in the selected compartment.



Complete the remaining provisioning or cloning steps, as specified in Provision an Autonomous Database Instance, Clone an Autonomous Database Instance, or Clone an Autonomous Database from a Backup.

Change the Elastic Pool Shape

Shows the steps for the pool leader to change the elastic pool shape for an existing elastic pool.

> **Note:**
>
> Only a pool leader can modify the pool shape.

To change the shape of an elastic pool (update the pool size):

1. On the Autonomous Database Details page select **Manage resource allocation** to display the **Manage resource allocation** area.
2. In the **Pool ECPU count** field, select a value that is different than the current value.

   By default each instance in an elastic pool is automatically assigned a maintenance window. By selecting a **Pool ECPU count** that is 1024 greater, you have the option of assigning a custom 2-hour maintenance window during which the leader and all elastic pool members are patched together. To select a custom maintenance window for your elastic pool, file a Service Request at Oracle Cloud Support.

   See About Scheduled Maintenance and Patching for more information.
3. Click **Apply**.

   When you click **Apply**, the **Lifecycle state** changes to **Updating**. After the **Lifecycle state** changes to **Available** the changes apply immediately.

> **Note:**
>
> Decreasing the CPU allocation, **Pool ECPU count**, to a value that cannot accommodate all the members of the elastic pool is not allowed.

For example, for an elastic pool with a **Pool ECPU count** of 256 ECPUs and a pool capacity of 1024 ECPUs: If the elastic pool contains eight (8) Autonomous Database instances with 80 ECPUs each for a total of 640 ECPUs, the elastic pool leader cannot decrease the **Pool ECPU count** to 128 ECPUs. In this case, if the pool size were reduced to 128 ECPUs, the pool capacity would be 512 ECPUs, which is less than the total allocation for the pool members (640 ECPUs).

List Elastic Pool Members

Shows the steps for the pool leader to list elastic pool members.

To list elastic pool members:

1. On the elastic pool leader's Autonomous Database details page, under **Resources**, click **Elastic pool members**.
2. The elastic pool members area shows a list of elastic pool members.

This shows the list of elastic pool members for the leader Autonomous Database instance.

If you click  at the end of a row in the list, you can select an action to perform for the member. The possible actions are:

- **View details**: Shows the member's Oracle Cloud Infrastructure Console
- **Copy OCID**: Copies the member's Autonomous Database instance OCID.
- **Remove from pool**: Brings up a dialog where you can confirm to remove the Autonomous Database instance from the pool.

Offload Queries to Member Refreshable Clones

When you have heavy read workloads, where a relatively small amount of data is scanned by multiple queries, you can offload queries (reads) either from an elastic pool leader or from an elastic pool member to Refreshable Clones.

- About Query Offloading
  Provides information about using query offload with an elastic pool and describes the query offload features.
- Enable Query Offload
  Describes how to enable query offload for an elastic pool leader or for an elastic pool member.
- Disable Query Offload
  Describes how to disable query offload for an elastic pool leader or for an elastic pool member.
- Query Offload from PL/SQL
  When query offload is enabled, queries are offloaded to Refreshable Clones even if they are embedded within PL/SQL. The PL/SQL can be stored procedures, functions, packages, or anonymous blocks.
- Query Offload from a Scheduler Job
  When query offload is enabled, queries from within an Oracle Scheduler job action are offloaded to Refreshable Clones. The queries are offloaded whether the job runs in the foreground or in the background.
- Notes for Offloading Queries
  Provides additional notes for the query offload feature.

About Query Offloading

Provides information about using query offload with an elastic pool and describes the query offload features.

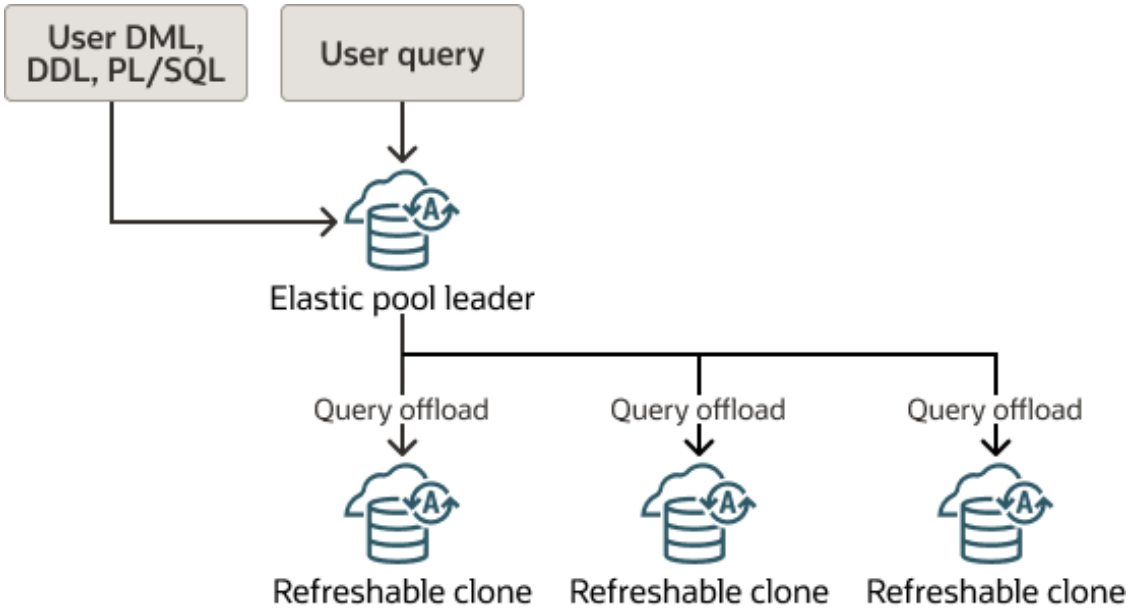
Elastic pool query offload provides performance benefits by allowing one or more Refreshable Clones to handle queries for either an elastic pool leader or for an elastic pool member. This feature also allows you to add Refreshable Clones to accommodate increasing query (read) demand. Offloading queries allows your application to scale horizontally, where you can add Refreshable Clones to maintain overall system performance as needed to satisfy your query request volume.

When query offload is enabled, queries are submitted either to the elastic pool leader or to an elastic pool member and one or more available Refreshable Clones become candidates for query offloading. In addition, when more Refreshable Clones are added, query offload dynamically adjusts to make use of the new resources.

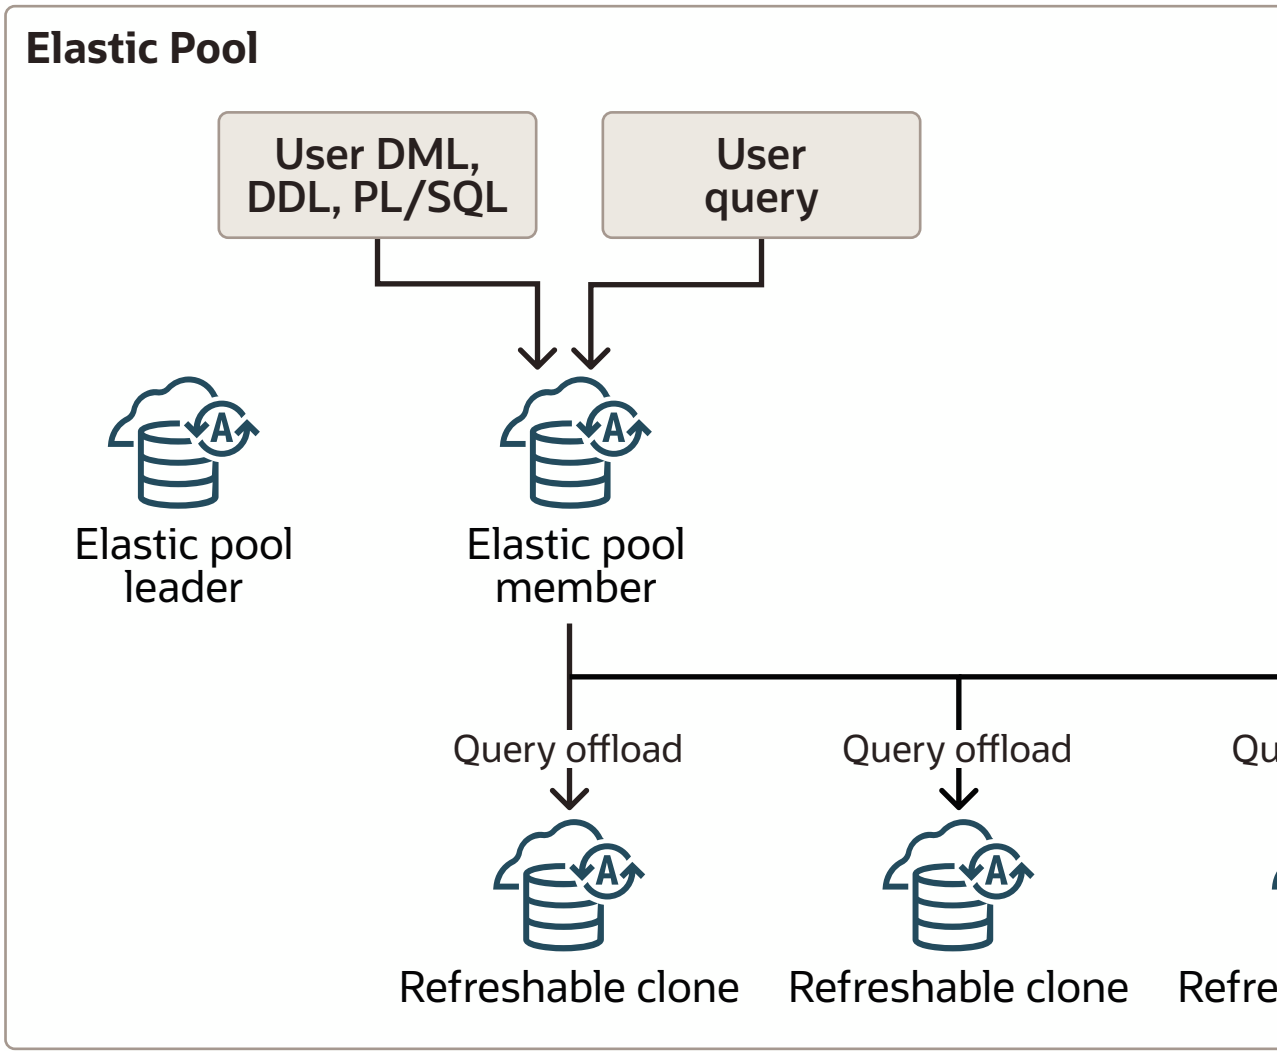
One use case for query offload is to enable the feature during peak hours to take load off of either the elastic pool leader or an elastic pool member. During quiet hours, you can disable query offload to perform maintenance operations such as refreshing the Refreshable Clones.

By default query offload considers queries from any session. Alternatively you can offload queries from a list of sessions that you specify by module or action name.

The following figure shows offloading queries from the elastic pool leader.



The following figure shows offloading queries from an elastic pool member.



Data on the Refreshable Clones is up to date based on the last refresh time for each refreshable clone. This means when query offload is enabled you perform all DDL, DML and PL/SQL operations either on the elastic pool leader or on the elastic pool member that is offloading queries. Then, after a Refreshable Clone is refreshed, the changes are reflected on the Refreshable Clone.

See Use Refreshable Clones with Autonomous Database for more information.

Query Offload Features

- **Dynamic Addition**: Refreshable Clones may be added as members of the elastic pool at any time. Query offload dynamically adjusts to make use of new members.
- **Dynamic Removal**: Refreshable Clones may be removed as members of the elastic pool at any time. Query offload dynamically adjusts to stop offloading queries to a Refreshable Clone that has been removed from the elastic pool.
- **Session Based Sticky Offload**: Query offload is sticky within a session, meaning if a query in a session has been offloaded to a particular Refreshable Clone, then query offload uses the same Refreshable Clone to offload subsequent queries in the same session.
- **One-to-One Service Mapping**: There is a one-to-one mapping between the service used in either in the elastic pool leader or the elastic pool member by the original query and the

service used in a Refreshable Clone for any offloaded query. For example, if a session is connected to the MEDIUM service, then query offload for that session also uses the MEDIUM service on the Refreshable Clone.

- **Determine if Session is Offloaded**: Using a `SYS_CONTEXT` query you can determine if a session runs on the elastic pool leader or member, or is offloaded to a Refreshable Clone.
- **Query Offload from PL/SQL**: Queries are offloaded to a Refreshable Clone even if they are embedded within PL/SQL. The PL/SQL can be stored procedures, functions, packages, or anonymous blocks.
- **Query Offload from Scheduler Job**: Queries from within an Oracle Scheduler job actions are offloaded. Oracle Scheduler job queries are offloaded whether the job runs in the foreground or in the background.

DBA_PROXY_SQL Views

Use the `DBA_PROXY_SQL` views `DBA_PROXY_SQL_ACTIONS` and `DBA_PROXY_SQL_MODULES` to display the list of modules or actions that are configured for query offload. You must query these views from a session that is not enabled for query offload.

Enable Query Offload

Describes how to enable query offload for an elastic pool leader or for an elastic pool member.

The following are requirements for enabling query offload:

- You can enable query offload for an elastic pool leader or for an elastic pool member with no Refreshable Clones. After you enable query offloading you can add Refreshable Clones and the query offload feature dynamically adjusts to make use of the refreshable clones.
- A Refreshable Clone that is a candidate for query offload must:
  - Have the elastic pool leader as its source database and be in the same region as the elastic pool leader.

    or

    Have an elastic pool member as its source database and be in the same region as the elastic pool member.
  - Be an elastic pool member.

To enable query offload:

1. Verify that the Autonomous Database instance is an elastic pool leader or an elastic pool member:
   a. Use the following query to verify that an instance is an elastic pool leader:

      ```
      SELECT sys_context('userenv', 'is_elastic_pool_leader') FROM DUAL;
      ```

      This query should return `YES`.
   b. Use the following query to verify that an instance is an elastic pool member:

      ```
      SELECT sys_context('userenv', 'is_elastic_pool_member') FROM DUAL;
      ```

      This query should return `YES`.

> **Note:**
>
> There may be a delay of up to 15 minutes for the `sys_context` value to reflect the current value if the elastic pool has recently been created or modified.

2. Enable query offload.

   There are two choices: you can enable query offload for queries from any session or limit query offload to the sessions you specify by module or action name.

   - To enable query offload for all sessions run `DBMS_PROXY_SQL.ENABLE_READ_ONLY_OFFLOAD` without parameters. For example:

     ```
     EXEC DBMS_PROXY_SQL.ENABLE_READ_ONLY_OFFLOAD;
     ```

   - To enable query offload for specific sessions by module name or action name, run `DBMS_PROXY_SQL.ENABLE_READ_ONLY_OFFLOAD` with the `module_name` or `action_name` parameters.

     See Enable Query Offload for Named Modules or Actions for more information.

When query offload is enabled for a session you can find the name of the Refreshable Clone to which queries are offloaded. For example, from the session run this query:

```
SELECT sys_context('userenv', 'con_name') FROM DUAL;
```

If queries are not being offloaded to a Refreshable Clone this query shows the name of the elastic pool leader (or of the elastic pool member that is not a refreshable clone).

See SYS_CONTEXT for more information.

See ENABLE_READ_ONLY_OFFLOAD Procedure for more information.

- Enable Query Offload for Named Modules or Actions
  Describes how to enable query offload for sessions with named modules or actions.

Enable Query Offload for Named Modules or Actions

Describes how to enable query offload for sessions with named modules or actions.

To enable query offload for specific sessions by module name or action name:

1. Verify that the instance is an elastic pool leader or an elastic pool member.
   a. Use the following query to verify that an instance is an elastic pool leader:

      ```
      SELECT sys_context('userenv', 'is_elastic_pool_leader') FROM DUAL;
      ```

      This query should return `YES`.

   b. Use the following query to verify that an instance is an elastic pool member:

      ```
      SELECT sys_context('userenv', 'is_elastic_pool_member') FROM DUAL;
      ```

      This query should return `YES`.

> **Note:**
>
> There may be a delay of up to 15 minutes for the `sys_context` value to reflect the current value if the elastic pool has recently been created or modified.

2. Use the routines `SET_ACTION` or `SET_MODULE` in `DBMS_APPLICATION_INFO` to set the module name and or the action name in the current session.

   For example

```
CREATE or replace PROCEDURE add_employee(
  name VARCHAR2,
  salary NUMBER,
  manager NUMBER,
  title VARCHAR2,
  commission NUMBER,
  department NUMBER) AS
BEGIN
  DBMS_APPLICATION_INFO.SET_MODULE(
    module_name => 'add_employee',
    action_name => 'insert into emp');
  INSERT INTO emp
    (ename, empno, sal, mgr, job, hiredate, comm, deptno)
    VALUES (name, emp_seq.nextval, salary, manager, title, SYSDATE,
            commission, department);
  DBMS_APPLICATION_INFO.SET_MODULE(null,null);
END;
```

   See DBMS_APPLICATION_INFO for more information.

3. Enable query offload and specify the eligible sessions by module name or by action name (or by both module name and action name).

   For example:

```
DECLARE
   mod_values clob := to_clob('["mod1", "mod2"]');
   act_values clob := to_clob('["act1", "act2"]');
BEGIN
   DBMS_PROXY_SQL.ENABLE_READ_ONLY_OFFLOAD(
      module_name => mod_values,
      action_name => act_values);
END;
/
```

When query offload is enabled for a session, from the session you can find the name of the Refreshable Clone to which queries are offloaded. For example:

```
SELECT sys_context('userenv', 'con_name') FROM DUAL;
```

If queries are not being offloaded to a Refreshable Clone this query shows the name of the elastic pool leader (or of the elastic pool member that is not a refreshable clone).

See SYS_CONTEXT for more information.

The views `DBA_PROXY_SQL_ACTIONS` and `DBA_PROXY_SQL_MODULES` display the list of modules or actions that are configured for query offload. You must query these views from a session that is not enabled for query offload. See DBA_PROXY_SQL Views for more information.

See ENABLE_READ_ONLY_OFFLOAD Procedure for more information.

Disable Query Offload

Describes how to disable query offload for an elastic pool leader or for an elastic pool member.

On the instance that is offloading queries, disable query offload.

For example:

```
EXEC DBMS_PROXY_SQL.DISABLE_READ_ONLY_OFFLOAD;
```

See DISABLE_READ_ONLY_OFFLOAD Procedure for more information.

Query Offload from PL/SQL

When query offload is enabled, queries are offloaded to Refreshable Clones even if they are embedded within PL/SQL. The PL/SQL can be stored procedures, functions, packages, or anonymous blocks.

For example, you can demonstrate a query being offloaded to a Refreshable Clone when you enable query offload for a specific module and action. In this example, assume that the leader's data has been updated and the Refreshable Clone has not yet been updated (so the Refreshable Clone's data is out of date and different). In this example, with different values on the elastic pool leader and the Refreshable Clone, you can see when data is coming from either the leader or from the Refreshable Clone.

1. Enable query offload and specify the eligible sessions by module name and action name. For example, '`mod1`', '`act1`'.

```
DECLARE
   mod_values clob := to_clob('["mod1"]');
   act_values clob := to_clob('["act1"]');
BEGIN
   DBMS_PROXY_SQL.ENABLE_READ_ONLY_OFFLOAD(
      module_name => mod_values,
      action_name => act_values);
END;
/
```

2. Define the function `f1` to test query offload.

```
SQL> create or replace function f1 (n number)
  2  return number
  3  as
  4    l_cnt number;
  5  begin
  6    select sum(c1) into l_cnt from u2.tab1;
  7    return l_cnt;
  8  end;
  9  /
```

```
Function created.
```

3. Run the function `f1`. The query with function `f1` runs on the elastic pool leader (or on an elastic pool member that has offload query enabled). This session runs with unspecified module and action names that do not match those specified for query offload in Step 1.

```
SQL> -- expected to fetch from Elastic Pool Leader and returns value 40
SQL> select f1(3) from dual;
     F1(3)
----------
        40

1 row selected.
```

4. Specify the module name and action name for the session to match the names specified in Step 1. The queries from this session are now eligible to be offloaded to a Refreshable Clone.

```
SQL> exec dbms_application_info.set_module('mod1', 'act1');

PL/SQL procedure successfully completed.
```

5. Run the function `f1` again. In this case the module name and action name match the names specified in Step 1. The query in function `f1` is offloaded to a Refreshable Clone.

```
SQL> -- Expected to fetch from Refreshable Clone and returns value 10
SQL> select f1(3) from dual;

     F1(3)
----------
        10

1 row selected.
```

6. Specify the module and action names as something other than `mod1` and `act1`. The queries from this session are no longer eligible to be offloaded to a Refreshable Clone, due to the module and action names not matching those specified in Step 1.

```
SQL> exec dbms_application_info.set_module('random', 'random');

PL/SQL procedure successfully completed.
```

7. Run the query with function `f1`. Because the module and action name for this session do not match those set in Step 1, this query is not offloaded and runs on the elastic pool leader.

```
SQL> -- expected to fetch from Elastic Pool Leader and returns value 40
SQL> select f1(3) from dual;
     F1(3)
----------
        40

1 row selected.
```

Query Offload from a Scheduler Job

When query offload is enabled, queries from within an Oracle Scheduler job action are offloaded to Refreshable Clones. The queries are offloaded whether the job runs in the foreground or in the background.

For example, similar to the previous example for PL/SQL, you can demonstrate a query being offloaded to a Refreshable Clone when you enable query offload for an Oracle Scheduler job. In this example, assume that the leader's data has been updated and the Refreshable Clone has not yet been updated (so the Refreshable Clone's data is out of date and different). In this example, with different values on the elastic pool leader and the Refreshable Clone, you can see when data is coming from either the leader or from the Refreshable Clone.

```
SQL> create or replace procedure sproc1 (n number)
  2  as
  3    l_cnt number;
  4  begin
  5    select sum(c1) into l_cnt from u2.tab1;
  6    dbms_output.put_line('l_cnt is ' || l_cnt);
  7  end;
  8  /

Procedure created.

Offload Queries from Elastic Pool Leader to Member Refreshable Clones

SQL> create or replace procedure sproc1 (n number)
  2  as
  3    l_cnt number;
  4  begin
  5    select sum(c1) into l_cnt from u2.tab1;
  6    dbms_output.put_line('l_cnt is ' || l_cnt);
  7  end;
  8  /

Procedure created.

SQL>
SQL> BEGIN
  2    dbms_scheduler.create_job(job_name    => 'PROXYTEST1',
  3                              job_type    => 'PLSQL_BLOCK',
  4                              job_action  => 'BEGIN sproc1(3); END;',
  5                              enabled     => FALSE);
  6  END;
  7  /

PL/SQL procedure successfully completed.

SQL>
SQL> -- Job runs in the foreground in the current session
SQL> -- expected to fetch from Leader: 40
SQL> exec dbms_scheduler.RUN_JOB(job_name => 'PROXYTEST1',
use_current_session => true);
l_cnt is 40
```

```
PL/SQL procedure successfully completed.

SQL>
SQL> exec dbms_proxy_sql.enable_read_only_offload;

PL/SQL procedure successfully completed.

SQL> show con_name

CON_NAME
------------------------------
CDB1_PDB1
SQL>
SQL> set serveroutput on
SQL> -- Job runs in the foreground in the current session
SQL> -- expected to fetch from Refreshable Clone: 10
SQL> exec dbms_scheduler.RUN_JOB(job_name => 'PROXYTEST1',
use_current_session => true);
l_cnt is 10

PL/SQL procedure successfully completed.

SQL>
SQL> -- Job runs in the background and is expected to fetch from Refreshable
Clone: 10
SQL> exec dbms_scheduler.RUN_JOB(job_name => 'PROXYTEST1',
use_current_session => false);

PL/SQL procedure successfully completed.
```

Notes for Offloading Queries

Provides additional notes for the query offload feature.

Notes for query offload:

- **Addition or Removal of a Refreshable Clone**: When a Refreshable Clone is added as an elastic pool member, there can be a delay of up to fifteen minutes for the addition to be reflected in the elastic pool leader. A newly added Refreshable Clone is not considered as a target for query offloads until the leader is aware of the newly added elastic pool member.

  Similarly, when a Refreshable Clone is removed from the elastic pool, there can be a delay of up fifteen minutes for the removal to be reflected in the elastic pool leader. The removed Refreshable Clone is considered as a target for query offloads until the leader knows that the member has been removed from the elastic pool.
- **Refresh of Refreshable Clone**: When a Refreshable Clone is being refreshed, queries that are offloaded to the Refreshable Clone may be delayed.
- **Public and Private Endpoint Support**: Query offload is supported for instances on a public endpoint and for instances on a private endpoint.
- **Disconnected Refreshable Clone**: If a Refreshable Clone becomes disconnected, it is no longer eligible as a target for query offload. There can be a delay of up to fifteen minutes for the elastic pool leader to detect that a Refreshable Clone has become disconnected and to stop offloading queries to it.

To ensure that a Refreshable Clone remains available for query offloading, either enable automatic refreshes or periodically manually refresh the Refreshable Clone.

See Refresh a Refreshable Clone on Autonomous Database for more information.

- **Case Sensitive Module Name and Action Name**: When you offload queries from a list of sessions that you specify by module or action name, the module name and action name set with `DBMS_APPLICATION_INFO` are case sensitive. `DBMS_PROXY_SQL.ENABLE_READ_ONLY_OFFLOAD` adheres to this case sensitive behavior. For example, when a session has a lowercase module name, or mixed case module name, the case must match in the parameter values when you enable query offload with `DBMS_PROXY_SQL.ENABLE_READ_ONLY_OFFLOAD` and you include the `module_name` or `action_name` parameters.

  You can check the module name and action name for the current session using `DBMS_APPLICATION_INFO.READ_MODULE`:

```
set serveroutput on;
declare
  l_mod varchar2(50);
  l_act varchar2(50);
BEGIN
DBMS_APPLICATION_INFO.READ_MODULE(l_mod, l_act);
  DBMS_OUTPUT.PUT_LINE('l_mod: ' || l_mod);
  DBMS_OUTPUT.PUT_LINE('l_act: ' || l_act);
END;
/
```

- **Typical Use Case**: One use case for query offload is to enable the feature during peak hours to take load off of the Elastic Pool Leader (or off of the Elastic Pool Member where query offload is enabled). During quiet hours, you can disable query offload to perform maintenance operations such as refreshing the Refreshable Clones.

Remove Pool Members from an Elastic Pool

As an elastic pool member you can remove yourself from an elastic pool. As an elastic pool leader you can remove pool members from an elastic pool.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then, depending on your workload click one of: Autonomous Data Warehouse, Autonomous JSON Database, or Autonomous Transaction Processing.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Display name** column.

As a pool member, you can remove your instance from an elastic pool:

1. On the **Details** page, under **Resource allocation**, in the **Elastic pool** field click **Leave pool**.

   This shows the **Leave pool** confirmation dialog.
2. In the **Leave pool** confirmation dialog, enter the database name.
3. Click **Leave**.

When you click **Leave**, the **Lifecycle state** changes to **Updating**. After the **Lifecycle state** changes to **Available** the changes apply immediately.

- As Pool Leader Remove Members from an Elastic Pool
  An elastic pool leader can remove pool members from an elastic pool.
- Notes for Leaving an Elastic Pool
  Provides information about resources when a member or the leader leaves an elastic pool.

As Pool Leader Remove Members from an Elastic Pool

An elastic pool leader can remove pool members from an elastic pool.

1. On the Autonomous Database Details page, under **Resources**, click **Elastic pool members**.

   This shows the **Elastic pool members** area, with a list showing details for each instance that is an elastic pool member.

2. Click  at the end of a row for an instance you want to remove and in the drop down list select **Remove from pool**.

   This shows the **Remove from pool** confirmation dialog.

3. Click **Leave** to confirm.

   When you click **Leave**, the **Lifecycle state** changes to **Updating**. After the **Lifecycle state** changes to **Available** the changes apply immediately.

Notes for Leaving an Elastic Pool

Provides information about resources when a member or the leader leaves an elastic pool.

- When a pool member or the leader leaves an elastic pool, auto scaling is disabled. After leaving the elastic pool you can enable auto-scaling for the instance.
- When a pool member leaves an elastic pool, the elastic pool has more resources available. For example, if the elastic pool were fully allocated up to the pool capacity, and an instance with 10 ECPUs leaves the pool, the elastic pool would have 10 available ECPUs.
- Billing for an Autonomous Database instance that leaves an elastic pool returns to individual instance billing, based on the compute and storage resources that the individual instance uses:
  - If a pool member with 2 ECPUs or more leaves the pool, the individual instance's ECPU allocation remains and the instance is billed for that number of ECPUs.
  - If a pool member with 1 ECPU leaves the pool, the ECPU allocation is automatically set to 2 ECPUs and the instance is billed for 2 ECPUs going forward, unless it's scaled up.

Terminate an Elastic Pool

Shows the steps to terminate an elastic pool. Only the pool leader can terminate an elastic pool.

> **Note:**
>
> Terminating an elastic pool is only allowed when there are no pool members in the elastic pool.

- Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation menu click **Oracle Database** and then click **Autonomous Database**.
- On the Autonomous Databases page select the Autonomous Database that is the pool leader from the links under the **Display Name** column.

To terminate an elastic pool:

1. On the Autonomous Database Details page, select **Manage resource allocation**.
2. In the **Elastic pool** area, select **Terminate pool**.
3. Click **Apply** to terminate the elastic pool.

When you click **Apply**, the Lifecycle state changes to **Updating**. After the Lifecycle state changes to **Available** the changes apply immediately.

Elastic Pool Notes

Provides information about elastic pool usage.

Notes for elastic pool usage:

- Starting and stopping Autonomous Database instances that are members of an elastic pool does not depend on the state of the leader. That is, you can independently stop and start each instance that is part of an elastic pool, including the leader and members of the elastic pool that are not the leader.
- By default each instance in an elastic pool is automatically assigned a maintenance window. By selecting a pool shape that is 1024 ECPUs or greater, you have the option of assigning a custom 2-hour maintenance window during which the leader and all elastic pool members are patched together. To select a custom maintenance window for your elastic pool, file a Service Request at Oracle Cloud Support.
- Provisioning a pool leader or a member is subject to service limits enforced at the tenancy or compartment levels.
- In an elastic pool, the pool leader's licensing selections determine the license requirements for the entire pool. If you want to use BYOL for an elastic pool and set a **BYOL License Limit**, you must set the limit value on the pool leader's instance. The license selections for pool members, including any **BYOL License Limit** does not apply while an instance is a member of a pool. The License selections, including any **BYOL License Limit** come into effect only if a pool member leaves the elastic pool.

For example, if an elastic pool includes a member instance that uses Autonomous Data Guard with an 8 ECPU Primary with the BYOL license type set, the Primary and Standby instance's licensing selection is set as follows:

- When the Primary is a pool member, the license type of the Primary can be set but it is does not apply for the Primary; the Primary uses the license type of the elastic pool leader. The license type selection on the Primary will apply when and if the Primary leaves the elastic pool.
- For a local Standby, the Autonomous Data Guard Standby is in the elastic pool. In this case, the local Standby uses the same license type as the elastic pool leader.
- For a cross-region or a cross-tenancy Standby, the Autonomous Data Guard Standby is not in the elastic pool. In this case, the cross-region Standby or cross tenancy Standby uses the license type specified on the Autonomous Data Guard Primary, and has its own BYOL license limit that can be set, if required.

Migrate to Autonomous Database

Database migration is the process of migrating your database from a source database to a target database.

To migrate your database to an Autonomous Database, Oracle provides different migration options.

- Migrate Oracle Databases to Autonomous Database
- Migrate Applications from Non-Oracle SQL to Oracle SQL
  Application migration involves moving an application from one environment to another. The migration process includes a variety of tasks, including migrating SQL statements from a non-Oracle database to Oracle SQL.

Migrate Oracle Databases to Autonomous Database

Oracle provides options to migrate your database to Autonomous Database.
You can choose one of these options depending on the amount of data you want to migrate, downtime requirements during the migration, network bandwidth between your source database and your Oracle Cloud Infrastructure region, and whether you want manual control over the migration.

- Migration Prerequisites
  Before starting your migration, it is recommended to run the Cloud Premigration Advisor Tool (CPAT), which helps you evaluate your source database for compatibility with the Autonomous Database.
- Recommended Migration Methods
  Following are the recommended migration methods to migrate your database into an Autonomous Database.
- Migrate Existing SQL Execution Plans to Autonomous Database to Mitigate the Risk of SQL Performance Regressions
  Describes how to reduce the risk of SQL performance regressions when you migrate to Autonomous Database.

Migration Prerequisites

Before starting your migration, it is recommended to run the Cloud Premigration Advisor Tool (CPAT), which helps you evaluate your source database for compatibility with the Autonomous Database.

CPAT identifies potential actions you may need to take before or during the migration, prioritizing their importance and suggesting resolutions. Some migration tools and services automatically run this advisor.

See Cloud Premigration Advisor Tool for more information.

Recommended Migration Methods

Following are the recommended migration methods to migrate your database into an Autonomous Database.

Online Migration

If your source database needs to be always online and you can only take minimal downtime for the migration, you can use one of the online migration services and tools below, independent of your database size. Note that both options require the source database to be on version 11.2.0.4 or later. If your source database is using an older version, you need to upgrade it first if you want to use one of the online migration methods.

Following are the online migration methods:

- **OCI Database Migration**: OCI Database Migration, a managed cloud service, provides validated, cross-version, fault-tolerant, and incremental Oracle Database migrations for online and offline use cases. You can choose this option if you want a fully managed cloud service that handles the migration for you.

  See OCI Database Migration for more information.

- **Zero Downtime Migration (ZDM)**: Zero Downtime Migration (ZDM) is a tool with a command line interface that you install and run on a host you provision. Zero Downtime Migration provides online and offline migration options. You can use this option if you want more control over the migration.

  See Zero Downtime Migration for more information.

> **Note:**
>
> These methods also run the Cloud Premigration Advisor Tool (CPAT) to check your source database for compatibility with the Autonomous Database and use Oracle Data Pump and Oracle GoldenGate to carry out the migration for you.

Offline Migration

If you can take your source database offline for the migration, you can use one of the following migration options in addition to OCI Database Migration and Zero Downtime Migration (ZDM).

> **Note:**
>
> This is a manual method; you must orchestrate the export/import manually.

Following is the offline migration method:

- **Oracle Data Pump**: Oracle Data Pump technology enables very high-speed movement of data and metadata from one database to another. You can export your source database or individual schemas as dump files and import them into the Autonomous Database. You

can store the dump files on the Object Store or a network filesystem (NFS) that you can attach to your source database and to your Autonomous Database.

See Overview of Oracle Data Pump for the product documentation for Oracle Data Pump.

See Import Data Using Oracle Data Pump on Autonomous Database for more information about how to export and import your source schemas into the Autonomous Database.

Other Migration Methods

In addition to the recommended migration tools and services, you can use the following offline data migration methods:

> **Note:**
>
> These methods require more manual steps, coding, and advanced knowledge of each method. You must also migrate your source schema metadata using your own `DDL` scripts or tools such as Oracle Data Pump before initiating the data load using these methods. These options may be preferable for small data sets for which you want to avoid installing additional tools or using cloud migration services.

Following are the other offline migration methods:

- **Database Links**: You can create database links from your Autonomous Database to your source database and insert data into your tables by running queries over the database link. You can use multiple sessions to load multiple tables concurrently to improve the performance of the data load.

  See Create Database Links from Autonomous Database to Another Autonomous Database for more information about creating database links from your Autonomous Database to your source databases.

- **Unload and Load Using Files**: You can unload data from your source database into files and load those files into your target tables in your Autonomous Database. Autonomous Database supports loading data from different file formats including CSV, JSON, Parquet, and more. You can place the files on the Object Store or a network filesystem (NFS) attached to your Autonomous Database.

  See Load Data from Files in the Cloud for more information about how to load data from files in the object store.

  See Attach Network File System to Autonomous Database for attaching NFS directories to your Autonomous Database.

  See Load Data or Query Data from Files in a Directory for more information about loading data from files on a filesystem.

- **Materialized Views**: Using database links from your Autonomous Database to your source database, you can also create materialized views in your target database. As changes happen on your source tables, you can refresh the materialized views to keep them updated during the migration.

  See Managing Read-Only Materialized Views for more information about using materialized views to replicate data.

> **Note:**
>
> Deciding the migration tool or service depends on multiple factors such as your source database, source data format, data volume, and complexity. To help you identify the most optimal solution for migrating your data to the Autonomous Database, Oracle provides an advisory utility called Oracle Cloud Migration Advisor. See Migrate Oracle Databases to OCI for more information about this utility.

Migrate Existing SQL Execution Plans to Autonomous Database to Mitigate the Risk of SQL Performance Regressions

Describes how to reduce the risk of SQL performance regressions when you migrate to Autonomous Database.

- About Migrating to Autonomous Database with Real-time SQL Plan Management (SPM)
  Describes using Real-time SQL Plan Management when you migrate from a source Oracle Database to Autonomous Database.
- Enable Automatic SQL Tuning Set on the Source Oracle Database
  Before you migrate to Autonomous Database enable Automatic SQL Tuning Set (ASTS) on your source Oracle Database.
- Migrate Your Data to Autonomous Database
  After you capture a sufficient amount of SQL on the source Oracle Database with Automatic SQL Tuning Set (ASTS) enabled, perform the migration to Autonomous Database.
- Export Automatic SQL Tuning Set from Your Source Oracle Database
  After you perform the migration to Autonomous Database, export the Automatic SQL Tuning Set (ASTS) from your source Oracle Database.
- Import Automatic SQL Tuning Set to Autonomous Database
  After you perform the migration to Autonomous Database and you export the Automatic SQL Tuning Set (ASTS) from the source Oracle Database you are migrating, import the ASTS to your Autonomous Database.
- Verify Your Real-time SPM Setting on Autonomous Database
  Describes the steps to verify that real-time SPM is enabled on Autonomous Database.

About Migrating to Autonomous Database with Real-time SQL Plan Management (SPM)

Describes using Real-time SQL Plan Management when you migrate from a source Oracle Database to Autonomous Database.

When you migrate from a source Oracle Database to Autonomous Database you can use Real-time SQL Plan Management (SPM). This enables you to capture the SQL execution plans from your source database and move them into Automatic SQL Tuning Set (ASTS) on Autonomous Database so that plans continue to perform with the same or better performance after the migration. Real-time SPM allows plans to change on Autonomous Database, but if Real-time SPM sees poor performance, it may use a plan that came from the pre-migrated database if that plan provides better performance (Real-time SPM uses the plans in Automatic SQL Tuning Set (ASTS) only if they provide better performance).

Keep the following in mind for using Real-time SPM with a migration to Autonomous Database:

- SPM relies on repeatable SQL statements. SPM is not suitable for databases that use literal values in SQL statements or for SQL that is highly dynamic, such as in ad-hoc query

environments. However, if the SQL statements use literal values and the CURSOR_SHARING parameter is set to FORCE, SPM will work.

- In the source Oracle Database that you are migrating to Autonomous Database you capture application SQL statements in a SQL tuning set (SYS_AUTO_SYS). This can consume space in SYSAUX, but typically consumes no more than a few gigabytes (even for large systems). You can monitor the SYSAUX usage and increase the size of the tablespace, if necessary.
- Real-time SPM cannot prevent all performance regressions, but it can significantly reduce the risk of performance regressions happening due to SQL execution plan changes.

Perform the following steps to enable SQL Plan Management (SPM) on your source Oracle Database and to migrate the database to Autonomous Database:

1. Enable Automatic SQL Tuning Set on the Source Oracle Database
2. Migrate Your Data to Autonomous Database
3. Export Automatic SQL Tuning Set from Your Source Oracle Database
4. Import Automatic SQL Tuning Set to Autonomous Database
5. Verify Your Real-time SPM Setting on Autonomous Database

See the following for more information:

- CURSOR_SHARING
- What is Real-time SQL plan Management?
- Using the Automatic SQL Tuning Set in Oracle Database 19c

Enable Automatic SQL Tuning Set on the Source Oracle Database

Before you migrate to Autonomous Database enable Automatic SQL Tuning Set (ASTS) on your source Oracle Database.

ASTS needs to run for enough time to cover your workload and capture all or most of the SQL statements and their execution plans. So, consider enabling ASTS with lead time in advance of the migration to Autonomous Database. For example, for a financial or sales application, capture month-end or year-end processing.

On the source Oracle Database, as DBA user enable Automatic SQL Tuning Set (ASTS):

1. On the database you want to migrate, enable ASTS.

```
BEGIN
  dbms_auto_task_admin.enable(
    client_name => 'Auto STS Capture Task',
    operation   => NULL,
    window_name => NULL);
END;
/
```

   See DBMS_AUTO_TASK_ADMIN for more information.

2. Verify that the ASTS background task is enabled.

```
SELECT task_name, interval,status, last_schedule_time, enabled
    FROM   dba_autotask_schedule_control
```

```
        WHERE  dbid = sys_context('userenv','con_dbid') AND
               task_name = 'Auto STS Capture Task';
```

If you want to monitor the SQL that's being captured, view `DBA_SQLSET_STATEMENTS`. For example:

```
SELECT substr(sql_text,1,100) txt, executions
    FROM dba_sqlset_statements
    WHERE sqlset_name = 'SYS_AUTO_STS';
```

As required, you can monitor the size and free space of `SYSAUX`. For example:

```
SELECT sum(bytes)/(1024*1024*1024) size_gb
    FROM   dba_data_files
    WHERE tablespace_name = 'SYSAUX' GROUP BY tablespace_name;

SELECT sum(bytes)/(1024*1024*1024) free_gb
    FROM dba_free_space
    WHERE tablespace_name = 'SYSAUX' GROUP BY tablespace_name;
```

Migrate Your Data to Autonomous Database

After you capture a sufficient amount of SQL on the source Oracle Database with Automatic SQL Tuning Set (ASTS) enabled, perform the migration to Autonomous Database.

See Migrate Oracle Databases to Autonomous Database for options for migrating to Autonomous Database.

Export Automatic SQL Tuning Set from Your Source Oracle Database

After you perform the migration to Autonomous Database, export the Automatic SQL Tuning Set (ASTS) from your source Oracle Database.

1. On the source Oracle Database, create and populate a staging table for ASTS data.

   As a DBA user, create the staging table:

   ```
   BEGIN
      dbms_sqlset.create_stgtab('ASTS_TABLE');
      dbms_sqlset.pack_stgtab('SYS_AUTO_STS','SYS','ASTS_TABLE');
   END;
   /
   ```

   After this operation completes the staging table contains the SQL statements that were captured on the source Oracle Database.

2. Export the staging table.

   For example, export the staging table using Oracle Data Pump:

   ```
   CREATE DIRECTORY dpdir AS '/export_directory';
   expdp user/password@database tables=asts_table directory=directory
   dumpfile=filename
   ```

See DBMS_SQLSET for more information.

Import Automatic SQL Tuning Set to Autonomous Database

After you perform the migration to Autonomous Database and you export the Automatic SQL Tuning Set (ASTS) from the source Oracle Database you are migrating, import the ASTS to your Autonomous Database.

1. Import the staging file to your Autonomous Database.

   Use Oracle Data Pump to import the staging table into Autonomous Database. First upload the dump file you exported from your source Oracle Database to a Cloud Object Storage bucket and then import the dump file.

   For example, as the ADMIN user run these commands:

   a. Create the credential to access the Cloud Object Store bucket.

   ```
   BEGIN
     DBMS_CLOUD.CREATE_CREDENTIAL(
       credential_name => 'BUCKET_CREDENTIAL',
       username => 'oracleidentitycloudservice/aaaaa@bbbbb.com',
       password => 'password'
     );
   END;
   /
   ```

   See Create Credentials to Access Cloud Services for more information.

   b. Use Oracle Data Pump to import the dump file with the ASTS data to your Autonomous Database instance.

   ```
   impdp admin/password@db_adb_high \
        directory=data_pump_dir \
        credential=BUCKET_CREDENTIAL \
        dumpfile= https://namespace-string.objectstorage.us-
   ashburn-1.oci.customer-oci.com/n/namespace-string/b/bucketname/o/
   asts_staging.dmp \
        tables=asts_table
   ```

   See Import Data Using Oracle Data Pump on Autonomous Database for more information.

> **Note:**
>
> Make sure that the timezone file in Autonomous Database matches the value in the source Oracle Database where you created the dump file. If there is a timezone mismatch, the database raises the following error:
>
> ```
> ORA-39002: invalid operation' error raised by
> dbms_datapump.start_job'
> ```
>
> See Datetime Data Types and Time Zone Support and Manage Time Zone File Updates on Autonomous Database for more information.

2. In the migrated Autonomous Database, unpack the SQL statements from the staging table and load them into the destination ASTS.

```
BEGIN
   dbms_sqlset.unpack_stgtab('SYS_AUTO_STS','SYS',TRUE,'ASTS_TABLE');
END;
/
```

See DBMS_SQLSET for more information.

Verify Your Real-time SPM Setting on Autonomous Database

Describes the steps to verify that real-time SPM is enabled on Autonomous Database.

Real-time SPM is enabled by default on Autonomous Database. You can verify the Real-time SPM mode as follows:

```
SELECT parameter_value spm_mode
    FROM   dba_sql_management_config
    WHERE  parameter_name = 'AUTO_SPM_EVOLVE_TASK';
```

The mode `AUTO` (automatic) real-time SPM indicates real-time SPM is enabled.

If real-time SPM is not enabled, use the following command to enable it:

```
EXEC dbms_spm.configure('AUTO_SPM_EVOLVE_TASK', 'AUTO')
```

Migrate Applications from Non-Oracle SQL to Oracle SQL

Application migration involves moving an application from one environment to another. The migration process includes a variety of tasks, including migrating SQL statements from a non-Oracle database to Oracle SQL.

The `DBMS_CLOUD_MIGRATION` package facilitates the translation of SQL statements written for a non-Oracle database to Oracle SQL. This allows you to easily migrate applications to Autonomous Database, when the applications were developed for non-Oracle databases that use a different SQL flavor.

The `DBMS_CLOUD_MIGRATION` package supports migration of SQL code from the following non-Oracle databases:

- PostgreSQL
- MySQL
- SQL Server

The `DBMS_CLOUD_MIGRATION` package provides routines that allow you to:

- Translate individual SQL statements from a non-Oracle database to Oracle SQL.
- Translate files containing SQL code from a non-Oracle database to Oracle SQL, and generate a new file containing the translated Oracle SQL.
- Work interactively in a session, entering SQL statements from a non-Oracle system. In this case, the non-Oracle statements are translated and run interactively in your session.

> **Note:**
>
> To run the subprograms in `DBMS_CLOUD_MIGRATION` you must be logged in as the `ADMIN` user or have `EXECUTE` privilege on `DBMS_CLOUD_MIGRATION`.

Topics

- Migrate Applications from PostgreSQL Databases to Autonomous Database
  You can migrate SQL statements from PostgreSQL to Oracle SQL and run the statements on Autonomous Database.
- Migrate Applications from MySQL Databases to Autonomous Database
  You can migrate SQL statements from MySQL to Oracle SQL and run the statements on Autonomous Database.
- Migrate Applications from SQL Server Databases to Autonomous Database
  You can migrate SQL statements from SQL Server to Oracle SQL and run the statements on Autonomous Database.

Migrate Applications from PostgreSQL Databases to Autonomous Database

You can migrate SQL statements from PostgreSQL to Oracle SQL and run the statements on Autonomous Database.

- Translate PostgreSQL Statements to Oracle SQL
  You can translate SQL statements written in PostgreSQL to Oracle SQL and run the translated statements on Autonomous Database.
- Run PostgreSQL Statements in Autonomous Database
  You can interactively translate and run PostgreSQL statements in your Autonomous Database.
- Migrate PostgreSQL Files to Oracle SQL
  You can migrate a file containing PostgreSQL statements to a file containing Oracle SQL statements.
- Limitations for Migration and Translation of PostgreSQL Statements to Oracle SQL
  This section summarizes the limitations for migrating SQL statements from PostgreSQL to Oracle SQL.

Translate PostgreSQL Statements to Oracle SQL

You can translate SQL statements written in PostgreSQL to Oracle SQL and run the translated statements on Autonomous Database.

Use `DBMS_CLOUD_MIGRATION.MIGRATE_SQL` to translate the PostgreSQL statement to Oracle SQL. There are procedure and function variants of `DBMS_CLOUD_MIGRATION.MIGRATE_SQL`.

Migrate PostgreSQL statement to Oracle SQL with MIGRATE_SQL Procedure

The following example accepts the SQL statement written in PostgreSQL as input, translates the statement to Oracle SQL, assigns the translated SQL statement to `output_sql_result`, and prints the result:

```
SET SERVEROUTPUT ON
   declare output_sql_result CLOB;
BEGIN
```

```
  DBMS_CLOUD_MIGRATION.MIGRATE_SQL(
    original_sql => 'SELECT e.employee_id, e.last_name, e.salary FROM
employees AS e;',
    output_sql   => output_sql_result,
    source_db    => 'POSTGRES');
    DBMS_OUTPUT.PUT_LINE (output_sql_result);
END;
/

Output
------------------------------------------------------------
SELECT e.employee_id, e.last_name, e.salary FROM employees e;
```

The `original_sql` parameter specifies the PostgreSQL statement.

The `output_sql` parameter stores the translated SQL.

The `source_db` parameter specifies the PostgreSQL database name.

See MIGRATE_SQL Procedure and Function for more information.

Migrate PostgreSQL statement to Oracle SQL with MIGRATE_SQL Function

The following example shows the `DBMS_CLOUD_MIGRATION.MIGRATE_SQL` function within a `SELECT` statement. The function input is a PostgreSQL statement and the function returns the translated statement in Oracle SQL:

```
SELECT DBMS_CLOUD_MIGRATION.MIGRATE_SQL(
     'CREATE TABLE IF NOT EXISTS cars (brand VARCHAR(255), model
VARCHAR(255), year INT)',
     'POSTGRES') AS output FROM DUAL;

OUTPUT
--------------------------------------------------------------------------------
create table cars (brand VARCHAR2(255), model VARCHAR2(255), year NUMBER(10);
```

See MIGRATE_SQL Procedure and Function for more information.

Notes for running `DBMS_CLOUD_MIGRATION.MIGRATE_SQL`:

- You may encounter an error during the translation if the input SQL statement is not supported in Oracle SQL. See Limitations for Migration and Translation of PostgreSQL **Statements** to Oracle SQL for more information.
- The `DBMS_CLOUD_MIGRATION.MIGRATE_SQL` subprograms only accept one SQL statement as input. So, only a single SQL statement can be translated per call.

Run PostgreSQL Statements in Autonomous Database

You can interactively translate and run PostgreSQL statements in your Autonomous Database.

Use the `ENABLE_TRANSLATION` procedure to enable real-time translation of SQL statements written in PostgreSQL. After you enable translation in a session, PostgreSQL statements are automatically translated and run as Oracle SQL statements, and you can see the results.

For example, after you enable translation by running `ENABLE_TRANSLATION` you can interactively do the following in a session:

- Create the tables. For example, create the tables `MOVIE` and `INVENTORY`.
- Insert data into tables.
- Query tables.
- Perform JOIN operations on tables. For example, you can perform a left outer join on tables.

To enable translation with PostgreSQL and run commands:

1. Connect to your Autonomous Database using a SQL client.

   See Connect to Autonomous Database for more information.

   > **Note:**
   >
   > `DBMS_CLOUD_MIGRATION.ENABLE_TRANSLATION` is not supported in Database Actions and is not supported with the Oracle APEX Service.

2. Run `DBMS_CLOUD_MIGRATION.ENABLE_TRANSLATION` to enable real-time SQL translation in your session.

   ```
   BEGIN
    DBMS_CLOUD_MIGRATION.ENABLE_TRANSLATION('POSTGRES');
   END;
   /
   ```

   Use the following query to verify the SQL translation language for your session:

   ```
   SELECT SYS_CONTEXT('USERENV','SQL_TRANSLATION_PROFILE_NAME') FROM DUAL;
   ```

   See ENABLE_TRANSLATION Procedure for more information.

3. Enter PostgreSQL statements. For example:

   ```
   CREATE TABLE movie (film_id int, title varchar(255));

   Table MOVIE created.
   ```

   This automatically translates and runs the PostgreSQL `CREATE TABLE` statement.

   You can verify using the `DESC` command. For example:

   ```
   DESC movie;

   Name        Null?     Type
   -------     -----     -------------
   FILM_ID               NUMBER(38)
   TITLE                 VARCHAR2(255)
   ```

   The `MOVIE` table is created and the datatypes of each of the columns are automatically converted to Oracle datatypes.

You may encounter an error during the translation if the input SQL statement is not supported. See Limitations for Migration and Translation of PostgreSQL **Statements** to Oracle SQL for more information.

4. Insert data into the `MOVIE` table. For example:

```
INSERT INTO movie (film_id, title) VALUES (123, 'Tangled');

1 row inserted.

INSERT INTO movie (film_id, title) VALUES (234, 'Frozen');

1 row inserted.
```

Verify the data insertion by querying the `MOVIE` table. For example:

```
SELECT * FROM movie;

FILM_ID  TITLE
-------- ---------
123      Tangled
234      Frozen
```

5. Create an `INVENTORY` table:

```
CREATE TABLE inventory (film_id int, inventory_id int);

Table INVENTORY created.
```

You can verify this step with the `DESC` command. For example:

```
DESC inventory;

Name         Null?     Type
-------      -----     -------------
FILM_ID                NUMBER(38)
INVENTORY_ID           NUMBER(38)
```

The `INVENTORY` table is created and the datatypes of each of the columns are automatically converted to Oracle datatypes.

6. Insert data into the `INVENTORY` table. For example:

```
INSERT INTO inventory(film_id, inventory_id) VALUES (123, 223);

1 row inserted.

INSERT INTO inventory(film_id, inventory_id) VALUES (234, 334);

1 row inserted.
```

Verify the data insertion by querying `INVENTORY`. For example:

```
SELECT * FROM inventory;

FILM_ID  INVENTORY_ID
-------- -------------
123      223
234      334
```

7. Perform a left outer join on tables `MOVIE` and `INVENTORY`:

```
SELECT m.film_id, m.title, inventory.inventory_id
    FROM movie AS m LEFT JOIN inventory
    ON inventory.film_id = m.film_id;

  FILM_ID  TITLE       INVENTORY_ID
---------- ---------- ------------
       234 Frozen      334
       123 Tangled     223
```

This example performs a `LEFT OUTER JOIN` on the `movie` and `inventory` tables. The `AS` keyword for table aliases in the `FROM` clause is not supported in Oracle SQL. The query is first translated to Oracle SQL and then run in your session.

8. Use `DBMS_CLOUD_MIGRATION.DISABLE_TRANSLATION` procedure to disable real-time SQL language translation for your session.

```
BEGIN
 DBMS_CLOUD_MIGRATION.DISABLE_TRANSLATION;
END;
/
```

This returns an error if SQL language translation is not enabled for your session.

See DISABLE_TRANSLATION Procedure

You can query the `V$MAPPED_SQL` view to list the PostgreSQL statements that are translated and mapped in memory to Oracle SQL statements.

For example:

```
SELECT v.*
    FROM v$mapped_sql v, dba_objects o
    WHERE v.sql_translation_profile_id = o.object_id
        AND o.object_name = 'POSTGRES'
        AND o.object_type = 'TRANSLATION PROFILE';
```

See V$MAPPED_SQL for more information.

Migrate PostgreSQL Files to Oracle SQL

You can migrate a file containing PostgreSQL statements to a file containing Oracle SQL statements.

The `DBMS_CLOUD_MIGRATION.MIGRATE_FILE` procedure translates SQL statements in a PostgreSQL file in Object Storage and generates a new file containing Oracle SQL.

As a prerequisite, upload one or more PostgreSQL files with a `.sql` extension to a location on Object Storage. The following examples use the file `postgrestest.sql` that is uploaded to Object Storage. See Put data into object storage for more information.

To migrate PostgreSQL files to Oracle SQL:

1. Connect to your Autonomous Database instance.

   See Connect to Autonomous Database for more information.

2. Configure access to Cloud Object Storage using a resource principal or by creating a credential object.

   This step provides access to the Cloud Object Storage where you put the files you are migrating:

   - **Use a resource principal to access Cloud Object Storage**:

     Define policies and roles to access Oracle Cloud Infrastructure Object Storage with a resource principal. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources and Use Resource Principal with DBMS_CLOUD for more information.

     If your files reside in another vendor's Cloud Object Storage, define the appropriate policies and roles to access the resource. See Configure Policies and Roles to Access Resources for more information.

   - **Create a credential to access Cloud Object Storage:**

     You can create a credential to access Cloud Object Storage.

     For example:

     ```
     BEGIN
       DBMS_CLOUD.CREATE_CREDENTIAL (
         credential_name => 'OBJ_STORE_CRED',
         username        => 'user1@example.com',
         password        => 'password'
       );
     END;
     /
     ```

     See CREATE_CREDENTIAL Procedure for more information.

3. Optionally, you can list the files in Object Storage. For example:

   ```
   VAR function_list CLOB;
   SELECT object_name FROM DBMS_CLOUD.LIST_OBJECTS
      (credential_name => 'OCI$RESOURCE_PRINCIPAL', location_uri => 'https://
   objectstorage.region.oraclecloud.com/n/namespace/b/bucket/o/files');

   OBJECT_NAME
   ----------------------
   postgrestest.sql
   ```

   Alternatively, if you create a credential instead of the resource principal, `OCI$RESOURCE_PRINCIPAL`, specify the credential name in the `credential_name` parameter.

   See LIST_OBJECTS Function for more information.

4. Run `DBMS_CLOUD_MIGRATION.MIGRATE_FILE` to migrate the PostgreSQL file to Oracle SQL:

```
BEGIN
 DBMS_CLOUD_MIGRATION.MIGRATE_FILE (
     credential_name => 'OCI$RESOURCE_PRINCIPAL',
     location_uri    => 'https://objectstorage.region.oraclecloud.com/n/
namespace/b/bucket/o/files/postgrestest.sql',
     source_db       => 'POSTGRES'
    );
END;
/
```

The `credential_name` parameter specifies the credential to access the Cloud Object Storage URI. The user that runs `DBMS_CLOUD_MIGRATION.MIGRATE_FILE` must have `EXECUTE` privilege for the credential object used to access the Object Storage URI. That is, the credential you specify with the `credential_name` parameter. If you use a credential instead of a resource principal, specify the credential name in the `credential_name` parameter.

The `location_uri` parameter specifies the source file URI. The format of the URI depends on the Cloud Object Storage service you are using. See DBMS_CLOUD URI Formats for more information.

The `source_db` parameter specifies the PostgreSQL as database language. Use the value `POSTGRES` to translate PostgreSQL files to Oracle SQL.

Running this command translates the PostgreSQL file `postgrestest.sql` to Oracle SQL and generates a new file with the name *original_filename*`_oracle.sql`.

For this example, running `DBMS_CLOUD_MIGRATION.MIGRATE_FILE` with the input file `postgrestest.sql` generates `postgrestest_oracle.sql`. After the translation step the procedure uploads `postgrestest_oracle.sql` to Object Storage.

Optionally, use the `target_uri` parameter to specify the location where the translated file is uploaded. The default value for this parameter is `NULL`, which means the translated file is uploaded to the same location as specified in the `location_uri` parameter.

See MIGRATE_FILE Procedure for more information.

5. Verify that the output file was generated.

```
SELECT object_name FROM DBMS_CLOUD.LIST_OBJECTS (
    credential_name => 'OCI$RESOURCE_PRINCIPAL', location_uri => 'https://
objectstorage.region.oraclecloud.com/n/namespace/b/bucket/o/files');

OBJECT_NAME
--------------------
postgrestest.sql
postgrestest_oracle.sql
```

Alternatively, if you create a credential instead of the resource principal, `OCI$RESOURCE_PRINCIPAL`, specify the credential name in the `credential_name` parameter.

See LIST_OBJECTS Function for more information.

Run the following query to view the content of the `postgrestest_oracle.sql` file:

```
SELECT UTL_RAW.CAST_TO_VARCHAR2 (DBMS_CLOUD.GET_OBJECT(
   credential_name => 'OCI$RESOURCE_PRINCIPAL',
   object_uri => 'https://objectstorage.region.oraclecloud.com/n/namespace/b/
bucket/o/files'))
FROM dual;

UTL_RAW.CAST_TO_VARCHAR2(DBMS_CLOUD.GET_OBJECT(CREDENTIAL_NAME=>'CRED1',OBJECT
_U
--------------------------------------------------------------------------------
--
SELECT f.film_id, f.title, inventory_id
FROM film f LEFT JOIN inventory
ON inventory.film_id = f.film_id;
```

Alternatively, if you create a credential instead of the resource principal, `OCI$RESOURCE_PRINCIPAL`, specify the credential name in the `credential_name` parameter.

See GET_OBJECT Procedure and Function for more information.

Limitations for Migration and Translation of PostgreSQL **Statements** to Oracle SQL

This section summarizes the limitations for migrating SQL statements from PostgreSQL to Oracle SQL.

The following list of PostgreSQL statements are not supported when migrating to Autonomous Database:

- `CREATE DOMAIN`
- `CREATE EXTENSION`
- `CREATE DATABASE`
- `CREATE TYPE`
- `SET`

The following list of PostgreSQL statements are supported with restrictions:

- `ALTER TABLE`: Only `ALTER TABLE ADD CONSTRAINT` is supported when migrating to Autonomous Database.
- `DELETE`: The `RETURNING *` keyword in the `DELETE` statement is not supported in Autonomous Database. You must replace the `RETURNING *` clause with the `RETURNING INTO` clause. For example, `DELETE FROM tasks WHERE status = 'DONE' RETURNING *;`. See RETURNING INTO Clause for more information.
- `CREATE FUNCTION`: Following are not supported with `CREATE FUNCTION`:
  - The `SETOF` return type, you must replace `SETOF` with `CURSORS` or `COLLECTIONS` return types.
  - The `IMMUTABLE` clause.
  - The parameter declarations in the format `FUNCTION_NAME (DATATYPE, DATATYPE)`.
- `ALTER FUNCTION`: The `ALTER FUNCTION` function arguments, for example, `RENAME TO`, `OWNER TO`, `SET SCHEMA` are not supported.

Migrate Applications from MySQL Databases to Autonomous Database

You can migrate SQL statements from MySQL to Oracle SQL and run the statements on Autonomous Database.

- Translate MySQL Statements to Oracle SQL
  You can translate SQL statements written in MySQL to Oracle SQL and run the translated statements on Autonomous Database.
- Run MySQL Statements in Autonomous Database
  You can interactively translate and run MySQL statements in your Autonomous Database.
- Migrate MySQL Files to Oracle SQL
  You can migrate a file containing MySQL statements to a file containing Oracle SQL statements.
- Limitations for Migration and Translation of MySQL Statements to Oracle SQL
  This section summarizes the limitations for migrating SQL statements from MySQL to Oracle SQL.

Translate MySQL Statements to Oracle SQL

You can translate SQL statements written in MySQL to Oracle SQL and run the translated statements on Autonomous Database.

Use `DBMS_CLOUD_MIGRATION.MIGRATE_SQL` to translate a MySQL statement to Oracle SQL. There are procedure and function variants of `DBMS_CLOUD_MIGRATION.MIGRATE_SQL`.

Migrate MySQL Statement to Oracle SQL with MIGRATE_SQL Procedure

The following example accepts the SQL statement written in MySQL as input, translates the statement to Oracle SQL, assigns the translated SQL statement to `output_sql_result`, and prints the result:

```
SET SERVEROUTPUT ON
   declare output_sql_result CLOB;
BEGIN
  DBMS_CLOUD_MIGRATION.MIGRATE_SQL(
    original_sql => 'CREATE TABLE movie (movie_id INT, title VARCHAR(255));',
    output_sql   => output_sql_result,
    source_db    => 'MYSQL');
    DBMS_OUTPUT.PUT_LINE (output_sql_result);
END;
/

OUTPUT
-------------------------------------------------------------
CREATE TABLE movie (movie_id NUMBER(10), title VARCHAR2(255));
```

The `original_sql` parameter specifies the MySQL statement.

The `output_sql` parameter stores the translated SQL.

The `source_db` parameter specifies MySQL as the database name.

See MIGRATE_SQL Procedure and Function for more information.

Migrate MySQL Statement to Oracle SQL with MIGRATE_SQL Function

The following example shows the `DBMS_CLOUD_MIGRATION.MIGRATE_SQL` function within a `SELECT` statement. The function input is a MySQL statement and the function returns the translated statement in Oracle SQL:

```
SELECT DBMS_CLOUD_MIGRATION.MIGRATE_SQL(
     'CREATE TABLE movie (film_id INT, title VARCHAR(255));','MYSQL') AS
output FROM DUAL;

OUTPUT
--------------------------------------------------------------------------------
create table cars (brand VARCHAR2(255), model VARCHAR2(255);
```

See MIGRATE_SQL Procedure and Function for more information.

Notes for running `DBMS_CLOUD_MIGRATION.MIGRATE_SQL`:

- You may encounter an error during the translation if the input SQL statement is not supported in Oracle SQL. See Limitations for Migration and Translation of MySQL Statements to Oracle SQL for more information.
- The `DBMS_CLOUD_MIGRATION.MIGRATE_SQL` subprograms only accept one SQL statement as input. So, only a single SQL statement can be translated per call.

Run MySQL Statements in Autonomous Database

You can interactively translate and run MySQL statements in your Autonomous Database.

Use the `ENABLE_TRANSLATION` procedure to enable real-time translation of SQL statements written in MySQL. After you enable translation in a session, MySQL statements are automatically translated and run as Oracle SQL statements, and you can see the results.

For example, after you enable translation by running `ENABLE_TRANSLATION` you can interactively do the following in a session:

- Create the tables. For example, create the tables `MOVIE` and `INVENTORY`.
- Insert data into tables.
- Query tables.
- Perform JOIN operations on tables. For example, you can perform a left outer join on tables.

To enable translation with MySQL and run commands:

1. Connect to your Autonomous Database using a SQL client.

   See Connect to Autonomous Database for more information.

> **Note:**
>
> `DBMS_CLOUD_MIGRATION.ENABLE_TRANSLATION` is not supported in Database Actions and is not supported with the Oracle APEX Service.

2. Run `DBMS_CLOUD_MIGRATION.ENABLE_TRANSLATION` to enable real-time SQL translation in your session. Use the `MYSQL` parameter to translate from MySQL.

```
BEGIN
 DBMS_CLOUD_MIGRATION.ENABLE_TRANSLATION('MYSQL');
END;
/
```

Use the following query to verify the SQL translation language for your session:

```
SELECT SYS_CONTEXT('USERENV','SQL_TRANSLATION_PROFILE_NAME') FROM DUAL;
```

See ENABLE_TRANSLATION Procedure for more information.

3. Enter MySQL statements. For example:

```
CREATE TABLE movie (film_id int, title varchar(255));

Table MOVIE created.
```

This automatically translates and runs the MySQL `CREATE TABLE` statement.

You can verify using the `DESC` command. For example:

```
DESC movie;

Name        Null?     Type
-------     -----     -------------
FILM_ID               NUMBER(38)
TITLE                 VARCHAR2(255)
```

The `MOVIE` table is created and the datatypes of each of the columns are automatically converted to Oracle datatypes.

You may encounter an error during the translation if the input SQL statement is not supported. See Limitations for Migration and Translation of MySQL Statements to Oracle SQL for more information.

4. Insert data into the `MOVIE` table. For example:

```
INSERT INTO movie (film_id, title) VALUES (123, 'Tangled');

1 row inserted.

INSERT INTO movie (film_id, title) VALUES (234, 'Frozen');

1 row inserted.
```

Verify the data insertion by querying the `MOVIE` table. For example:

```
SELECT * FROM movie;

FILM_ID  TITLE
-------- ---------
```

```
123     Tangled
234     Frozen
```

5. Create an `INVENTORY` table:

```
CREATE TABLE inventory (film_id int, inventory_id int);

Table INVENTORY created.
```

You can verify this step with the `DESC` command. For example:

```
DESC inventory;

Name         Null?     Type
-------      -----     -------------
FILM_ID                NUMBER(38)
INVENTORY_ID           NUMBER(38)
```

The `INVENTORY` table is created and the datatypes of each of the columns are automatically converted to Oracle datatypes.

6. Insert data into the `INVENTORY` table. For example:

```
INSERT INTO inventory(film_id, inventory_id) VALUES (123, 223);

1 row inserted.

INSERT INTO inventory(film_id, inventory_id) VALUES (234, 334);

1 row inserted.
```

Verify the data insertion by querying `INVENTORY`. For example:

```
SELECT * FROM inventory;

FILM_ID  INVENTORY_ID
-------- -------------
123     223
234     334
```

7. Perform a left outer join on tables `MOVIE` and `INVENTORY`:

```
SELECT m.film_id, m.title, inventory_id
    FROM movie AS m LEFT JOIN inventory
    ON inventory.film_id = m.film_id;

  FILM_ID  TITLE       INVENTORY_ID
---------- ---------- ------------
       234 Frozen      334
       123 Tangled     223
```

This example performs a `LEFT OUTER JOIN` on the `movie` and `inventory` tables. The `AS` keyword for table aliases in the `FROM` clause is not supported in Oracle SQL. The query is first translated to Oracle SQL and then run in your session.

8. Use `DBMS_CLOUD_MIGRATION.DISABLE_TRANSLATION` procedure to disable real-time SQL language translation for your session.

```
BEGIN
 DBMS_CLOUD_MIGRATION.DISABLE_TRANSLATION;
END;
/
```

This returns an error if SQL language translation is not enabled for your session.

See DISABLE_TRANSLATION Procedure

You can query the `V$MAPPED_SQL` view to list the MySQL statements that are translated and mapped in memory to Oracle SQL statements.

For example:

```
SELECT v.*
    FROM v$mapped_sql v, dba_objects o
    WHERE v.sql_translation_profile_id = o.object_id
        AND o.object_name = 'MYSQL'
        AND o.object_type = 'TRANSLATION PROFILE';
```

See V$MAPPED_SQL for more information.

Migrate MySQL Files to Oracle SQL

You can migrate a file containing MySQL statements to a file containing Oracle SQL statements.

The `DBMS_CLOUD_MIGRATION.MIGRATE_FILE` procedure translates SQL statements in a MySQL file in Object Storage and generates a new file containing Oracle SQL.

As a prerequisite, upload one or more MySQL files with a `.sql` extension to a location on Object Storage. The following examples use the file `mysqltest.sql` that is uploaded to Object Storage. See Put data into object storage for more information.

To migrate MySQL files to Oracle SQL:

1. Connect to your Autonomous Database instance.

   See Connect to Autonomous Database for more information.

2. Configure access to Cloud Object Storage using a resource principal or by creating a credential object.

   This step provides access to the Cloud Object Storage where you put the files you are migrating:

   - **Use a resource principal to access Cloud Object Storage**:

     Define policies and roles to access Oracle Cloud Infrastructure Object Storage with a resource principal. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources and Use Resource Principal with DBMS_CLOUD for more information.

If your files reside in another vendor's Cloud Object Storage, define the appropriate policies and roles to access the resource. See Configure Policies and Roles to Access Resources for more information.

- **Create a credential to access Cloud Object Storage:**

  You can create a credential to access Cloud Object Storage.

  For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL (
    credential_name => 'OBJ_STORE_CRED',
    username        => 'user1@example.com',
    password        => 'password'
  );
END;
/
```

  See CREATE_CREDENTIAL Procedure for more information.

3. Optionally, you can list the files in Object Storage. For example:

```
VAR function_list CLOB;
SELECT object_name FROM DBMS_CLOUD.LIST_OBJECTS
   (credential_name => 'OCI$RESOURCE_PRINCIPAL', location_uri => 'https://
objectstorage.region.oraclecloud.com/n/namespace/b/bucket/o/files');

OBJECT_NAME
--------------------
mysqltest.sql
```

   Alternatively, if you create a credential instead of the resource principal, `OCI$RESOURCE_PRINCIPAL`, specify the credential name in the `credential_name` parameter.

   See LIST_OBJECTS Function for more information.

4. Run `DBMS_CLOUD_MIGRATION.MIGRATE_FILE` to migrate the MySQL file to Oracle SQL:

```
BEGIN
 DBMS_CLOUD_MIGRATION.MIGRATE_FILE (
     credential_name => 'OCI$RESOURCE_PRINCIPAL',
     location_uri    => 'https://objectstorage.region.oraclecloud.com/n/
namespace/b/bucket/o/files/mysqltest.sql',
     source_db       => 'MYSQL'
    );
END;
/
```

   The `credential_name` parameter specifies the credential to access the Cloud Object Storage URI. The user that runs `DBMS_CLOUD_MIGRATION.MIGRATE_FILE` must have `EXECUTE` privilege for the credential object used to access the Object Storage URI. That is, the credential you specify with the `credential_name` parameter. If you use a credential instead of a resource principal, specify the credential name in the `credential_name` parameter.

The `location_uri` parameter specifies the source file URI. The format of the URI depends on the Cloud Object Storage service you are using. See DBMS_CLOUD URI Formats for more information.

The `source_db` parameter specifies MySQL as the database language. Use the value `MYSQL` to translate MySQL files to Oracle SQL.

Running this command translates the MySQL file `mysqltest.sql` to Oracle SQL and generates a new file with the name *original_filename*`_oracle.sql`.

For this example, running `DBMS_CLOUD_MIGRATION.MIGRATE_FILE` with the input file `mysqltest.sql` generates `mysqltest_oracle.sql`. After the translation step the procedure uploads `mysqltest_oracle.sql` to Object Storage.

Optionally, use the `target_uri` parameter to specify the location where the translated file is uploaded. The default value for this parameter is `NULL`, which means the translated file is uploaded to the same location as specified in the `location_uri` parameter.

See MIGRATE_FILE Procedure for more information.

5. Verify that the output file was generated.

```
SELECT object_name FROM DBMS_CLOUD.LIST_OBJECTS (
    credential_name => 'OCI$RESOURCE_PRINCIPAL', location_uri => 'https://
objectstorage.region.oraclecloud.com/n/namespace/b/bucket/o/files');

OBJECT_NAME
--------------------
mysqltest.sql
mysqltest_oracle.sql
```

Alternatively, if you create a credential instead of the resource principal, `OCI$RESOURCE_PRINCIPAL`, specify the credential name in the `credential_name` parameter.

See LIST_OBJECTS Function for more information.

Run the following query to view the content of the `mysqltest_oracle.sql` file:

```
SELECT UTL_RAW.CAST_TO_VARCHAR2 (DBMS_CLOUD.GET_OBJECT(
   credential_name => 'OCI$RESOURCE_PRINCIPAL',
   object_uri => 'https://objectstorage.region.oraclecloud.com/n/namespace/b/
bucket/o/files'))
FROM dual;

UTL_RAW.CAST_TO_VARCHAR2(DBMS_CLOUD.GET_OBJECT(CREDENTIAL_NAME=>'CRED1',OBJECT
_U
--------------------------------------------------------------------------------
--
DROP TABLE movie;
DROP TABLE inventory;
CREATE TABLE movie (film_id NUMBER(10), title VARCHAR2(255));
INSERT INTO movie (film_id, title) VALUES (123, 'Tangled');
INSERT INTO movie (film_id, title) VALUES (234, 'Frozen');
CREATE TABLE movie (film_id NUMBER(10), inventory_id NUMBER(10));
INSERT INTO movie (film_id, inventory_id) VALUES (123, 223);
INSERT INTO inventory (film_id, inventory_id) VALUES (234, 334);
SELECT * FROM movie;
```

Alternatively, if you create a credential instead of the resource principal, `OCI$RESOURCE_PRINCIPAL`, specify the credential name in the `credential_name` parameter.

See GET_OBJECT Procedure and Function for more information.

Limitations for Migration and Translation of MySQL Statements to Oracle SQL

This section summarizes the limitations for migrating SQL statements from MySQL to Oracle SQL.

Note the following restrictions when migrating from MySQL database flavor to Oracle SQL:

- User-Defined Global Variables: MySQL user-defined global variables declared using `@var_name` format are not supported when translating to Oracle SQL.
- Quoted Identifiers: MySQL identifiers that are enclosed with backticks (`), such as `` `var_name` ``, are not supported when translating to Oracle SQL. Only identifiers that are without quotes or are enclosed in double quotation marks are supported.
- Delimiters:
  - For functions and procedures, only `$$` and `//` delimiters are supported.
  - For SQL statements only `;` delimiter is supported.

The following list of MySQL functions are not supported:

- `ADDTIME(datetime, time)`
- `AES_DECRYPT(crypt_str, key_str)`
- `AES_ENCRYPT(str, key_str)`
- `BIN(num)`
- `CONV(num, from_base, to_base)`
- `CONVERT_TZ(datetime, from, to)`
- `CRC32(exp)`
- `DATE_ADD(date, interval)`

Migrate Applications from SQL Server Databases to Autonomous Database

You can migrate SQL statements from SQL Server to Oracle SQL and run the statements on Autonomous Database.

- Translate SQL Server Statements to Oracle SQL
  You can translate SQL statements written in SQL Server to Oracle SQL and run the translated statements on Autonomous Database.
- Run SQL Server Statements in Autonomous Database
  You can interactively translate and run SQL Server statements in your Autonomous Database.
- Migrate SQL Server Files to Oracle SQL
  You can migrate a file containing SQL Server statements to a file containing Oracle SQL statements.
- Limitations for Migration and Translation of SQL Server Statements to Oracle SQL
  This section summarizes the limitations for migrating SQL statements from SQL Server to Oracle SQL.

Translate SQL Server Statements to Oracle SQL

You can translate SQL statements written in SQL Server to Oracle SQL and run the translated statements on Autonomous Database.

Use `DBMS_CLOUD_MIGRATION.MIGRATE_SQL` to translate a SQL Server statement to Oracle SQL. There are procedure and function variants of `DBMS_CLOUD_MIGRATION.MIGRATE_SQL`.

Migrate SQL Server Statement to Oracle SQL with MIGRATE_SQL Procedure

The following example accepts the SQL statement written in SQL Server as input, translates the statement to Oracle SQL, assigns the translated SQL statement to `output_sql_result`, and prints the result:

```
SET SERVEROUTPUT ON
   declare output_sql_result CLOB;
BEGIN
  DBMS_CLOUD_MIGRATION.MIGRATE_SQL(
    original_sql => 'CREATE TABLE [dbo].[movie] ([film_id] [int], [title]
[varchar](20));',
    output_sql   => output_sql_result,
    source_db    => 'SQLSERVER');
    DBMS_OUTPUT.PUT_LINE (output_sql_result);
END;
/

Output
----------------------------------------------------------------
CREATE TABLE movie (film_id NUMBER(10), title VARCHAR2(20));
```

The `original_sql` parameter specifies the SQL Server SQL statement.

The `output_sql` parameter stores the translated SQL.

The `source_db` parameter specifies the SQL Server database.

See MIGRATE_SQL Procedure and Function for more information.

Migrate SQL Server Statement to Oracle SQL with MIGRATE_SQL Function

The following example shows the `DBMS_CLOUD_MIGRATION.MIGRATE_SQL` function within a `SELECT` statement. The function input is a SQL Server statement and the function returns the translated statement in Oracle SQL:

```
SELECT DBMS_CLOUD_MIGRATION.MIGRATE_SQL(
     'CREATE TABLE [dbo].[movie] ([film_id] [int], [title] [varchar]
(20));','SQLSERVER') AS OUTPUT
     FROM DUAL;

OUTPUT
--------------------------------------------------------------------------------
CREATE TABLE movie (film_id NUMBER(10), title VARCHAR2(20));
```

Use the `SQLSERVER` parameter to translate from Microsoft SQL Server.

See MIGRATE_SQL Procedure and Function for more information.

Notes for running DBMS_CLOUD_MIGRATION.MIGRATE_SQL:

- You may encounter an error during the translation if the input SQL statement is not supported in Oracle SQL. See Limitations for Migration and Translation of SQL Server Statements to Oracle SQL for more information.
- The DBMS_CLOUD_MIGRATION.MIGRATE_SQL subprograms only accept one SQL statement as input. So, only a single SQL statement can be translated per call.

Run SQL Server Statements in Autonomous Database

You can interactively translate and run SQL Server statements in your Autonomous Database.

Use the ENABLE_TRANSLATION procedure to enable real-time translation of SQL statements written in SQL Server. After you enable translation in a session, non-Oracle SQL statements are automatically translated and run as Oracle SQL statements, and you can see the results.

For example, after you enable translation by running ENABLE_TRANSLATION you can interactively do the following in a session:

- Create the tables. For example, create the tables MOVIE and INVENTORY.
- Insert data into tables.
- Query tables.
- Perform JOIN operations on tables. For example, you can perform a left outer join on tables.

To enable translation with SQL Server statements and run commands:

1. Connect to your Autonomous Database using a SQL client.

   See Connect to Autonomous Database for more information.

   > **Note:**
   >
   > DBMS_CLOUD_MIGRATION.ENABLE_TRANSLATION is not supported in Database Actions and is not supported with the Oracle APEX Service.

2. Run DBMS_CLOUD_MIGRATION.ENABLE_TRANSLATION to enable real-time SQL translation in your session. Use the SQLSERVER parameter to translate from Microsoft SQL Server.

   ```
   BEGIN
    DBMS_CLOUD_MIGRATION.ENABLE_TRANSLATION('SQLSERVER');
   END;
   /
   ```

   Use the following query to verify the SQL translation language for your session:

   ```
   SELECT SYS_CONTEXT('USERENV','SQL_TRANSLATION_PROFILE_NAME') FROM DUAL;
   ```

   See ENABLE_TRANSLATION Procedure for more information.

3. Enter SQL Server SQL statements. For example:

```
CREATE TABLE [dbo].[movie] ([film_id] [int], [title] [varchar](20));

Table [dbo].[movie] created.
```

This automatically translates and runs the SQL Server `CREATE TABLE` statement.

You can verify using the `DESC` command. For example:

```
DESC movie;

Name        Null?     Type
-------     -----     -------------
FILM_ID               NUMBER(38)
TITLE                 VARCHAR2(255)
```

The `MOVIE` table is created and the datatypes of each of the columns are automatically converted to Oracle datatypes.

4. Insert data into the `MOVIE` table. For example:

```
INSERT INTO movie([film_id], [title]) VALUES (123, 'Tangled');

1 row inserted.

INSERT INTO movie([film_id], [title]) VALUES (234, 'Frozen');

1 row inserted.
```

Verify the data insertion by querying the `MOVIE` table. For example:

```
SELECT * FROM movie;

FILM_ID  TITLE
-------- ---------
123      Tangled
234      Frozen
```

5. Create an `INVENTORY` table:

```
CREATE TABLE [dbo].[inventory] ([film_id] [int], [title] [inventory_id]
(20));


Table INVENTORY created.
```

You can verify this step with the `DESC` command. For example:

```
DESC inventory;

Name        Null?     Type
-------     -----     -------------
```

```
FILM_ID                NUMBER(38)
INVENTORY_ID           NUMBER(38)
```

The `INVENTORY` table is created and the datatypes of each of the columns are automatically converted to Oracle datatypes.

6. Insert data into the `INVENTORY` table. For example:

```
INSERT INTO inventory([film_id], [inventory_id]) VALUES (123, 223);

1 row inserted.

INSERT INTO inventory(film_id, inventory_id) VALUES (234, 334);

1 row inserted.
```

Verify the data insertion by querying `INVENTORY`. For example:

```
SELECT * FROM inventory;

FILM_ID  INVENTORY_ID
-------- -------------
123      223
234      334
```

7. Perform a left outer join on tables `MOVIE` and `INVENTORY`:

```
SELECT m.film_id, m.title, inventory.inventory_id
    FROM movie AS m LEFT JOIN inventory
    ON inventory.film_id = m.film_id;

  FILM_ID  TITLE       INVENTORY_ID
---------- ---------- ------------
       234 Frozen      334
       123 Tangled     223
```

This example performs a `LEFT OUTER JOIN` on the `movie` and `inventory` tables. The query is first translated to Oracle SQL and then run in your session.

8. Use `DBMS_CLOUD_MIGRATION.DISABLE_TRANSLATION` procedure to disable real-time SQL language translation for your session.

```
BEGIN
 DBMS_CLOUD_MIGRATION.DISABLE_TRANSLATION;
END;
/
```

This returns an error if SQL language translation is not enabled for your session.

See DISABLE_TRANSLATION Procedure

You can query the `V$MAPPED_SQL` view to list the SQL Server statements that are translated and mapped in memory to Oracle SQL statements.

For example:

```
SELECT v.*
    FROM v$mapped_sql v, dba_objects o
    WHERE v.sql_translation_profile_id = o.object_id
        AND o.object_name = 'SQLSERVER'
        AND o.object_type = 'TRANSLATION PROFILE';
```

See V$MAPPED_SQL for more information.

Migrate SQL Server Files to Oracle SQL

You can migrate a file containing SQL Server statements to a file containing Oracle SQL statements.

The `DBMS_CLOUD_MIGRATION.MIGRATE_FILE` procedure translates SQL statements in a SQL Server file in Object Storage and generates a new file containing Oracle SQL.

As a prerequisite, upload one or more SQL Server files with a `.sql` extension to a location on Object Storage. The following examples use the file `mssqltest.sql` that is uploaded to Object Storage. See Put data into object storage for more information.

To migrate SQL Server files to Oracle SQL:

1. Connect to your Autonomous Database instance.

   See Connect to Autonomous Database for more information.

2. Configure access to Cloud Object Storage using a resource principal or by creating a credential object.

   This step provides access to the Cloud Object Storage where you put the files you are migrating:

   - **Use a resource principal to access Cloud Object Storage**:

     Define policies and roles to access Oracle Cloud Infrastructure Object Storage with a resource principal. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources and Use Resource Principal with DBMS_CLOUD for more information.

     If your files reside in another vendor's Cloud Object Storage, define the appropriate policies and roles to access the resource. See Configure Policies and Roles to Access Resources for more information.

   - **Create a credential to access Cloud Object Storage:**

     You can create a credential to access Cloud Object Storage.

     For example:

     ```
     BEGIN
       DBMS_CLOUD.CREATE_CREDENTIAL (
         credential_name => 'OBJ_STORE_CRED',
         username        => 'user1@example.com',
         password        => 'password'
       );
     END;
     /
     ```

     See CREATE_CREDENTIAL Procedure for more information.

3. Optionally, you can list the files in Object Storage. For example:

```
VAR function_list CLOB;
SELECT object_name FROM DBMS_CLOUD.LIST_OBJECTS
   (credential_name => 'OCI$RESOURCE_PRINCIPAL', location_uri => 'https://
objectstorage.region.oraclecloud.com/n/namespace/b/bucket/o/files');

OBJECT_NAME
--------------------
mssqltest.sql
```

Alternatively, if you create a credential instead of the resource principal, `OCI$RESOURCE_PRINCIPAL`, specify the credential name in the `credential_name` parameter.

See LIST_OBJECTS Function for more information.

4. Run `DBMS_CLOUD_MIGRATION.MIGRATE_FILE` to migrate the SQL Server file to Oracle SQL:

```
BEGIN
 DBMS_CLOUD_MIGRATION.MIGRATE_FILE (
     credential_name => 'OCI$RESOURCE_PRINCIPAL',
     location_uri    => 'https://objectstorage.region.oraclecloud.com/n/
namespace/b/bucket/o/files/mssqltest.sql',
     source_db       => 'SQLSERVER'
    );
END;
/
```

The `credential_name` parameter specifies the credential to access the Cloud Object Storage URI. The user that runs `DBMS_CLOUD_MIGRATION.MIGRATE_FILE` must have `EXECUTE` privilege for the credential object used to access the Object Storage URI. That is, the credential you specify with the `credential_name` parameter. If you use a credential instead of a resource principal, specify the credential name in the `credential_name` parameter.

The `location_uri` parameter specifies the source file URI. The format of the URI depends on the Cloud Object Storage service you are using. See DBMS_CLOUD URI Formats for more information.

The `source_db` parameter specifies SQL Server as the database language. Use the value `SQLSERVER` to translate SQL Server files to Oracle SQL.

Running this command translates the SQL Server file `mssqltest.sql` to Oracle SQL and generates a new file with the name *original_filename*`_oracle.sql`.

For this example, running `DBMS_CLOUD_MIGRATION.MIGRATE_FILE` with the input file `mssqltest.sql` generates `mssqltest_oracle.sql`. After the translation step the procedure uploads `mssqltest_oracle.sql` to Object Storage.

Optionally, use the `target_uri` parameter to specify the location where the translated file is uploaded. The default value for this parameter is `NULL`, which means the translated file is uploaded to the same location as specified in the `location_uri` parameter.

See MIGRATE_FILE Procedure for more information.

5. Verify that the output file was generated.

```
SELECT object_name FROM DBMS_CLOUD.LIST_OBJECTS (
    credential_name => 'OCI$RESOURCE_PRINCIPAL', location_uri => 'https://
objectstorage.region.oraclecloud.com/n/namespace/b/bucket/o/files');

OBJECT_NAME
--------------------
mssqltest.sql
mssqltest_oracle.sql
```

Alternatively, if you create a credential instead of the resource principal, `OCI$RESOURCE_PRINCIPAL`, specify the credential name in the `credential_name` parameter.

See LIST_OBJECTS Function for more information.

Run the following query to view the content of the `mssqltest_oracle.sql` file:

```
SELECT UTL_RAW.CAST_TO_VARCHAR2 (DBMS_CLOUD.GET_OBJECT(
   credential_name => 'OCI$RESOURCE_PRINCIPAL',
   object_uri => 'https://objectstorage.region.oraclecloud.com/n/namespace/b/
bucket/o/files'))
FROM dual;

UTL_RAW.CAST_TO_VARCHAR2(DBMS_CLOUD.GET_OBJECT(CREDENTIAL_NAME=>'CRED1',OBJECT
_U
--------------------------------------------------------------------------------
--
DROP TABLE movie;
CREATE TABLE movie (film_id NUMBER(10), title VARCHAR2(20));

INSERT INTO movie (film_id, title) VALUES (123, 'Tangled');
INSERT INTO movie (film_id, title) VALUES (234, 'Frozen');

CREATE TABLE inventory(film_id NUMBER(10), inventory_id NUMBER(10));
INSERT INTO inventory(film_id, inventory_id) VALUES (123, 223);
INSERT INTO inventory(film_id, inventory_id) VALUES (234, 334);

SELECT * FROM movie;
SELECT * FROM inventory;
SELECT m.film_id, m.title, inventory_id FROM movie m LEFT JOIN inventory ON
inventory.film_id = m.film_id;
```

Alternatively, if you create a credential instead of the resource principal, `OCI$RESOURCE_PRINCIPAL`, specify the credential name in the `credential_name` parameter.

See GET_OBJECT Procedure and Function for more information.

Limitations for Migration and Translation of SQL Server Statements to Oracle SQL

This section summarizes the limitations for migrating SQL statements from SQL Server to Oracle SQL.

The following list of SQL Server functions are not supported when migrating from SQL Server database flavor to Oracle SQL:

- `ISJSON()`

- `OPENJSON()`

Use Database In-Memory with Autonomous Database

The Database In-Memory feature available on Autonomous Database can significantly improve performance for real-time analytics and mixed workloads.

- About Database In-Memory
  Oracle Database In-Memory (Database In-Memory) is a suite of features that greatly improves performance for real-time analytics and mixed workloads.
- Enabling Objects for In-Memory Population
  You can enable and disable objects for population in the In-Memory column store.
- Sizing the In-Memory Column Store
  When your Autonomous Database instance has 16 or more ECPUs, the Database In-Memory feature is enabled by default, and up to 50% of the SGA can be used by In-Memory.
- Disabling the In-Memory Column Store
  You can disable the In-Memory column store for your Autonomous Database by modifying the maximum In-Memory size to 0.

About Database In-Memory

Oracle Database In-Memory (Database In-Memory) is a suite of features that greatly improves performance for real-time analytics and mixed workloads.

The Database In-Memory feature is useful when:

- You have enabled the In-Memory feature for your on-premises database and are now migrating from the on-premises database to Autonomous Database.
- Queries in your database spend most of their run time scanning and filtering data, as well as performing joins and group-by aggregations.
- You need to improve the response time of your queries that are analytical in nature.

The In-Memory column store is the key feature of Database In-Memory. The In-Memory column store resides in the In-Memory Area, which is an optional portion of the System Global Area (SGA). When enabled for In-Memory, tables, partitions, sub-partitions and materialized views are populated in a columnar format in the In-Memory column store. This columnar format is optimized for very fast scans, which can enable analytic queries to run orders of magnitude faster than row format-based queries. See Oracle Database In-Memory for more information.

The In-Memory column store dynamically grows over time to accommodate the hottest `INMEMORY` objects. You can use the `PRIORITY` table clause or have continuous table scans to ensure that the In-Memory column store grows large enough for your workload. See Sizing the In-Memory Column Store for more information.

You can use the `INMEMORY` clause in DDL statements to enable In-Memory for the following objects:

- Tables (including external tables)
- Partitions or sub-partitions
- Materialized views

See Enabling Objects for In-Memory Population for more information.

Note the following for enabling Database In-Memory:

- Database In-Memory is only available for an Autonomous Database instance with:
  - The ECPU compute model
  - A minimum of 16 ECPUs
- If you restart the database, objects enabled for `INMEMORY` require re-population into the In-Memory column store. Depending on the priority setting, re-population may happen automatically or on first access. If you are monitoring performance you will notice slower performance until the object is repopulated.

Enabling Objects for In-Memory Population

You can enable and disable objects for population in the In-Memory column store.

Only objects with the `INMEMORY` clause are eligible for population into the In-Memory column store. Use DDL statements `CREATE TABLE` or `ALTER TABLE` to apply the `INMEMORY` clause.

By default, objects are created with the `NO INMEMORY` attribute, which means they are not eligible for population into the In-Memory column store.

Enabling an object as `INMEMORY` is specifying that an object can potentially reside in the In-Memory column store. In-Memory population is a separate step that occurs when the database reads existing row-format data, transforms it into columnar format, and then stores it in the In-Memory column store. By default, all of the columns in an object with the `INMEMORY` attribute are populated into the In-Memory column store.

Setting the `INMEMORY` attribute on an object implies that the object is a candidate for population in the In-Memory column store. The object is not immediately populated to the In-Memory column store. However, you can use the `PRIORITY` clause with the `INMEMORY` attribute to control the priority of population into the In-Memory column store. The `PRIORITY` clause can be specified with the following values:

- `CRITICAL`: Object is populated immediately after the database is opened.
- `HIGH` : Object is populated after all `CRITICAL` objects have been populated, if space remains available in the In-Memory column store.
- `MEDIUM`: Object is populated after all `CRITICAL` and `HIGH` objects have been populated, and space remains available in the In-Memory column store.
- `LOW`: Object is populated after all `CRITICAL`, `HIGH`, and `MEDIUM` objects have been populated, if space remains available in the In-Memory column store.
- `NONE`: Objects are only populated after they are scanned for the first time, if space is available in the In-Memory column store. The priority `NONE` is the default priority.

Example to create a table with the `INMEMORY` attribute:

```
CREATE TABLE im_emp (
    id  NUMBER,  name VARCHAR2(20),
    depno NUMBER,  sal   NUMBER,
    mgr   NUMBER,
    loc   VARCHAR2(20))
    INMEMORY;
```

This example creates the `im_emp` table with the `INMEMORY` attribute. The example uses the default priority of `NONE` for the `INMEMORY` clause, which means the table is populated only after it is scanned for the first time.

Example to use the `PRIORITY` clause with the `INMEMORY` attribute:

```
CREATE TABLE im_emp_1 (
    id  NUMBER,  name VARCHAR2(20),
    depno NUMBER,  sal   NUMBER,
    mgr   NUMBER,
    loc   VARCHAR2(20))
    INMEMORY PRIORITY CRITICAL;
```

Example to set the `INMEMORY` attribute for an existing table:

```
ALTER TABLE employees INMEMORY;
```

This example sets the in the `INMEMORY` attribute for the `employees` table.

You can also alter a table to set the `INMEMORY` attribute only for a subset of columns in the table. For example:

```
ALTER TABLE im_emp_tb INMEMORY NO INMEMORY(depno);
```

This example sets the in the `INMEMORY` attribute for the `im_emp_tb`, but it excludes the `depno` column.

See INMEMORY_TABLE_CLAUSE for more information.

After `INMEMORY` is enabled for a table, you can query the `V$IM_SEGMENTS` view to determine if data from the table is populated in the In-Memory column store. For example:

```
SELECT SEGMENT_NAME, POPULATE_STATUS
FROM   V$IM_SEGMENTS
WHERE  SEGMENT_NAME = 'IM_EMP';
```

See V$IM_SEGMENTS for more information.

After the object is populated in the In-Memory column store, the object is only evicted when the object is dropped or moved, or the object is updated with the `NO INMEMORY` attribute. For example:

```
ALTER TABLE im_emp NO INMEMORY;
```

See How In-Memory Population Works for more information.

Notes for enabling objects for In-Memory population:

- The database does not automatically populate In-Memory objects with the `PRIORITY` setting of `NONE`. However, you can manually populate these objects into In-Memory column store. See Populating the IM Column Store Manually and DBMS_INMEMORY.POPULATE Procedure for more information.

- There must be continuous table scans to ensure that the In-Memory column store will automatically grow large enough for your workload. If an object is partially populated, the table scan can use both In-Memory and Exadata Smart Scan.

Sizing the In-Memory Column Store

When your Autonomous Database instance has 16 or more ECPUs, the Database In-Memory feature is enabled by default, and up to 50% of the SGA can be used by In-Memory.

The memory for the In-Memory Area is not reserved upfront, and the initial size of the In-Memory Area is 0. The In-Memory Area grows gradually each time there is insufficient space for the In-Memory population.

Objects in the In-Memory column store are dynamically managed by Automatic In-Memory (AIM). Automatic In-Memory is by default enabled for your Autonomous Database.

AIM uses internal statistics to determine how frequently In-Memory objects and columns are accessed. If the In-Memory column store is full and other more frequently accessed segments would benefit from population in the In-Memory column store, then the In-Memory column store evicts inactive segments.

To ensure that the working data set is always populated, AIM automatically evicts cold (infrequently accessed) segments. See Automating Management of In-Memory Objects for more information.

The default maximum In-Memory size for your Autonomous Database can be up to 50% of the SGA size. However, you can modify the maximum In-Memory size for your Autonomous Database:

- By changing the ECPU count for your Autonomous Database.
- By using the `DBMS_INMEMORY_ADMIN.SET_SGA_PERCENTAGE` procedure to modify the maximum In-Memory size. You can specify a value in the range between 0 and 70. Specifying a zero (0) value disables the Database In-Memory. For example:

  ```
  BEGIN
      DBMS_INMEMORY_ADMIN.SET_SGA_PERCENTAGE(60);
  END;
  /
  ```

  This example modifies the maximum In-Memory size to 60% of the SGA size.

See SET_SGA_PERCENTAGE Procedure for more information.

Query the `V$INMEMORY_AREA` view to retrieve the information about space allocation inside the In-Memory area. For example:

```
SELECT * FROM V$INMEMORY_AREA;
```

See V$INMEMORY_AREA for more information.

Disabling the In-Memory Column Store

You can disable the In-Memory column store for your Autonomous Database by modifying the maximum In-Memory size to 0.

For example:

```
BEGIN
    DBMS_INMEMORY_ADMIN.SET_SGA_PERCENTAGE(0);
END;
/
```

This example modifies the maximum In-Memory size value to zero (0) and disables Database In-Memory for the instance.

See SET_SGA_PERCENTAGE Procedure for more information.

6
Reference

- Previous Feature Announcements
  Announcements for the important changes made to Oracle Autonomous Database Serverless.
- Sample Star Schema Benchmark (SSB) Queries and Analytic Views
  The SSB schema contains the tables: lineorder, customer, supplier, part, and dwdate. The following is a list of sample queries and analytic views you can use against the SSB schema. Note that you need to prefix the table names with the schema name SSB in your queries.
- Autonomous Database Supplied Package Reference
  This appendix provides information about the packages you use with Autonomous Database. The `DBMS_CLOUD` topic also covers the `DBMS_CLOUD REST APIs`.
- Autonomous Database for Experienced Oracle Database Users
  This appendix provides information on using Autonomous Database for experienced Oracle Database users with Autonomous Database Serverless.
- Migrating MySQL and Third-Party Databases to Autonomous Database
- SODA Collection Metadata
  It describes SODA Collection metadata on the database.
- Cloud Support and Identity Information
  This lists the procedure to file a service request using Oracle Cloud Support.

Previous Feature Announcements

Announcements for the important changes made to Oracle Autonomous Database Serverless.

See What's New for Oracle Autonomous Database Serverless for feature announcements in 2023.

- 2024 What's New
  Announcements for the noteworthy changes made to Oracle Autonomous Database Serverless in the year 2024.
- 2023 What's New
  Announcements for the noteworthy changes made to Oracle Autonomous Database Serverless in the year 2023.
- 2022 What's New
  Announcements for the noteworthy changes made to Oracle Autonomous Database Serverless in the year 2022.
- 2021 What's New
  Announcements for the noteworthy changes made to Oracle Autonomous Database Serverless Deployment in the year 2021.
- 2020 What's New
  Announcements for the noteworthy changes made to Oracle Autonomous Database Serverless Deployment in the year 2020.

- 2019 What's New
  Announcements for the noteworthy changes made to Oracle Autonomous Database Serverless Deployment in the year 2019.
- 2018 What's New
  Announcements for the noteworthy changes made to Oracle Autonomous Database Serverless Deployment in the year 2018.

2024 What's New

Announcements for the noteworthy changes made to Oracle Autonomous Database Serverless in the year 2024.

- December 2024
- November 2024
- October 2024
- September 2024
- August 2024
- July 2024
- June 2024
- May 2024
- April 2024
- March 2024
- February 2024
- January 2024

December 2024

Feature	Description
Immutable Backup Retention	The **Immutable backup retention** option allows you to lock the automatic backup retention period. If locked, Autonomous Database does not allow changes to the retention period for automatic backups. See About Backup and Recovery on Autonomous Database for more information.
AutoML Pipeline	AutoML Pipeline automates the three major stages of the machine learning pipeline: • Algorithm selection • Adaptive data reduction (feature and sample size selection) • Hyperparameter optimization See AutoML Pipeline for more information.
OML4Py Metrics	OML4Py Metrics provides metrics to evaluate the performance of machine learning models. See OML4Py Metrics for more information.
SQL Query Performance Event: `QueryPerformanceDegradationDetected`	The `QueryPerformanceDegradationDetected` information event is triggered when the average elapsed time of a query has exceeded ten times its average elapsed time over the past seven (7) days. See Information Events on Autonomous Database for more information.

Feature	Description
Refreshable Clone Events: `AutomaticRefreshEnd` and `AutomaticRefreshFailed`	The `AutomaticRefreshEnd` and `AutomaticRefreshFailed` information events are generated when a scheduled Automatic Refresh completes or when a scheduled Automatic Refresh fails. See Information Events on Autonomous Database for more information.

November 2024

Feature	Description
Data Marketplace in Data Studio	Data Marketplace provides an innovative platform to access and share datasets effortlessly with Autonomous Database. The Data Marketplace is a part of the Data Studio suite of tools that provides you access to load and query datasets for your organizational requirements. You can access and discover data sets curated by data providers. See The Data Marketplace Tool for more information.
PAR URL Table View with Column Level Group By and UI Options	You can access Pre-Authenticated Request (PAR) URL data from a browser and group the data by the values of cells in a column. You can also specify view options such as sorting, filtering, and cell color by column. See Generate a Table Hyperlink with UI Features Specified on Columns for more information.
Network Usage Events: `NetworkUsageInbound` and `NetworkUsageOutbound`	The `NetworkUsageInbound` and `NetworkUsageOutbound` information events are triggered when the average incoming network traffic or average outgoing network traffic to your Autonomous Database over the previous 24 hours has exceeded 50% of the average incoming or outgoing traffic. See Information Events on Autonomous Database for more information.
Maintenance Window Details on OCI Console	The Oracle Cloud Infrastructure Console Autonomous Database details page shows additional information on upcoming maintenance. When Autonomous Data Guard is enabled this includes maintenance information for a local standby database. See About Scheduled Maintenance and Patching for more information.
Microsoft SharePoint Support for Database Links with Oracle-Managed Heterogeneous Connectivity	Use Autonomous Database Oracle-managed heterogeneous connectivity to create database links to Microsoft SharePoint. When you use database links with Oracle-managed heterogeneous connectivity, Autonomous Database configures and sets up the connection. See Create Database Links to Non-Oracle Databases with Oracle-Managed Heterogeneous Connectivity for more information.
Customer-Managed Keys with Cross-Region Disaster Recovery	Replication of key vaults across regions for a cross-region Disaster Recovery Peer is now supported both with OCI Vault default vault and virtual private vault options. See Manage Master Encryption Keys in OCI Vault for more information.

October 2024

Feature	Description
Bring Your Own License (BYOL) Option Moved for Provisioning and Cloning	During provisioning or cloning the Bring Your Own License (BYOL) option is moved to advanced options in the Configure the database area. See Choose Bring Your Own License Option While Provisioning or Cloning for more information.

Feature	Description
Blocking Sessions Metric	The `BlockingSessions` metric shows the number of current blocking sessions. See Available Metrics: oci_autonomous_database for more information.
Control Lockdown Profile Errors	Autonomous Database uses lockdown profiles to restrict certain database functionality and features. The `LOCKDOWN_ERRORS` parameter allows you to suppress lockdown profile errors. See LOCKDOWN_ERRORS for more information.
NVIDIA GPU Compute Capabilities with Oracle Machine Learning Notebooks	Oracle Machine Learning Notebooks offers support for NVIDIA GPUs (Graphics Processing Unit) compute capabilities. The GPU capabilities in OML Notebooks allow you to run advanced machine learning algorithms such as deep learning models, transformers (embedding models) for generating vectors, and small LLMs. The GPU feature is enabled for Autonomous Database instances with Transaction Processing or Data Warehouse workloads with 16 ECPUs or more specified for the OML application. GPU resources are not available on Always Free or Autonomous Database for Developers instances. See Enable GPU Compute Capabilities in a Notebook through the Python Interpreter for more information.
Customer-Managed Encryption Keys	In addition to Oracle Cloud Infrastructure Vault, Autonomous Database can be configured to use encryption keys with the following key providers: • Azure Key Vault • AWS Key Management System (KMS) • Oracle Key Vault (OKV) See Manage Encryption Keys on Autonomous Database for more information.
Select AI with Hugging Face	Autonomous Database can interact with AI service providers. Select AI supports OpenAI, Anthropic, Azure OpenAI Service, CohereAI, Google Gemini, Hugging Face, and OCI Generative AI. This feature supports having LLMs work with Oracle database by generating SQL from natural language prompts. This allows you to talk to your database. See Use Select AI for Natural Language Interaction with your Database for more information.
Zero Trust Packet Routing	Oracle Cloud Infrastructure Zero Trust Packet Routing (ZPR) protects sensitive data from unauthorized access through intent-based security policies that you write for resources, such as an Autonomous Database on a private endpoint, that you assign security attributes to. See Use Zero Trust Packet Routing on Autonomous Database for more information.
Select AI with Anthropic	Autonomous Database can interact with AI service providers. Select AI supports OpenAI, Anthropic, Azure OpenAI Service, CohereAI, Google Gemini, Hugging Face, and OCI Generative AI. This feature supports having LLMs work with Oracle database by generating SQL from natural language prompts. This allows you to talk to your database. See Use Select AI for Natural Language Interaction with your Database for more information.

Feature	Description
Data Studio Enhancements	The Data Studio is updated with new features and changes in the following areas: • Data Load in Data Studio- Public APIs for Data Load Live Feed: The LiveFeed Data Studio feature now has public PL/SQL APIs. See DBMS_LIVE_FEED Package Reference for more information. • Data Studio Settings Changes: – **Support for non-Oracle Cloud Infrastructure (OCI) Cloud Services** Data Studio now supports loading and linking to data in AWS and Azure. In addition, it supports AI Profiles from AWS. – **Support for Policy checking in Data Studio** Data Studio utilizes many OCI services to enrich your experience. If your OCI credential is missing policies, these features fail, typically silently. The Data Studio Settings UI now allows you to test your credential and see which features are likely to work, and which require missing policies. – **Improved Day-one user experience** Several improvements around Settings experiences have been added, so users are successful in their tasks. See Data Studio Settings for more information. • Data Studio Analysis- Explain query using AI: The Data Analysis tool already had AI support to generate SQL from a text prompt. You can now utilize AI to explain the SQL queries the tool generates in text format. See Generate Natural Language explanation from SQL Query using AI Explain for more information. • Data Studio Analysis- Vector Search for Text: You can use Oracle AI Vector Search capabilities to search for relevant text from your source data on a specific column. The Vector Search helps quickly locate and match data that is similar to the text you provide as your search input. See Vector Search in the Data Analysis Tool for more information.

September 2024

Feature	Description
Implicit Partitioning for External Tables	You can create implicit partitioned external tables on your Autonomous Database from Hive style partitioned data or from simple folder partitioned data stored on your Cloud Object Store. See Query External Tables with Implicit Partitioning for more information.
Use Select AI to Generate Synthetic Data	You can generate synthetic data using LLMs that conforms to your schema for solution testing, proofs of concept, and other uses. Synthetic data can be used for better testing of your applications when real data is not available and can also be used to populate an Autonomous Database clone or a metadata clone. See Synthetic Data Generation for more information.
Use Select AI Retrieval Augmented Generation (RAG)	Select AI enables SQL access to generative AI using Large Language Models (LLMs) and embedding models. This includes support for natural language to SQL query generation and retrieval augmented generation, among other features. See Use Select AI for Natural Language Interaction with your Database for more information.

Feature	Description
Select AI with Google Gemini	Autonomous Database can interact with AI service providers. Select AI supports Google Gemini, OCI Generative AI, Azure OpenAI Service, OpenAI, and CohereAI. This feature supports having LLMs work with Oracle database by generating SQL from natural language prompts. This allows you to talk to your database. See Use Select AI for Natural Language Interaction with your Database for more information.
Undelete Database	With the ECPU compute model, you have the option to undelete a terminated Autonomous Database instance. See Undelete a Terminated Autonomous Database Instance for more information.

August 2024

Feature	Description
Use Database In-Memory with Autonomous Database	You can now declare your tables as in-memory to enable the Database In-Memory feature on Autonomous Database. This feature can significantly improve performance for real-time analytics and mixed workloads. See Use Database In-Memory with Autonomous Database for more information.
Disaster Recovery Peer Lag Metric	The `PeerLag` metric shows the total time in seconds by which the Disaster Recovery peer lags behind its Primary database. See Available Metrics: oci_autonomous_database for more information.
Create Credential Cloud Shell script enhancements	The `adb-create-cred.sh` Create Credential Cloud Shell script used to create OCI Native Credentials and Auth Token in your Autonomous Database supports new arguments: `username` and `credential`. See Create Credentials to Access Cloud Services in Cloud Shell for more information.
Upgrade by Cloning from Oracle Database 19c to Oracle Database 23ai	When you clone an Autonomous Database instance and the source database version is Oracle Database 19c, in regions where Oracle Database 23ai is available, you can choose either Oracle Database 19c or Oracle Database 23ai for the clone. See Upgrade to Oracle Database 23ai by Creating a Full Clone for more information.
View Network Information on the OCI Console	On the Oracle Cloud Infrastructure Console you can view the network information for your Autonomous Database. See View Network Information on the OCI Console for more information.
Export and Import Operations Run as Jobs in Data Transforms	You can now use the Jobs page to see the status of the export or import process that is currently running. See Export Objects for more information.

July 2024

Feature	Description
Migrate Applications from Non-Oracle Databases to Autonomous Database	You can translate SQL statements written for a non-Oracle SQL database to Oracle SQL in your Autonomous Database. See Migrate Applications from PostgreSQL Databases to Autonomous Database for more information.
Autonomous Database 23ai Available in the Paid Tier	Autonomous Database with Oracle Database 23ai in the Paid tier is available in all commercial public cloud regions. See Provision an Autonomous Database Instance and Clone an Autonomous Database Instance for more information.
Update for Oracle-Managed Heterogeneous Connectivity Database Links	Autonomous Database support for Oracle-managed heterogeneous connectivity with database links for Azure, the two optional parameters include: `auth_method` and `domain` to support NTLM/AD authentication. Autonomous Database support for Oracle-managed heterogeneous connectivity with database links for Snowflake, the optional parameters include: `role`, `schema`, and `warehouse`. These values specify a different schema, role, or warehouse value, other than the default. See Create Database Links to Non-Oracle Databases with Oracle-Managed Heterogeneous Connectivity for more information.
OCI Full Stack Disaster Recovery	You can enable the Oracle Cloud Infrastructure Full Stack Disaster Recovery on Autonomous Database. See Use OCI Full Stack Disaster Recovery with Autonomous Database for more information.
Updated Object Storage URL Format using Dedicated Endpoints for commercial (OC1) realms	Autonomous Database supports Oracle Cloud Infrastructure Object Storage URI formats for commercial (OC1) realms. See Import Data Using Oracle Data Pump on Autonomous Database and Move Data with Data Pump Export to Object Store, for Oracle Data Pump examples.
Pre-Authenticated Request (PAR) URL Enhancements	Pre-Authenticated Request (PAR) URL enhancements include: • Bind variables are supported when generating a PAR URL and when accessing data. See Use Table Hyperlinks for Read Only Data Access on Autonomous Database for further information. • Request a PAR URL that uses HIGH, MEDIUM or LOW service-level guarantees and resources when accessing an object or SQL statement. See Generate a Table Hyperlink for a Table or a View and Generate a Table Hyperlink with a Select Statement for details. • The expiration time and access count for a PAR URL can be extended. See Generate a Table Hyperlink for a Table or a View and Generate a Table Hyperlink with a Select Statement for details.
Use `RESULT_CACHE` Hint to Cache Data for a Cloud Links Consumer	You can enable caching on a data set consumer to cache the results of a query or for a query fragment that accesses Cloud Link data. See Enable Caching for a Cloud Link Consumer for more information.
Updated Object Storage URL Format in DBMS_CLOUD	Autonomous Database supports the updated Oracle Cloud Infrastructure Object Storage URI formats in DBMS_CLOUD for: • OCI Native • Swift • S3 Compatible • Pre-Authenticated Requests (PAR) See DBMS_CLOUD URI Formats for further information.

Feature	Description
Text Indexing and Search on JSON Files in Object Storage	You can build a JSON search index on documents residing in files stored in Object Storage, which allows you to search fields in your JSON documents using `JSON_TEXTCONTAINS`, and use wildcards in your search. See Use JSON Search on Documents in Object Storage for more information.
Documentation Addition: Go Application Connection Information	You can connect Go applications to your Autonomous Database instance either with a wallet (mTLS) or without a wallet (TLS). See Connect Go Applications to Autonomous Database for more information.
Database Actions Data Studio Updates	The Data Studio in Database Actions is updated with new features and changes in the following areas: • New Catalog page: The new catalog page features faceted search for easy filtering, presentation by entity type, better defaults for catalog features, and many other small improvements. See The Catalog Tool. • Data Load support for Boolean data type: The Data Load tool now supports loading Boolean data from source files into the new Oracle Database 23ai Boolean data type. See Loading Data. • Data Load from Image Files: The Data Load tool supports loading images of tabular data into relational tables. The supported file formats are `PDF`, `Tif` and `JPEG` formats. See Use OCI Document Understanding to extract Tables from Images. • Data Load support for Vector data type: The Data Load tool supports loading into the 23ai `VECTOR` data type. See Loading Data. • Support data load from Iceberg files: Data Studio can now link and create an external table on or load Iceberg tables from the object store. See Load Apache Iceberg Tables. • Rectify Geometry support for data load from GeoJSON files: Data Studio supports `RECTIFY` feature of GeoJSON spatial issues along with expression support for area and length. See Add Spatial Expression for GeoJSON columns. • Support for Language Detection using OCI AI Services: You can now add expressions to your target table to identify the language used in a column. See Use OCI Language Service Capabilities in Data Studio. • Data Studio enhancements for AWS Glue Administration: You can now add a new Glue Catalog and sync with a Glue catalog from the Data Studio tool. See Create Amazon Web Services (AWS) Credentials.

June 2024

Feature	Description
Autonomous Database 23ai Available in the Paid Tier	Autonomous Database with Oracle Database 23ai in the Paid tier is available in all commercial public cloud regions. See Provision an Autonomous Database Instance and Clone an Autonomous Database Instance for more information.
Elastic Pool Usage Details in Oracle Cloud Infrastructure Usage Report and in `OCI_USAGE_DATA` View	You can obtain a detailed breakdown of elastic pool usage in the Oracle Cloud Infrastructure usage report and this information is also shown in the `OCI_USAGE_DATA` view. See About Elastic Pool Billing for more information.
Documentation Addition: Provide Steps to Reduce Latency for Database Connections with Autonomous Data Guard	Provides information on the steps to take to configure an Autonomous Data Guard standby environment to reduce latency for database connections when you connect after a failover or after a switchover (when the standby becomes the Primary). See Steps to Reduce Latency for Database Connections for Databases with Autonomous Data Guard for more information.
Bring Your Own License (BYOL) Input Option for BYOL Flexibility	Enable the option **BYOL ECPU limit** to select the BYOL flexible license option. See Choose Bring Your Own License Option While Provisioning or Cloning and Choose Bring Your Own License on Autonomous Database (ECPU Compute Model) for more information.
Automate Creation of OCI Credentials using a Cloud Shell Script	Use the Cloud Shell to run a script that creates OCI Native Credentials and Auth Token in your Autonomous Database. See Example: Create OCI Native Credentials and Run in Autonomous Database for more information.
Public Access Allowed for an Autonomous Database on a Private Endpoint	When you want to configure an Autonomous Database to use a private endpoint and you also want to allow connections from allowed public IPs or public VCNs, you can select the **Allow public access** option. This adds a public endpoint for a database that is configured on a private endpoint. See Use a Private Endpoint with Public Access Allowed for more information.
TCP Fast Open (TFO) to Connect Clients to Autonomous Database	Use TCP Fast Open (TFO) connections when your application is sensitive to network latency and you want to decrease the network latency between your application and the database. TFO connections can improve application performance by decreasing network latency for client connections. See Use TCP Fast Open (TFO) to Connect Clients to Autonomous Database for more information.
Oracle APEX 24.1	Autonomous Database uses Oracle APEX Release 24.1. See Create Applications with Oracle APEX in Autonomous Database for more information.
Documentation Addition: Load Data from Google Cloud Storage	Provides an example for loading data into an Autonomous Database instance from Google Cloud Storage. See Load Data into Autonomous Database from Google Cloud Storage for more information.
Shrink Storage for Tables with LONG Columns	Autonomous Database shrink operations will succeed for tables with LONG columns. See Shrink Storage for more information.

Feature	Description
Cross Tenancy Autonomous Standby	You can add a cross tenancy Autonomous Data Guard standby database. This creates an Autonomous Data Guard standby that resides in a different tenancy. The cross tenancy standby can be in the same region (local) or in a different region (cross-region) See Use a Cross Tenancy Autonomous Data Guard Standby Database for more information.
Oracle Spatial AI Features for OML	Oracle Spatial AI is a set of capabilities in OML4Py for detecting patterns and making predictions from geospatial data. It includes features for end-to-end workflows, including spatial machine learning pipelines. See About Oracle Spatial AI for more information.
YouTube Support for Database Links with Oracle-Managed Heterogeneous Connectivity	Use Autonomous Database Oracle-managed heterogeneous connectivity to create database links to YouTube. When you use database links with Oracle-managed heterogeneous connectivity, Autonomous Database configures and sets up the connection. See Create Database Links to Non-Oracle Databases with Oracle-Managed Heterogeneous Connectivity for more information.
Data Transforms Updates	Data Transforms is updated with new features and changes in the following areas: • Export and Import Objects You can move Data Transforms objects such as projects, connections, data loads, data flows, workflows, and schedules from a Development environment to a Production environment. See Export Objects for more information. • Variables You can now create variables and use them as steps within data flows and workflows. See Create a Variable for more information. • Purge jobs Data Transforms allows you to purge older jobs to keep the repository clean and optimize performance. See Create and Manage Jobs for more information. • Load Views to target servers In addition to tables, Data Load allows you to load views from the source server. Views are stored as tables in the target server. • Support for the new Object Storage URL format Data Transforms includes support for the Object Storage URL format when you create an Object Storage connection. See Create an Oracle Object Storage Connection for more information. • OAuth 2.0 support for NetSuite connections You can now use OAuth 2.0 authentication to connect to Oracle NetSuite data sources. See Create and Use an Oracle NetSuite Connection for more information.

May 2024

Feature	Description
Multiple Peers for Cross-Region Disaster Recovery	With the ECPU billing model, you have the option to create one or more cross-region standby databases for Autonomous Data Guard and multiple cross-region Backup-Based Disaster Recovery peers. See Use Standby Databases with Autonomous Data Guard for Disaster Recovery and Use Backup-Based Disaster Recovery for more information.
Table Format with PAR URLs	You can access PAR URLs with a browser and view the returned data in table format. See Use Table Hyperlinks to Access Data in Table Format for more information.
Cross-Tenancy Access for IAM Users to Autonomous Database	Identity and Access Management (IAM) cross-tenancy user access to an Autonomous Database instance is similar to a single tenancy scenario except that tenancy information is required for mappings and token requests and a policy is required in both tenancies to allow this cross tenancy database resource access See Accessing a Database Cross-Tenancy Using an IAM Integration for more information.
External Calls with a Customer-Managed Wallet	When your Autonomous Database instance is on a private endpoint you can use a customer-managed wallet with procedures in `UTL_HTTP`, `DBMS_LDAP`, `UTL_SMTP`, or `UTL_TCP`. You can also use a customer-managed wallet when the scheduler sends SMTP email notifications for various events related to scheduler jobs. See Make External Calls Using a Customer-Managed Wallet for more information.
Allow Different ACLs on Cross-Region Disaster Recovery Peer	You can independently modify network ACLs on a remote disaster recovery peer database. See Manage Remote Peer Network ACLs and Update Remote Peer Network ACLs for a Backup-Based Disaster Recovery Peer for more information.
Documentation Addition: Best Practices for Low Latency Connections	Provides information on steps you can take to reduce the latency for connections between an application and Autonomous Database. See Best Practices for Low Latency Connections with Autonomous Database for more information.
Always Free with Oracle Database 23ai	The Oracle Database versions available for Always Free Autonomous Database are: Oracle Database 19c or Oracle Database 23ai. See Always Free Autonomous Database for more information.
Autonomous Database Free Container Image with Oracle Database 23ai	The Oracle Database versions available for Autonomous Database Free Container Image are: Oracle Database 19c or Oracle Database 23ai. See Use the Autonomous Database Free Container Image for more information.

April 2024

Feature	Description
Cross-Region Cloud Links	You can use Cloud Links in multiple regions, where the source region contains the data set's source database and one or more remote regions contain refreshable clones of the source database. See Register or Unregister a Data Set in a Different Region for more information.
Test Workloads Against an Upcoming Patch	Using the workload auto replay feature you can automatically capture a workload from a production database that is on the regular patch level and replay the workload on a target refreshable clone that is on the early patch level. This feature allows you to test an upcoming patch by running an existing workload that is in production against a patch before the patch reaches production. See Test Workloads Against an Upcoming Patch for more information.
Cloning Permissions	You can control and limit cloning operations on Autonomous Database with IAM permissions. See Cloning Permissions in IAM Permissions and API Operations for Autonomous Database for more information.

Feature	Description
Data Studio Updates	The Data Studio is updated with new features and changes in the following areas: • Data Load Dashboard: A new Data Load Dashboard is available which displays summaries of the recent load jobs and their status and view details about the data you load. See The Data Load Page for more information. • Sentiment Analysis and Key Phrase Extraction on text-based columns: You can now perform Sentiment Analysis and Key Phrase Extraction on the Data Load tool to analyze text-based data. See Use OCI Language Service Capabilities in Data Studio for more information. • Data Load support to access Amazon Web Services (AWS) Glue Catalogs: The Data Load tool of Data Studio allows you to synchronize Autonomous Database with Amazon Web Service (AWS) Glue Data Catalog metadata. See Register AWS Glue Catalog for more information. • Data Load support for registration and synchronization of OCI Data Catalogs: The Data Load tool enables you to register and synchronize OCI Data Catalogs to create external tables in Autonomous Database. See Register OCI Data Catalog for more information. • Data Load support to load GeoJSON files: The Data Load tool supports loading GeoJSON files in Autonomous Database. See Use GeoJSON in Data Load for more information. • Data Load support to enable loading into other schemas: The Data Load tool allows an ADMIN to create and load data into other schemas. See Loading Data From Local Files for more information. • Enable Resource Principal using Data Studio: You can enable Resource Principals using Data Studio to authenticate and access Oracle Cloud Infrastructure resources. See Manage Resource Principal using Data Studio for more information.
Manage Oracle Cloud Operator Access	Oracle Cloud Operators do not have authorization to access your data or any other information in your database schemas. When access to your database schemas is required to troubleshoot or mitigate an issue, you can allow a cloud operator to access your Autonomous Database schemas for a limited time. See Manage Oracle Cloud Operator Access for more information.
Apache Iceberg External Tables	Autonomous Database supports querying of Apache Iceberg tables stored in Amazon Web Services (AWS) or in Oracle Cloud Infrastructure Object Storage. See Query Apache Iceberg Tables for more information.

Feature	Description
Cloud Tables	You can create Cloud Tables where table data resides on Oracle managed Cloud Storage and the table data does not consume database storage. See Use Cloud Tables to Store Logging and Diagnostic Information for more information.
Cloud Links Offload Targets Without Specifying Database IDs	Cloud Links support with the optional `offload_targets` parameter lets you specify that access is offloaded to refreshable clones. This parameter now supports the `ANY` keyword, where all data set requests from consumers are offloaded to a refreshable clone. See Register a Data Set with Offload Targets for Data Set Access for more information.
Live Data Sharing Using Database Actions in Autonomous Database	Use Data Sharing to share data, as of the latest database commit, with Autonomous Databases in the same region. A Data Sharing recipient always views the latest data. See Overview of the Data Share Tool for more information.
Updated Version of Autonomous Database Free Container Image	A new version of the Autonomous Database Free Container Image is available with updates and new features. Use the Autonomous Database Free Container Image to run Autonomous Database in a container in your own environment, without requiring access to Oracle Cloud Infrastructure Console or to the internet. See Use the Autonomous Database Free Container Image for more information.

March 2024

Feature	Description
Refreshable Clone Automatic Refresh Option	When you enable the automatic refresh option Autonomous Database automatically refreshes a refreshable clone with data from the source database. See Refreshable Clone with Automatic Refresh Enabled and Edit Automatic Refresh Policy for Refreshable Clone for more information.
Invoke External Procedures	You can invoke external procedures written in the C language from Autonomous Database as SQL functions. You do not install external procedures on Autonomous Database. To use an external procedure, the procedure is hosted remotely on a VM running in an Oracle Cloud Infrastructure Virtual Cloud Network (VCN). See Invoke External Procedures as SQL Functions for more information.
Network File System (NFS) NFSv4 Support	You can attach a Network File System to a directory location in your Autonomous Database. This allows you to load data from Oracle Cloud Infrastructure File Storage in your Virtual Cloud Network (VCN) or from any other Network File System in on-premises data centers. Depending on the version of the Network File System you want to access, both NFSv3 and NFSv4 are supported. See Access Network File System from Autonomous Database for more information.
Support for Regular Expressions in DBMS_CLOUD Procedures	The format option `regexuri` is supported with the following `DBMS_CLOUD` procedures: `COPY_COLLECTION`, `COPY_DATA`, `CREATE_EXTERNAL_TABLE`, `CREATE_EXTERNAL_PART_TABLE`, and `CREATE_HYBRID_PART_TABLE`. See DBMS_CLOUD Package Format Options for more information

Feature	Description
Encrypt or Decrypt Data	For greater security you can encrypt data that you export to Object Storage. When data on Object Storage is encrypted you can decrypt the data you import or when you use the data in an external table. See Encrypt Data While Exporting or Decrypt Data While Importing for more information.
`JOB_QUEUE_PROCCESSES` Initialization Parameter	You can set the `JOB_QUEUE_PROCESSES` initialization parameter. Setting this parameter to 0 disables non-Oracle supplied Scheduler jobs. See JOB_QUEUE_PROCESSES for more information.
Pre-Authenticated Request URLs Access Information	You can generate and manage Pre-Authenticated Request (PAR) URLs for data on Autonomous Database and Autonomous Database provides views that allow you to monitor PAR URL usage. See Monitor and View Table Hyperlink Usage for more information.
`UTL_HTTP` allows `HTTP_PROXY` Connections	With `UTL_HTTP`, when the Autonomous Database instance is on a private endpoint you can call `UTL_HTTP.set_proxy` and both `HTTPS` and `HTTP_PROXY` connections are allowed. See Submit an HTTP Request to a Private Host and PL/SQL Package Notes for Autonomous Database for more information.

February 2024

Feature	Description
Pre-Authenticated Request URLs for Read Only Data Access	You can generate and manage Pre-Authenticated Request (PAR) URLs for data on Autonomous Database. See Use Table Hyperlinks for Read Only Data Access on Autonomous Database for more information.
Replicate Backups to a Cross-Region Disaster Recovery Peer	You can enable replication of automatic backups from a primary database to a cross-region disaster recovery peer. Up to seven days of automatic backups are available in the remote region when this feature is enabled. See Replicating Backups to a Cross-Region Autonomous Data Guard Standby and Replicating Backups to a Cross-Region Backup Based Disaster Recovery Peer for more information.
Select AI with OCI Generative AI	Autonomous Database can interact with AI service providers. Select AI now supports OCI Generative AI, Azure OpenAI Service, OpenAI, and CohereAI. This feature supports having LLMs work with Oracle database by generating SQL from natural language prompts. This allows you to talk to your database. See Use Select AI for Natural Language Interaction with your Database for more information.

January 2024

Feature	Description
ECPU Compute Model for all Workload Types	Autonomous Database offers two compute models when you create or clone an instance: ECPU and OCPU. Previously only the workload types Data Warehouse and Transaction Processing supported the ECPU compute model. Now all workload types support the ECPU compute model, including JSON and APEX. See Compute Models in Autonomous Database for more information.
Manage Time Zone File Updates	Autonomous Database provides several options to automatically update time zone files. See Manage Time Zone File Updates on Autonomous Database for more information.
Custom Kerberos Service Name Support for Kerberos Authentication	You can use a custom Kerberos service name when you configure Kerberos to authenticate Autonomous Database users. Using a custom service name you can use the same Keytab files on multiple Autonomous Database instances when you enable Kerberos authentication (you do not need to create and upload different Keytab files for each Autonomous Database instance). See Configure Kerberos Authentication with Autonomous Database for more information.
Clone Database to Change Workload Type	You can create a clone of an Autonomous Database and select a different workload type for the cloned database. See Clone Autonomous Database to Change Workload Type for more information.

2023 What's New

Announcements for the noteworthy changes made to Oracle Autonomous Database Serverless in the year 2023.

- December 2023
- November 2023
- October 2023
- September 2023
- August 2023
- July 2023
- June 2023
- May 2023
- April 2023
- March 2023
- February 2023
- January 2023

December 2023

Feature	Description
Cross Tenancy Cloning	You can clone an Autonomous Database instance from one tenancy, source tenancy, to a different tenancy (destination tenancy). The cross tenancy cloning option is only available using the CLI or the Autonomous Database REST APIs. This option is not available using the Oracle Cloud Infrastructure Console. See Cross Tenancy and Cross-Region Cloning for more information.
Invoke Azure Remote Functions	You can invoke Azure remote functions in your Autonomous Database as SQL functions. See Steps to Invoke Azure Functions as SQL Functions for more information.
Create Database Links to Oracle RAC Database	You can create database links from an Autonomous Database to a target Oracle RAC database and specify multiple hostnames. See Create Database Links from Autonomous Database to an Oracle Database on a Private Endpoint for more information.
New Cloud Links Views	Obtain information on the Cloud Link specific privileges granted to all users or to the current user. See Monitor and View Cloud Links Information for more information.
Documentation Addition: Load Data from AWS S3	Provides an example for loading data into an Autonomous Database instance from AWS S3. See Load Data into Autonomous Database from AWS S3 for more information.

November 2023

Feature	Description
Oracle APEX 23.2	Autonomous Database uses Oracle APEX Release 23.2. See Create Applications with Oracle APEX in Autonomous Database for more information.
Select AI with Azure OpenAI Service	Autonomous Database can interact with AI service providers. Select AI now supports the Azure OpenAI Service, OpenAI, and CohereAI. One way LLMs can work with Oracle database is by generating SQL from natural language prompts that allow you to talk to your database. See Use Select AI for Natural Language Interaction with your Database for more information.
Break Glass Access with `SAAS_ADMIN` User	Autonomous Database supports break glass access for SaaS providers. Break glass access allows a SaaS operations team, when explicitly authorized by a SaaS customer, to access a customer's database to perform critical or emergency operations. See Break Glass Access for SaaS on Autonomous Database for more information.
Database Links Without a Wallet (TLS)	You can create database links from one Autonomous Database instance to a publicly accessible Autonomous Database without a wallet (TLS). See Create Database Links from Autonomous Database to a Publicly Accessible Autonomous Database without a Wallet (TLS) for more information.

October 2023

Feature	Description
Vault Secret Credential Objects	You can use vault secret credentials to access cloud resources or to access other databases (use vault secret credentials anywhere where username/password type credentials are required). Supported vaults are: • Oracle Cloud Infrastructure Vault • Azure Key Vault • AWS Secrets Manager • GCP Secret Manager See Use Vault Secret Credentials for more information.
Documentation Addition: Billing Information	Documentation includes summary and detailed Autonomous Database billing information. See How Is Autonomous Database Billed? for more information.
Documentation Addition: Oracle Database Migration Information	Documentation includes information on migrating your Oracle Database to Autonomous Database. See Migrate Oracle Databases to Autonomous Database for more information.
ServiceNow Support for Database Links with Oracle-Managed Heterogeneous Connectivity	Autonomous Database support for Oracle-managed heterogeneous connectivity makes it easy to create database links to ServiceNow. When you use database links with Oracle-managed heterogeneous connectivity, Autonomous Database configures and sets up the connection. See Create Database Links to Non-Oracle Databases with Oracle-Managed Heterogeneous Connectivity for more information.
Invoke Oracle Cloud Infrastructure Functions or AWS Lambda Remote Functions	You can invoke Oracle Cloud Infrastructure Functions and AWS Lambda remote functions in your Autonomous Database as SQL functions. See Invoke User Defined Functions for more information.

September 2023

Feature	Description
Use Select AI to Generate SQL from Natural Language Prompts	Autonomous Database can interact with AI service providers including: OpenAI and CohereAI to provide access to Generative AI capability from Oracle Database that uses Large Language Models (LLMs). One way LLMs can work with Oracle database is by generating SQL from natural language prompts that allow you to talk to your database. See Use Select AI for Natural Language Interaction with your Database for more information.
Cloud Links Usability and Security	Cloud Links provide a cloud-based method to collaborate between Autonomous Databases. Cloud Links now allow you to set a data set owner. Cloud Links can be configured to require an additional authorization step, where a data set only allows specified databases to access the data set. When you register a data set, you can also specify refreshable clones to offload access to the data set. See Use Cloud Links for Read Only Data Access on Autonomous Database for more information.

Feature	Description
Elastic Pools	Use an elastic pool to consolidate your Autonomous Database instances, in terms of their allocation of compute resources, and to provide up to 87% cost savings. See Use and Manage Elastic Pools on Autonomous Database for more information.
Free Container Image	Use the Autonomous Database Free Container Image to run Autonomous Database in a container in your own environment, without requiring access to the Oracle Cloud Infrastructure Console or to the internet. See Use the Autonomous Database Free Container Image for more information.
ECPU-based Databases with Transaction Processing Workload Type Support per-Gigabyte Increments of Storage	You may now provision, clone, and scale using Gigabyte (GB) increments of storage for ECPU-based Autonomous Databases provisioned with the Transaction Processing workload type. The minimum storage allowed for an Autonomous Database instance in ECPU compute model with Transaction Processing workload is 20 GB. See Provision an Autonomous Database Instance for more information.
Google Analytics Support for Database Links with Oracle-Managed Heterogeneous Connectivity	Use Autonomous Database Oracle-managed heterogeneous connectivity to create database links to Google Analytics. When you use database links with Oracle-managed heterogeneous connectivity, Autonomous Database configures and sets up the connection. See Create Database Links to Non-Oracle Databases with Oracle-Managed Heterogeneous Connectivity for more information.

August 2023

Feature	Description
`DBMS_SHARE` Package	You can use the subprograms and views in the `DBMS_SHARE` package to share data with external systems. See DBMS_SHARE Package for more information.
Logical Partition Change Tracking (LPCT)	Logical Partition Change Tracking enables you to create logical partitions on base tables. It evaluates the staleness of the base tables for individual logical partitions without using a materialized view log or requiring any of the tables used in the materialized view to be partitioned. Logical Partition Change Tracking provides the capability to leverage the user-supplied logical partitioning information of base tables of materialized view for a more fine-grained, partition-level tracking of stale data for both refresh and rewrite purposes. While classical Partitioning Change Tracking (PCT) relies on the physical partitioning of tables, LPCT has no dependency on tables being physically partitioned; you can use LPCT with both partitioned and nonpartitioned tables. See Logical Partition Change Tracking and Materialized Views for more information.
Update OCPU to ECPU Compute Model	You can update an Autonomous Database instance from the OCPU billing model to the ECPU billing model. See Update to ECPU Billing Model on Autonomous Database for more information.

Feature	Description
Choose Backup Retention Period with ECPU Compute Model	With the ECPU billing model you have the option to select the backup retention period for automatic backups, with a retention period between 1 day and up to 60 days. See Edit Automatic Backup Retention Period on Autonomous Database for more information.
AWS Glue Data Catalog Integration	Autonomous Database allows you to synchronize with Amazon Web Services (AWS) Glue Data Catalog metadata. You can query data stored in S3 from Autonomous Database without having to manually derive the schema for the external data sources and create external tables. See Query External Data with AWS Glue Data Catalog for more information.

July 2023

Feature	Description
Track Table Changes with Flashback Time Travel	Use Flashback Time Travel to view past states of database objects or to return database objects to a previous state without using point-in-time media recovery. See Tracking Table Changes with Flashback Time Travel for more information.
Database Actions Quick Links	Database Actions on the Oracle Cloud Infrastructure Console provides quick links to select Database Actions cards or to select the Database Actions Launchpad. See Access Database Actions as ADMIN for more information.
Object Storage and REST Endpoint Management	`DBMS_CLOUD` supports a number of preconfigured Object Store and REST endpoints. `DBMS_CLOUD` also allows you to access additional customer-managed endpoints. See DBMS_CLOUD Endpoint Management for more information.
`detectfieldorder` DBMS_CLOUD Format Option	The `detectfieldorder` format option specifies that the fields in external data files are in a different order than the columns in the table. Specify this option with `DBMS_CLOUD.COPY_DATA`, `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`, `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE`, or `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE` to detect the order of fields using the first row of each external data file and map it to the columns of the table. See DBMS_CLOUD Package Format Options for more information.
Vault Secret Credential Objects	You can create vault secret credentials with secrets stored in Oracle Cloud Infrastructure Vault. You can use vault secret credentials to access cloud resources or to access other databases (use vault secret credentials anywhere where username/password type credentials are required). See Use Vault Secret Credentials for more information.
Export to Directory with `DBMS_CLOUD.EXPORT_DATA`	Export files as text or as Oracle Data Pump dump files to a directory. See Export Data to a Directory for more information.

June 2023

Feature	Description
Oracle Cloud Infrastructure Signing Key Based Authentication for Oracle Data Pump Import	Oracle Data Pump import using `impdp` supports credentials created using Oracle Cloud Infrastructure key based attributes. See Import Data Using Oracle Data Pump on Autonomous Database for more information.
Configurable Automatic Failover Data Loss Limit for a Local Standby Database	For Autonomous Data Guard you can specify an automatic failover data loss limit between 0 and 3600 seconds. Autonomous Data Guard performs automatic failover to a local standby database when a local standby database is available and the system can guarantee either the default zero data loss RPO or up to the data loss limit you specify. See Automatic Failover with a Standby Database for more information.
Oracle Workspace Manager	Oracle Workspace Manager provides an infrastructure that enables applications to create workspaces and group different versions of table row values in different workspaces. Use Oracle Workspace Manager to version-enable one or more user tables in the database. See Using Oracle Workspace Manager on Autonomous Database for more information.
Automatic Failover Begin and Automatic Failover End Events	The `AutonomousDatabase-AutomaticFailoverBegin` and `AutonomousDatabase-AutomaticFailoverEnd` events are generated when an automatic failover begins and ends. These events are only be triggered if you are using Autonomous Data Guard. See About Critical Events on Autonomous Database for more information.
Operator Access Event	The `OperatorAccess` event is generated when operator access is detected for the database. Autonomous Database operations team never access your data unless you explicitly grant permission through a Service Request for a specified duration. See Information Events on Autonomous Database for more information.
Built-in Tool Availability Service Level Objective (SLO)	Oracle will use commercially reasonable efforts to have the listed built-in tools meet the Availability objectives as defined in Service Level Objective (SLO) documentation. See Availability Service Level Objectives (SLOs) for more information.
CMU-AD with Active Directory Servers on a Private Endpoint	There are now two options for configuring Autonomous Database with Centrally Managed Users (CMU) with Microsoft Active Directory: • Active Directory (AD) servers publicly accessible: the Active Directory servers are accessible from Autonomous Database through the public internet. • Active Directory (AD) servers reside on a private endpoint: the Active Directory servers reside on a private endpoint and are not accessible from Autonomous Database through the public internet. See Use Microsoft Active Directory with Autonomous Database for more information.

May 2023

Feature	Description
Oracle APEX 23.1	Autonomous Database uses Oracle APEX Release 23.1. See Creating Applications with Oracle APEX in Autonomous Database for more information.
Google BigQuery Support for Database Links with Oracle-Managed Heterogeneous Connectivity	Autonomous Database support for Oracle-managed heterogeneous connectivity to create database links to Google BigQuery. When you use database links with Oracle-managed heterogeneous connectivity, Autonomous Database configures and sets up the connection. See Create Database Links to Non-Oracle Databases with Oracle-Managed Heterogeneous Connectivity for more information.
Log File Options for `DBMS_CLOUD`	`DBMS_CLOUD` provides `format` parameter options for logging customization: `enablelogs`, `logprefix`, `logdir`, and `logretention`. See DBMS_CLOUD Package Format Options for more information.
`RESULT_CACHE_MODE` Parameter is Modifiable at Session and System Level	Set the `RESULT_CACHE_MODE` parameter to specify which queries are eligible to store result sets in the result cache. Only query execution plans with the result cache operator will attempt to read from or write to the result cache. See RESULT_CACHE_MODE for more information.
Stripe Financial Views	Stripe is an online payment processing and credit card processing platform for businesses. The Stripe views allow you to query views created on top of Stripe APIs with `DBMS_CLOUD` to get Stripe information such as products, invoices, plans, accounts, subscriptions, and customers. See Stripe Views on Autonomous Database for more information.
Oracle GoldenGate Parallel Replicat (Integrated Mode)	Integrated Parallel Replicat (iPR) supports GoldenGate replication apply for procedural replication, Automatic CDR, and DML handlers. See Deciding Which Apply Method to Use and About Parallel Replicat for more information.
Persistent Messaging in `DBMS_PIPE`	The `DBMS_PIPE` package has extended functionality on Autonomous Database to support persistent messaging, where messages are stored in Cloud Object Store. See Use Persistent Messaging with Messages Stored in Cloud Object Store for more information.
Access Oracle Cloud Infrastructure Logging Data with Autonomous Database Views	Using the Oracle Cloud Infrastructure Logging Interface you can access log data from an Autonomous Database instance, in relational format. You can query log data in Oracle Cloud Infrastructure across all compartments and regions. See Oracle Cloud Infrastructure Logging Interface Views for more information.

April 2023

Feature	Description
Oracle Spatial Features for Geocoding Address Data and Reverse Geocoding Location Data	Oracle Spatial on Autonomous Database includes features for geocoding address data and for reverse geocoding longitude/latitude data to a street address. See Using Oracle Spatial with Autonomous Database for more information.

Feature	Description
Kerberos Authentication for CMU-AD	You can configure Autonomous Database to use Kerberos authentication for CMU with Microsoft Active Directory users. This configuration allows CMU Active Directory (CMU-AD) users to access an Autonomous Database instance using Kerberos credentials.
User Defined Notification Handler for Scheduler Jobs	Database Scheduler provides an email notification mechanism to track the status of periodically running or automated jobs. In addition to this, the Database Scheduler also supports user-defined PL/SQL Scheduler job notification handler procedure. Adding a scheduler job notification handler procedure allows you to monitor scheduled or automated jobs running in your Autonomous Database. See User Defined Notification Handler for Scheduler Jobs for more information.
Database Links with Customer-Managed Heterogeneous Connectivity to Non-Oracle Databases on a Private Endpoint	You can create database links from an Autonomous Database to an customer-managed Oracle Database Gateway to access Non-Oracle databases that are on a private endpoint. See Create Database Links with Customer-Managed Heterogeneous Connectivity to Non-Oracle Databases on a Private Endpoint for more information.
Oracle Machine Learning Notebooks Early Adopter	Oracle Machine Learning Notebooks Early Adopter is an enhanced web-based notebook platform for data engineers, data analyst, R and Python users, and data scientists. You can write code, text, create visualizations, and perform data analytics including machine learning. Notebooks work with interpreters in the back-end. In Oracle Machine Learning, notebooks are available in a project within a workspace, where you can create, edit, delete, copy, move, and even save notebooks as templates. See Get Started with Notebooks Early Adopter for Data Analysis and Data Visualization for more information.
Salesforce Support for Database Links with Oracle-Managed Heterogeneous Connectivity	Autonomous Database support for Oracle-managed heterogeneous connectivity to create database links to Salesforce databases. When you use database links with Oracle-managed heterogeneous connectivity, Autonomous Database configures and sets up the connection. See Create Database Links to Non-Oracle Databases with Oracle-Managed Heterogeneous Connectivity for more information.
Snapshot Standbys for Cross-Region Disaster Recovery	A disaster recovery cross-region peer can be converted to a snapshot standby. This converts either an Autonomous Data Guard cross-region standby database or a cross-region Backup-Based Disaster Recovery peer to a read-write database for up to two days. See Convert Cross Region Peer to Snapshot Standby for more information.
Singleton Pipe in `DBMS_PIPE`	Singleton Pipe is an addition to the `DBMS_PIPE` package that allows you to cache a custom message. Using a Singleton Pipe you can send and retrieve a custom message and share the message across multiple database sessions with concurrent reads. See Caching Messages with Singleton Pipes for more information.
Migrate Databases Across Regions with Minimal Downtime Using Autonomous Data Guard	You may now terminate, as well as perform all other Primary database actions on a remote Autonomous Data Guard standby when it becomes the Primary database, after a switchover or after a failover. This enables you to use a remote peer for migrating your database to the remote region. See Using Standby Databases with Autonomous Data Guard for Disaster Recovery for more information.

Feature	Description
Send Microsoft Teams Notifications	Use Microsoft Teams notifications. You can send messages, alerts, or the output of a query from Autonomous Database to a Microsoft Teams Channel. See Send Microsoft Teams Notifications from Autonomous Database for more information.
Send Email Using `DBMS_CLOUD_NOTIFICATION`	You can send email to a public SMTP endpoint using the Oracle Cloud Infrastructure Email Delivery service. See Send Email from Autonomous Database Using DBMS_CLOUD_NOTIFICATION for more information.

March 2023

Feature	Description
Cloud Links	Cloud Links provide a cloud-based method to collaborate between Autonomous Databases. Using Cloud Links you can access read-only data based on the cloud identities of Autonomous Database instances. The scope of collaboration can be the region where an Autonomous Database resides, individual tenancies, compartments, or particular Autonomous Database instances. See Using Cloud Links for Read Only Data Access on Autonomous Database for more information.
Backup-Based Disaster Recovery	Backup-Based Disaster Recovery uses backups to instantiate a peer database at the time of switchover or failover. For local Backup-Based Disaster Recovery, existing local backups are utilized and there are no additional costs. Cross-Region Backup-Based Disaster Recovery is also available and incurs additional costs. See Using Backup-Based Disaster Recovery for more information.
Documentation Addition: Oracle Machine Learning with Autonomous Database	Provides information on options for using Oracle Machine Learning with Autonomous Database. See Machine Learning with Autonomous Database for more information.
Export Data to Cloud Object Storage as Parquet	Use `DBMS_CLOUD.EXPORT_DATA` to export data as text. The text format export options are CSV, JSON, Parquet, or XML. See Export Data to Object Store as Text (CSV, JSON, Parquet, or XML) for more information.
Oracle Real Application Testing Capture and Replay	Use Capture and Replay from one Autonomous Database instance to another Autonomous Database instance. This enables you to compare workloads across different Autonomous Database instances. These Autonomous Database instances may vary at patch levels, database versions, or regions. See Using Oracle Real Application Testing for more information.
Availability Domain Information	You can obtain tenancy details, including the Availability Domain (AD) for an Autonomous Database instance. See Obtain Tenancy Details for more information.
Provision an Autonomous Database Instance and Save as Stack	The provisioning option Save as Stack allows you to create and deploy an Autonomous Database instance using Resource Manager. See Provision Autonomous Database and Overview of Resource Manager for more information.

Feature	Description
Oracle Client for Microsoft Tools (OCMT)	Oracle Client for Microsoft Tools (OCMT) is a graphical user interface (GUI) native Microsoft Software installer (MSI) that simplifies ODP.NET setup and provides Autonomous Database connectivity to multiple Microsoft data tools. See Connect Power BI and Microsoft Data Tools to Autonomous Database for more information.
Long-Term Backups	You can create long-term backups on Autonomous Database with a retention period between three months and up to ten years. When you create a long-term backup you can create a one-time backup or set a schedule to automatically create backups that are taken weekly, monthly, quarterly, or annually (yearly). See About Backup and Recovery on Autonomous Database and Create Long-Term Backups on Autonomous Database for more information.
Use Directories to Load Data with `DBMS_CLOUD` Procedures	As an alternative to an object store location URI, you can specify a directory with `DBMS_CLOUD` procedures to load or unload data from files in a directory, including directories created on attached network file systems. See Load Data from Directories in Autonomous Database for more information.

February 2023

Feature	Description
Maximum Storage Space Metric	The `StorageMax` metric shows the maximum amount of storage reserved for the database. See Autonomous Database Metrics and Dimensions for more information.
Oracle Cloud Infrastructure Vault Secret for ADMIN Password	When you create or clone an Autonomous Database instance or when you reset the ADMIN password, you can use an Oracle Cloud Infrastructure vault secret to specify the ADMIN password. See Use Oracle Cloud Infrastructure Vault Secret for ADMIN Password for more information.
Inactive Connections Detected Event	The `InactiveConnectionsDetected` event is generated when the number of inactive database connections detected is above a certain ratio, compared to all database connections for the Autonomous Database instance. Subscribing to this event can help you keep track of unused connections. See About Information Events on Autonomous Database for more information.
Enable Automatic Time Zone File Updates	Autonomous Database lets you choose to enable the feature that automatically updates time zone files on your Autonomous Database instance. See Manage Time Zone File Version on Autonomous Database for more information.

Feature	Description
Database Relocate and Failed Logon Events	The `FailedLoginWarning` event is generated when the number of failed login attempts reaches 3 * number of total users in the last three (3) hours. The `InstanceRelocateBegin` and `InstanceRelocateEnd` events are triggered when an Autonomous Database instance is relocated to another Exadata infrastructure due to server maintenance, hardware refresh, a hardware issue, or as part of a resource scale up. See About Critical Events on Autonomous Database for more information.

January 2023

Feature	Description
ECPU Compute Model	Autonomous Database offers two compute models when you create or clone an instance: ECPU and OCPU. See Compute Models in Autonomous Database for more information.
Configure Built-in Tools	Autonomous Database includes built-in tools that you can enable and disable when you provision or clone a database, or at any time for an existing database. See Managing Autonomous Database Built-in Tools for more information.
Configure a Custom Private IP Address for a Private Endpoint	You can configure a private endpoint when you provision or clone an Autonomous Database instance, or if your instance is configured to use a public endpoint you can change the configuration to use a private endpoint. Optionally when you configure a private endpoint you can enter a custom private IP address. Providing a custom IP address allows you to preserve an existing private IP address when you migrate from another system. See Configure Private Endpoints for more information.
Cross-Region Clone from Backup	Autonomous Database provides the option to create a clone in another region when you are cloning a database and you select **Clone from a backup**. In this case, you choose your preferred region to select the region where you want to create the clone. See Clone an Autonomous Database from a Backup for more information.
Send Slack Notifications	You can send slack notifications from Autonomous Database. This allows you to send messages to a Slack channel or send the results of a query to a Slack channel. See Send Slack Notifications from Autonomous Database for more information.

2022 What's New

Announcements for the noteworthy changes made to Oracle Autonomous Database Serverless in the year 2022.

- December 2022
- November 2022
- October 2022

- September 2022
- August 2022
- July 2022
- June 2022
- May 2022
- April 2022
- March 2022
- February 2022
- January 2022

December 2022

Feature	Description
Data Pipelines for Loading and Exporting	Data pipelines allow you to repeatedly load data from object store, or export data to object store. Load pipelines provide continuous incremental data loading from external sources (as data arrives on object store it is loaded to a database table). Export pipelines provide continuous incremental data exporting to object store (as new data appears in a database table it is exported to object store). See Using Data Pipelines for Continuous Load and Export for more information.
Credential Objects for use with `UTL_HTTP`, `UTL_SMTP`, and `DBMS_LDAP`	You can use credential objects to set authentication for use with `UTL_HTTP`, `UTL_SMTP` or `DBMS_LDAP`. See Use Credential Objects to set SMTP Authentication, Use Credential Objects to Set HTTP Authentication, and PL/SQL Packages for more information.
Medium and High Service Limits with Auto Scaling	The number of concurrent statements for the high and medium services increases by three times when OCPU Auto Scaling is enabled. See Service Concurrency for more information.
Multiple Data Catalog Support	Instead of each Autonomous Database instance supporting a single connection to a Data Catalog instance, this feature allows connections to multiple Data Catalog instances. See About Querying with Data Catalog for more information.

November 2022

Feature	Description
Cloud Code Repository Branch Management and Schema Export and Install	Additions and improvements for Cloud Code Repository in the `DBMS_CLOUD_REPO` package for (Git) repository branch management and schema export and install. See Using and Managing a Cloud Code Repository with Autonomous Database for more information.
Log File Prefix and Retention Format Options for `COPY_DATA`	`DBMS_CLOUD.COPY_DATA` supports the format parameter options: `logretention` and `logprefix`. See DBMS_CLOUD Package Format Options for more information.

Feature	Description
Cross-Region Autonomous Data Guard: Remove Second Hostname from Database Connection Strings	The `wallet.zip` and the connection strings provided for a database with cross-region Autonomous Data Guard enabled no longer contain both the primary and standby database hostnames in a single connection string or wallet. To avoid connection retry delays and optimize your disaster recovery setup, it is recommended to use the Primary database wallet or connection string to connect to the Primary database, and the remote database wallet or connection string to connect to the remote database (when the remote database is Available to connect to, after a switchover or failover). See Cross-Region Autonomous Data Guard Connection Strings and Wallets for more information.
Data Transforms in Database Actions	Data Transforms is a built-in data integration tool and is accessible from Database Actions. Data Transforms is a simple-to-use, drag-and-drop, no-code tool to load and transform data from heterogeneous sources to Autonomous Database. Use Data Transforms for all data integration needs such as building data warehouses and feeding analytics applications. See The Data Transforms Page for more information.
Access Network File System (NFS) Directories	You can attach a Network File System to a directory location in your Autonomous Database. This allows you to load data from Oracle Cloud Infrastructure File Storage in your Virtual Cloud Network (VCN) or from any other Network File System in on-premises data centers. See Access Network File System from Autonomous Database for more information.
Customer Managed Keys with Cross-Region Autonomous Data Guard	Autonomous Database fully supports using customer-managed keys with a remote standby database with Autonomous Data Guard. See Autonomous Data Guard with Customer Managed Keys and Notes for Customer-Managed Keys with Autonomous Data Guard for more information.
Cross-Region Refreshable Clones	Create one or more clones in regions other than the region of your primary (source) database. Clones in remote regions can be refreshed from the source database. See Using Refreshable Clones with Autonomous Database for more information.
Use Text Index to Query Text on Object Storage	You can build an Oracle Text index on object store files. This allows you to search the object store for text and use wildcards with your search. See Query Text in Object Storage for more information.
Documentation Addition: Data Lake Capabilities with Autonomous Database	A new section provides information on using Autonomous Database as a data lakehouse. See Using Data Lake Capabilities with Autonomous Database for more information.
Bulk Operations for Files in the Cloud	The PL/SQL package `DBMS_CLOUD` offers parallel execution support for bulk file upload, download, copy, and transfer activities, which streamlines the user experience and delivers optimal performance for bulk file operations. See Bulk Operations for Files in the Cloud for more information.

October 2022

Feature	Description
Clone from Backup Select Latest Available Timestamp Option	When you clone an Autonomous Database instance you can select clone from latest backup. This option selects the latest available backup timestamp as the clone source. You can select this option if a database becomes unavailable and you want to create a clone, or for any reason when you want to create a clone based on the most recent backup. See Clone Autonomous Database from a Backup for more information.
Use Customer-Managed Encryption Keys Located in a Different Tenancy	You can use customer-managed master encryption keys with the Vault and Keys in a different tenancy. The Vault and the Autonomous Database instance can be in a different tenancy, but they must be in the same region. See Use Customer-Managed Encryption Key Located in a Remote Tenancy for more information.
Use Google Service Account to Access GCP Resources	You can use a Google service account to access Google Cloud Platform (GCP) resources from an Autonomous Database instance. See Use Google Service Account to Access Google Cloud Platform Resources for more information.
`SYSDATE_AT_DBTIMEZONE` Parameter is Modifiable at Session and System Level	Depending on the value of `SYSDATE_AT_DBTIMEZONE`, you see either the date and time based on the default Autonomous Database time zone, Coordinated Universal Time (UTC), or based on the time zone that you set in your database. See SYSDATE_AT_DBTIMEZONE for more information.
View Oracle Cloud Operations Actions	The `DBA_OPERATOR_ACCESS` view stores information on the actions that Oracle Cloud Infrastructure cloud operations performs on your Autonomous Database. This view will not show results if Oracle Cloud Infrastructure cloud operations hasn't performed any actions or run any statements in your Autonomous Database instance. See View Oracle Cloud Infrastructure Operations Actions for more information.

September 2022

Feature	Description
Expanded list of Databases with Oracle-Managed Heterogeneous Connectivity	Autonomous Database support for Oracle-managed heterogeneous connectivity makes it easy to create database links to non-Oracle databases. When you use database links with Oracle-managed heterogeneous connectivity, Autonomous Database configures and sets up the connection to the non-Oracle database. The list of supported non-Oracle databases is expanded to include Hive and MongoDB. See Create Database Links to Non-Oracle Databases with Oracle-Managed Heterogeneous Connectivity for more information.
Oracle APEX Upgrade Events	The `APEXUpgradeBegin` and `APEXUpgradeEnd` events are generated when you are using Oracle APEX and your Autonomous Database instance begins an upgrade or completes an upgrade to a new Oracle APEX release. See About Events Based Notification and Automation on Autonomous Database for more information.

Feature	Description
Oracle MTS (OraMTS) Recovery Service	Use Oracle MTS (OraMTS) Recovery Service to resolve in-doubt Microsoft Transaction Server transactions on Autonomous Database. See Using OraMTS Recovery Feature on Autonomous Database for more information.

August 2022

Feature	Description
Documentation Addition: Updated Python Connection Information	You can connect Python applications to your Autonomous Database instance either with a wallet (mTLS) or without a wallet (TLS), using the python-oracledb driver. See Connect Python Applications to Autonomous Database for more information.
Autonomous Data Guard SLA	Autonomous Database provides a higher 99.995% availability SLA service commitment with Autonomous Data Guard enabled with a standby database. Database availability SLA is calculated based on the "Monthly Uptime Percentage" described in the Oracle PaaS and IaaS Public Cloud Services Pillar Document, document under **Delivery Policies**.
Oracle Machine Learning for R (OML4R)	Oracle Machine Learning for R (OML4R) is a component of the Oracle Machine Learning family of products, which integrates R with Autonomous Database. Use Oracle Machine Learning for R to: • Perform data exploration and data preparation while seamlessly leveraging Oracle Database as a high-performance computing environment. • Run user-defined R functions on database spawned and controlled R engines, with system-supported data-parallel and task-parallel capabilities. • Access and use powerful in-database machine learning algorithms from R language. See About Oracle Machine Learning for R for more information.
Synchronization with Data Catalog Enhancements	Use the `DBMS_DCAT$SYNC_LOG` view to easily access the log for the latest Data Catalog synchronization execution. See DBMS_DCAT SYNC_LOG View for more information.
Regional Availability Metrics	You can view regional availability metrics by data center for Oracle Autonomous Database. See Monitor Regional Availability of Autonomous Databases for more information.

July 2022

Feature	Description
Use Oracle Java	Autonomous Database supports Oracle JVM. Oracle JVM is a standard, Java-compatible environment that runs any pure Java application. See Using Oracle Java on Autonomous Database for more information.

Feature	Description
Oracle APEX 22.1	Autonomous Database uses Oracle APEX Release 22.1. See Creating Applications with Oracle APEX in Autonomous Database for more information.
Wallet Rotation with a Grace Period	Autonomous Database allows you to rotate wallets for an Autonomous Database instance or for all instances that a cloud account owns in a region, with a grace period of 1 hour to 24 hours. See Rotate Wallets for Autonomous Database for more information.
Change Display Name	You can change the display name for an Autonomous Database instance. See Update the Display Name for an Autonomous Database Instance for more information.
Default Database Name	The default database name and the matching default display name supplied when you provision or clone an instance is a generated 16-character string. See Provision Autonomous Database and Clone an Autonomous Database Instance for more information.
Oracle Real Application Testing: Database Replay	You can use Oracle Real Application Testing Database Replay to capture workload from an on-premises or other cloud service database and replay it on an Autonomous Database instance. This enables you to compare workloads between an on-premises database or other cloud service database and an Autonomous Database. See Using Oracle Real Application Testing - Database Replay for more information.
Export Data with Column Names in Header Row	You can optionally write column names as the first line in output files when you use the `EXPORT_DATA` procedure with CSV output. See EXPORT_DATA Procedure and DBMS_CLOUD Package Format Options for EXPORT_DATA with Text Files (CSV, JSON, and XML for more information.

June 2022

Feature	Description
Character Set Support	The Autonomous Database default database character set is Unicode AL32UTF8 and the default national character set is AL16UTF16. When you provision a database, depending on the workload type, you can select a database character set and a national character set. See Choose a Character Set for Autonomous Database for more information.
Support for Oracle LogMiner	Autonomous Database supports using Oracle LogMiner. See Oracle LogMiner for more information.
Synchronization with Data Catalog Enhancements	Synchronize Data Catalog partitioned metadata with Autonomous Database to create partitioned external tables. See the DBMS_DCAT RUN_SYNC procedure.
Database Actions Changes	Database Actions includes new features. See Changes in Oracle Database Actions for more information.

Feature	Description
Network Security Group (NSG) Configuration with Private Endpoint	For a private endpoint, the incoming and outgoing connections are limited by the combination of ingress and egress rules defined in NSGs and in the Security Lists associated with the private endpoint's VCN. Adding an NSG is now optional; when you do not include an NSG the ingress and egress rules defined in the Security Lists for the VCN still apply. See Configuring Network Access with Private Endpoints for more information.
Simplified Configuration Steps with CMU for Microsoft Active Directory	You can configure Autonomous Database to authenticate and authorize Microsoft Active Directory users. The configuration steps are simplified for enabling Autonomous Database with Centrally Managed Users (CMU). See Use Microsoft Active Directory with Autonomous Database for more information.

May 2022

Feature	Description
Azure Active Directory (Azure AD) Integration	Azure Active Directory (Azure AD) users can connect to an Autonomous Database instance using Azure OAuth2 access tokens. See Use Azure Active Directory (Azure AD) with Autonomous Database for more information.
OCPU Limit for License Type Bring Your Own License (BYOL)	The license type Bring Your Own License (BYOL), Oracle Database Standard Edition sets the maximum number of OCPUs that you can use to 8. See View and Update Your License and Oracle Database Edition on Autonomous Database for more information.
Kerberos Authentication Support	You can use Kerberos to authenticate Autonomous Database users. See Configure Kerberos Authentication for more information.
Autonomous Data Guard Standby State	Autonomous Data Guard shows a standby database in "Standby" state. See Autonomous Database Standby Database State for more information.
Database Actions Additions	Database Actions provides all the functionality that is available on the Autonomous Database Service Console. The Service Console will be deprecated soon. See Service Console Replacement with Database Actions for more information on where to find Service Console functionality in Database Actions.
Autonomous Data Guard Support for Terraform	Autonomous Data Guard supports operation from Terraform scripts See the following for more information: • Use the API • oci_database_autonomous_database • Oracle Cloud Infrastructure Provider • Terraform Registry Data Source: Data Source: oci_database_autonomous_database
Use Azure Service Principal	You can use an Azure service principal with Autonomous Database to access Azure resources without having to create and store your own credential objects in the database. See Use Azure Service Principal to Access Azure Resources for more information.

Feature	Description
Monitor Autonomous Database Availability	You can monitor availability information for your Autonomous Database instance form the Oracle Cloud Infrastructure Console or by viewing metrics. See Monitor Autonomous Database Availability and View Metrics for an Autonomous Database Instance for more information.
Database Names Length Limit	The length limit for database names when you create or clone your database has been increased from 14 to 30 characters. This change provides more flexibility, giving you the option to create longer database names. See Provision Autonomous Database or Clone an Autonomous Database Instance for more information.
Egress Rules for all Outbound Connections with Private Endpoints	When you define a private endpoint for your Autonomous Database instance you can provide enhanced security by setting a database property to enforce that all outgoing connections to a target host are subject to and limited by the private endpoint's egress rules. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

April 2022

Feature	Description
Display Autonomous Database Availability	You can monitor Autonomous Database availability from the Oracle Cloud Infrastructure Console or using metrics. See Monitor Autonomous Database Availability for more information.
OML4Py SQL API for Embedded Python Execution	OML4Py embedded Python execution on Autonomous Database supports a SQL API in addition to the REST API. Embedded execution allows you to run user-defined Python functions in database-spawned and managed Python engines, along with data-parallel and task-parallel automation. See SQL API for Embedded Python Execution with On-premises Database for more information.
Create Database Links from Autonomous Database to Oracle Databases on a Private Endpoint without a Wallet	You can create database links from an Autonomous Database to a target Oracle Database that is on a private endpoint and connect without a wallet (TCP). See Create Database Links from Autonomous Database to Oracle Databases on a Private Endpoint without a Wallet for more information.
Availability Service Level Objectives (SLOs)	Describes the Service Level Objectives (SLOs) for Oracle Autonomous Database. See Availability Service Level Objectives (SLOs) for more information.
Synchronization with Data Catalog Enhancements	The `DBMS_DCAT RUN_SYNC` procedure has new parameters to give you more control over what objects are synchronized and how the synchronization is performed. See CREATE_SYNC_JOB Procedure and RUN_SYNC Procedure for more information.
Data Catalog Views	Discover all accessible data catalogs in the current region or across all regions. See ALL_DCAT_GLOBAL_ACCESSIBLE_CATALOGS View and ALL_DCAT_LOCAL_ACCESSIBLE_CATALOGS View for more information.

March 2022

Feature	Description
Oracle APEX 21.2	Autonomous Database uses Oracle APEX Release 21.2. See Creating Applications with Oracle APEX in Autonomous Database for more information.
OML4Py Client Access	A new client package for OML4Py provides a “universal" client in that users can access Autonomous Database from a standalone Python engine on Linux to connect to Autonomous Database as well as 19c and newer Oracle Database releases. This supports working with OML4Py on Autonomous Database, local third-party IDEs, and other notebook environments, like JupyterLab. It further allows switching between on-premises and Autonomous Database instances using the same client package. See Install OML4Py Client for Linux for Use With Autonomous Database for more information.
Oracle APEX Upgrade Available Event	The `APEXUpgradeAvailable` event is generated when you are using Oracle APEX and a new release becomes available. See About Events Based Notification and Automation on Autonomous Database for more information.
Database Links with Oracle-Managed Heterogeneous Connectivity	Autonomous Database support for Oracle-managed heterogeneous connectivity makes it easy to create database links to non-Oracle databases. When you use database links with Oracle-managed heterogeneous connectivity, Autonomous Database configures and sets up the connection to the non-Oracle database. See Create Database Links to Non-Oracle Databases with Oracle-Managed Heterogeneous Connectivity for more information.
Metadata Columns in External Tables	To identify which file each row is coming from in your external tables, you can query the columns named `file$name` and `file$path` to find the source file name and the Object Store path URL for that file. See External Table Metadata Columns for more information.
External Tables with Partitioning Specified in Source Files	When you use external file partitioning with the partitions specified in external source files, Autonomous Database analyzes Cloud Object Store path information to determine the partition columns and data types. See Query External Tables with Partitioning Specified in Source Files for more information.
Identity and Access Management (IAM) Authentication Additional Features	You can configure Autonomous Database to use Oracle Cloud Infrastructure Identity and Access Management (IAM) authentication and authorization to allow IAM users to access an Autonomous Database with IAM credentials. New additions for IAM support include the following: defining global user mapping, defining global roles mapping, support for resource principal usage with IAM, and proxy user support. See Use Identity and Access Management (IAM) Authentication with Autonomous Database for more information.
Storage Auto Scaling	With **Storage auto scaling** enabled the Autonomous Database can expand to use up to three times the reserved base storage. If you need additional storage, the database automatically uses the reserved storage without any manual intervention required. Storage auto scaling is disabled by default. See Storage Auto Scaling for more information.

Feature	Description
Documentation Addition: Get Help, Search Forums, and Contact Support	When you use Autonomous Database, sometimes you need to get help from the community or to talk to someone in Oracle support. This documentation addition provides information about getting help by viewing and posting questions on forums and using Oracle Cloud Support to create a support request. See Get Help, Search Forums, and Contact Support for more information.

February 2022

Feature	Description
Track Oracle Cloud Infrastructure Resources, Cost and Usage	Autonomous Database provides details on Oracle Cloud Infrastructure resources, cost and usage. See Track Oracle Cloud Infrastructure Resources, Cost and Usage Reports with Autonomous Database Views for more information.
Critical Events with Customer-Managed Keys	Autonomous Database events are triggered when an instance is using customer-managed keys and the database becomes inaccessible or when the database recovers from being inaccessible. See About Critical Events on Autonomous Database for more information.
Time Zone Handling in Calls to `SYSDATE` and `SYSTIMESTAMP`	The initialization parameter `SYSDATE_AT_DBTIMEZONE` parameter enables special handling in a session for the date and time value returned in calls to `SYSDATE` and `SYSTIMESTAMP`. Depending on the value of `SYSDATE_AT_DBTIMEZONE`, you see either the date and time based on the default Autonomous Database time zone, Coordinated Universal Time (UTC), or based on the time zone that you set in your database. See SYSDATE_AT_DBTIMEZONE for more information.

January 2022

Feature	Description
Use a Cloud Code Repository	Autonomous Database provides routines to manage and store files in Cloud Code (Git) Repositories. The supported Cloud Code Repositories are: GitHub Repository, AWS CodeCommit, and Azure Repos. See Using and Managing a Cloud Code Repository with Autonomous Database for more information.
Oracle Machine Learning Notebooks Repository	Oracle Machine Learning Notebooks repositories are stored in a schema on the Autonomous Database instance. This enables database instance-specific backup and restore of Oracle Machine Learning Notebooks as well as support for cross-region Autonomous Data Guard. See What's New in Oracle Machine Learning on Autonomous Database for more information.
Oracle Machine Learning Notebooks Import and Export of Jupyter Format Notebooks	Oracle Machine Learning Notebooks supports importing and exporting Jupyter format notebooks. This makes it easier for users to interoperate with other notebook environments. See Export a Notebook and Import a Notebook for more information.

Feature	Description
Add My IP Address Button	When you provision an Autonomous Database instance and configure network access with Access Control Lists (ACLs) or when you update network access, click **Add My IP Address** to add your current IP address to the ACL entry. See Configuring Network Access with Access Control Rules (ACLs) for more information.
Reconnect a Disconnected Refreshable Clone	You can reconnect a refreshable clone to the source database. This allows you to use a refreshable clone as a test database and run DML and make changes while the database is disconnected. When you are done with your testing you can reconnect to the source database, which refreshes the clone to the point where it was when you disconnected. See Reconnect a Refreshable Clone to the Source Database for more information.
GitHub Raw URL support in `DBMS_CLOUD`	Use GitHub Raw URLs with `DBMS_CLOUD` APIs to access source files that reside on a GitHub Repository. See GitHub Raw URL Format for more information.
View Maintenance Status and Timezone Version Notifications	The `DB_NOTIFICATIONS` view stores information about maintenance status notifications and timezone version upgrade notifications for your Autonomous Database instance. See View Maintenance Status and Timezone Version Notifications for more information.
One-way TLS Connections on Linux x64 Clients using Oracle Call Interface without a Wallet	Using TLS authentication for Oracle Call Interface connections on Linux x64, a wallet is not required when the client program connects with Oracle Instant Client 19.13. Using TLS authentication for Oracle Call Interface connections on all other platforms, or without Oracle Instant Client 19.13, clients must provide a generic CA root certificate wallet. See Prepare for Oracle Call Interface, ODBC, and JDBC OCI Connections Using TLS Authentication for more information.

2021 What's New

Announcements for the noteworthy changes made to Oracle Autonomous Database Serverless Deployment in the year 2021.

- December 2021
- November 2021
- October 2021
- September 2021
- August 2021
- July 2021
- June 2021
- May 2021
- April 2021
- March 2021
- February 2021
- January 2021

December 2021

Feature	Description
Oracle Database API for MongoDB	Oracle Database API for MongoDB lets applications interact with collections of JSON documents in Autonomous Database using MongoDB commands. See Using Oracle Database API for MongoDB for more information.
Oracle Data Miner Access to Autonomous Database	Oracle Data Miner in SQL Developer can be used with Oracle Autonomous Database. Oracle Data Miner is a SQL Developer extension that enables users to create, schedule, and deploy analytical workflows through a drag-and-drop user interface. Oracle Data Miner serves as a productivity tool for data scientists and an enabler for citizen data scientists with a no-code machine learning environment that automates some common machine learning steps. See Oracle Data Miner for more information.
Autonomous Data Guard Paired Regions	The Autonomous Data Guard paired region list is expanded with additional regions. Autonomous Data Guard paired regions are remote regions where you can create a cross-region standby database. See Autonomous Data Guard Paired Regions for more information.
`DBMS_PIPE` PL/SQL Package	The `DBMS_PIPE` package that lets two or more sessions in the same instance communicate is available in Autonomous Database.
Cross-Region Cloning	When you clone an Autonomous Database instance using the Oracle Cloud Infrastructure Console, you can select your preferred region for the clone. You can create the clone in the current region or in a remote region. The list of available regions to create a clone only shows the regions that you are subscribed to. See Clone an Autonomous Database Instance for more information.
Identity and Access Management (IAM) Authentication	You can configure Autonomous Database to use Oracle Cloud Infrastructure Identity and Access Management (IAM) authentication and authorization to allow IAM users to access an Autonomous Database with IAM credentials. See Use Identity and Access Management (IAM) Authentication with Autonomous Database for more information.
Schedule Start and Stop Times	When you enable an Auto Start/Stop Schedule for an Autonomous Database instance, the instance automatically starts and stops according to the schedule you specify. This allows you to reduce costs by scheduling shutdown periods for times when a system is not in use. See Schedule Start and Stop Times for an Autonomous Database Instance for more information.
PL/SQL API for Switching Consumer Groups	Use the `CS_SESSION` Package to switch the database service and consumer group of an existing session. See CS_SESSION Package for more information.
Export Data as CSV, JSON, or XML Data Files	You can export data from Autonomous Database as text in the following formats: CSV, JSON, or XML. See Move Data to Object Store as CSV, JSON, or XML Using EXPORT_DATA for more information.
Use Database Links with a Target Database on a Private Endpoint	Create database links from an Autonomous Database source to a target Oracle Database that is on a private endpoint. To create a database link to a target database on a private endpoint, the target database must be accessible from the source database's Oracle Cloud Infrastructure VCN. See Create Database Links from Autonomous Database to Oracle Databases on a Private Endpoint for more information.

November 2021

Feature	Description
Database Management	The **Associated Services** area on the Oracle Cloud Infrastructure Console includes a **Database Management** link. You can use Database Management service to monitor the health of a single Autonomous Database or a fleet of Autonomous Databases. See Use Database Management Service to Monitor Databases for more information.
Database Actions Access	Access Database Actions on the Oracle Cloud Infrastructure Console with **Database Actions** button. See Connect with Built-in Oracle Database Actions for more information.
DBMS_LDAP and UTL_TCP Packages	The DBMS_LDAP and UTL_TCP packages are available, with some restrictions. See Restrictions and Notes for Database PL/SQL Packages for more information.
Concurrent Operations with Lifecycle Management Actions	When you initiate operations that take some time to complete, including scaling the system (Scale Up/Down), these operations no longer prevent you from performing other operations. For example, if you perform certain database lifecycle management actions such as stopping the database during a long-running operation, depending on the operation, the long running operation is either canceled or paused. See Concurrent Operations on Autonomous Database for more information.

October 2021

Feature	Description
Import JSON Data from Cloud Object Store	You can import JSON data from Cloud Object Store into a table. This import method supports all the Cloud Object Stores supported by Autonomous Database, and you can use an Oracle Cloud Infrastructure resource principal to access your Oracle Cloud Infrastructure Object Store or Amazon Resource Names (ARNs) to access AWS Simple Storage Service (S3). See Create Credentials and Copy JSON Data into an Existing Table for more information.
Blockchain Tables Support	Blockchain tables are insert-only tables that organize rows into a number of chains. Blockchain tables protect data that records important actions, assets, entities, and documents from unauthorized modification or deletion by criminals, hackers, and fraud. Blockchain tables prevent unauthorized changes made using the database and detect unauthorized changes that bypass the database. See Managing Blockchain Tables for more information.
Immutable Tables Support	Immutable tables are read-only tables that prevent unauthorized data modifications by insiders and accidental data modifications resulting from human errors. See Managing Immutable Tables for more information.
Integrate with Oracle Cloud Infrastructure Data Catalog	OCI Data Catalog metadata is synchronized with Autonomous Database creating schemas and external tables. You are allowed to immediately query data available in object store. See Query External Data with Data Catalog for more information.

Feature	Description
Oracle APEX 21.1	Autonomous Database uses Oracle APEX Release 21.1. See Creating Applications with Oracle Application Express in Autonomous Database for more information.

September 2021

Feature	Description
TLS and Mutual TLS Connections	Autonomous Database by default supports Mutual TLS (mTLS) connections. You have the option to configure an Autonomous Database instance to support both mTLS and TLS connections. See the following for more information: • About Connecting to an Autonomous Database Instance • Update Network Options to Allow TLS or Require Only Mutual TLS (mTLS) Authentication on Autonomous Database
Vanity URLs for APEX, ORDS and Database Tools	You can configure a custom domain name or vanity URL for your APEX apps, Oracle REST Data Services (ORDS), and developer tools by placing an Oracle Cloud Infrastructure Load Balancer directly in front of your Autonomous Database. See the following for additional details, including prerequisites and step-by-step instructions: • Access Oracle APEX, Oracle REST Data Services, and Developer Tools Using a Vanity URL • Access Oracle REST Data Services, Oracle APEX, and Developer Tools Using a Vanity URL

August 2021

Feature	Description
Cross-Region Autonomous Data Guard	Autonomous Data Guard allows you to enable a cross-region standby (peer) database to provide data protection and disaster recovery for your Autonomous Database instance. When you enable Autonomous Data Guard, the system creates a standby database that is continuously updated with the changes from the primary database. You can enable Autonomous Data Guard with a standby in the current region, a local standby, or with a standby in a different region, a cross-region standby. You can also enable Autonomous Data Guard with both a local standby and a cross-region standby. See Using Standby Databases with Autonomous Database for Disaster Recovery for more information.
SQL Tracing for Database Session	Use SQL tracing to help you identify the source of an excessive database workload, such as a high load SQL statement in your application. See Perform SQL Tracing on Autonomous Database for more information.

Feature	Description
Export Data as JSON to Cloud Object Store	Use the procedure `DBMS_CLOUD.EXPORT_DATA` to export data to your Cloud Object Store as JSON, by specifying a query. This export method supports all the Cloud Object Stores supported by Autonomous Database, and you can use an Oracle Cloud Infrastructure resource principal to access your Oracle Cloud Infrastructure Object Store or Amazon Resource Names (ARNs) to access AWS Simple Storage Service (S3). See Move Data to Object Store as JSON Data Using EXPORT_DATA for more information.
Documentation Addition: Load Data from a Source File with Fixed Width Data	The documentation provides an example for loading data from a fixed-width source file to an external table. See Example Loading Data from a Fixed Width File for more information.
Set Patch Level	When you provision or clone an Autonomous Database instance you can select a patch level to apply upcoming patches. There are two patch level options: Regular and Early. After your provision or clone an Autonomous Database instance and set the patch level to Early, maintenance windows for the instance are scheduled and applied one week before instances with the patch level set to Regular. The Early patch level allows you to use and to test upcoming patches before they are applied to all systems. See Set the Patch Level for more information.

July 2021

Feature	Description
View Patch Details	You can view Autonomous Database patch information, including a list of resolved issues. See View Patch Information for more information.
Include Developer Tools Links in Wallet Readme File	The wallet `README` file includes links for Autonomous Database tools and resources, including: Database Actions, Graph Studio, Oracle APEX, Oracle Machine Learning Notebooks, Autonomous Database Service Console, and SODA drivers. See Wallet README File for more information.

June 2021

Feature	Description
Change MEDIUM Service Concurrency Limit from Service Console	If your application requires a customized concurrency limit not available with the predefined services, you can modify the concurrency limit for the MEDIUM service from the Autonomous Database Service Console or using PL/SQL procedures. See Change MEDIUM Service Concurrency Limit for more information.

Feature	Description
Automatic Partitioning	Automatic partitioning analyzes and automates partition creation for tables and indexes of a specified schema to improve performance and manageability in Autonomous Database. Automatic partitioning, when applied, is transparent and does not require user interaction or maintenance. Automatic partitioning does not interfere with existing partitioning strategies and manually partitioned tables are excluded as candidates for automatic partitioning. See Manage Automatic Partitioning on Autonomous Database for more information.
Use Customer-Managed Encryption Keys	Autonomous Database provides two options for Transparent Data Encryption (TDE) to encrypt data in the database: • Oracle-managed encryption keys • Customer-managed encryption keys See Managing Encryption Keys on Autonomous Database for more information.
Autonomous Database RMAN Recovery Catalog	You can use Autonomous Database as a Recovery Manager (RMAN) recovery catalog. A recovery catalog is a database schema that RMAN uses to store metadata about one or more Oracle databases. See Autonomous Database RMAN Recovery Catalog for more information.
Read-Only Mode for Sessions	You can set the Autonomous Database operation mode to read-only for a session. In read-only mode users for the session can only run queries. See Change Autonomous Database Operation Mode for a Session for more information.
Documentation Addition: `CS_RESOURCE_MANAGER` Package	The `CS_RESOURCE_MANAGER` package provides procedures to list and update consumer group parameters, and to revert parameters to default values. See CS_RESOURCE_MANAGER Package for more information.

May 2021

Feature	Description
Transparent Application Continuity (TAC) on Autonomous Database	Autonomous Database provides application continuity features for making connections to the database. You enable Application Continuity on Autonomous Database in one of two configurations: Application Continuity(AC) or Transparent Application Continuity (TAC). Application Continuity hides outages for thin Java-based applications, Oracle Database Oracle Call Interface, and ODP.NET based applications with support for open-source drivers, such as Node.js, and Python. Transparent Application Continuity (TAC) transparently tracks and records session and transactional state so the database session can be recovered following recoverable outages. This is done with no reliance on application knowledge or application code changes. See Using Application Continuity on Autonomous Database for more information.

Feature	Description
Amazon Resource Names (ARNs) to Access AWS Resources	You can use Amazon Resource Names (ARNs) to access AWS resources with Autonomous Database. When you use ARN role based authentication with Autonomous Database, you can securely access AWS resources without creating and saving credentials based on long-term AWS IAM access keys. See Use Amazon Resource Names (ARNs) to Access AWS Resources for more information.
Resource Principal to Access Oracle Cloud Infrastructure Resources	When you use a resource principal with Autonomous Database, you or your tenancy administrator define the Oracle Cloud Infrastructure policies in a dynamic group that allows you to access resources. With a resource principal you do not need to create a credential object and Autonomous Database creates and secures the resource principal credentials you use to access the Oracle Cloud Infrastructure resources specified in the dynamic group. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.
Graph Studio	Graph Studio features include automated modeling to create graphs from database tables, an integrated notebook to run graph queries and analytics, and native graph and other visualizations. You can invoke nearly 60 pre-built graph algorithms and visualize your data with many visualization options. Graph Studio is a fully integrated, automated feature with Autonomous Database. See Using Oracle Graph with Autonomous Database for more information.

April 2021

Feature	Description
Change MEDIUM Service Concurrency Limit	If your application requires a customized concurrency limit not available with the predefined services, you can modify the concurrency limit for the MEDIUM service. See Change MEDIUM Service Concurrency Limit for more information.
DBMS_CLOUD REST API Results Cache	The DBMS_CLOUD REST API functions allow you to make HTTP requests and obtain and save results. Saving results in the cache allows you to view past results with the `SESSION_CLOUD_API_RESULTS` view. Saving and querying historical results of `DBMS_CLOUD` REST API requests can help you when you need to work with previous results in your applications. See DBMS_CLOUD REST API Results Cache for more information.
View and Manage Customer Contacts for Operational Issues and Announcements	When customer contacts are set, Oracle sends notifications to the specified email addresses to inform you of service-related issues. Contacts in the customer contacts list receive unplanned maintenance notices and other notices, including but not limited to notices for database upgrades and upcoming wallet expiration. See View and Manage Customer Contacts for Operational Issues and Announcements for more information.

Feature	Description
Always Free Autonomous JSON Database	You have the option to create a limited number of Always Free Autonomous JSON Databases that do not consume cloud credits. Always Free Autonomous JSON Databases can be created in Oracle Cloud Infrastructure accounts that are in a trial period, have paying status, or are always free. See Always Free Autonomous Database for more information.
Always Free Oracle APEX Application Development	You have the option to create a limited number of Always Free APEX Services that do not consume cloud credits. Always Free APEX Service can be created in Oracle Cloud Infrastructure accounts that are in a trial period, have paying status, or are always free. See Always Free Oracle APEX Application Development for more information.
Expiring Wallets Notification	Starting six weeks before the wallet expiration date Oracle sends notification emails each week, indicating the wallet expiration date. These emails provide notice before your wallet expires that you need to download a new wallet. You can also use the `WalletExpirationWarning` event to be notified when a wallet is due to expire. See Download Client Credentials (Wallets) for more information.
Documentation Addition: Oracle Graph	Oracle Graph with Autonomous Database enables you to create graphs from data in your Autonomous Database. With graphs you can analyze your data based on connections and relationships between data entities. See Using Oracle Graph with Autonomous Database for more information.
Documentation Addition: Oracle Spatial	Oracle Spatial with Autonomous Database allows developers and analysts to get started easily with location intelligence analytics and mapping services. See Using Oracle Spatial with Autonomous Database for more information.
Download Data Pump Dump File Set with Script	To support using Oracle Data Pump Import to import a dump file set to a target database, you can use a script that supports substitution characters to download all the dump files from your Object Store in a single command. See Download Dump Files, Run Data Pump Import, and Clean Up Object Store for more information.

March 2021

Feature	Description
Autonomous Database Tools for Database Actions	Autonomous Database provides the following data tools: • **Data Load**: helps you select data to load from your local computer, from tables in other databases, or from cloud storage. Then you can add the data you select to new or existing tables or views in Autonomous Database. See The Data Load Page for more information. • **Catalog**: helps you to see information about the entities in your Autonomous Database, and to see the effect that changing an object has on other objects. Catalog provides a tool for you to examine data lineage and understand the impact of changes. See The Catalog Page for more information. • **Data Insights**: crawls a table or business model and discovers hidden patterns, anomalies, and outliers in your data. Insights are automatically generated by built-in analytic functions. The results of the insight analysis appear as bar charts in the Data Insights dashboard. See The Data Insights Page for more information. • **Business Models**: describes the business entities that are derived from data in your Autonomous Database schema or from other sources. This allows you to create a semantic model on top of your data identifying hierarchies, measures, and dimensions. See The Business Models Page for more information.
Maintenance History	You can view Autonomous Database maintenance history to see details about past maintenance, such as the start time and stop time for past maintenance events. See About Autonomous Database Maintenance Windows and History for more information.
Oracle Machine Learning AutoML User Interface	AutoML User Interface (AutoML UI) is an Oracle Machine Learning interface that provides you with no-code automated machine learning. When you create and run an experiment in AutoML UI, it automatically performs algorithm and feature selection, as well as model tuning and selection, thereby enhancing productivity as well as model accuracy and performance. See Get Started with AutoML for more information.
Oracle Machine Learning Models	REST endpoints are available so that you can store Machine Learning models and create scoring endpoints for the models. The REST API for Oracle Machine Learning Services supports both Oracle Machine Learning models and ONNX format models. See REST API for Oracle Machine Learning Services for more information.
ADMIN Password Expiration Warning Event	The `AdminPasswordWarning` event produces an event when the Autonomous Database ADMIN password is expiring within 30 days or is expired. See About Events Based Notification and Automation on Autonomous Database for more information.

Feature	Description
Oracle Machine Learning for Python (OML4Py)	Oracle Machine Learning for Python (OML4Py) on Autonomous Database provides a Python API to explore and prepare data, and build and deploy machine learning models using the database as a high-performance compute engine. OML4Py is available through the Python interpreter in Oracle Machine Learning Notebooks. See Get Started with Notebooks for Data Analysis and Data Visualization for more information.

February 2021

Feature	Description
Oracle Database Actions	SQL Developer Web is now called Database Actions. Database Actions is a web-based interface that provides development tools, data tools, administration, and monitoring features for Autonomous Database. Using Database Actions you can load data and run SQL statements, queries, and scripts in a worksheet. See Connect with Built-in Oracle Database Actions for more information.
Query Big Data Service Hadoop (HDFS) Data	Big Data Service provides enterprise-grade Hadoop as a service, with end-to-end security, high performance, and ease of management and upgradeability. After deploying the Oracle Cloud SQL Query Server to Big Data Service, you can easily query data available on Hadoop clusters at scale from Autonomous Database using SQL. See Query Big Data Service Hadoop (HDFS) Data from Autonomous Database for more information.

January 2021

Feature	Description
Autonomous Database Events	Autonomous Database generates events that you can subscribe to with Oracle Cloud Infrastructure Events. There are two new Information events: `ScheduledMaintenanceWarning` and `WalletExpirationWarning`. See Use Autonomous Database Events for more information.
Private Endpoint Support for Tools (Oracle Machine Learning Notebooks are supported with Private Endpoints)	You can access Oracle APEX, Oracle SQL Developer Web, Oracle REST Data Services, and Oracle Machine Learning Notebooks with Private Endpoints. See Configuring Network Access with Private Endpoints for more information.

2020 What's New

Announcements for the noteworthy changes made to Oracle Autonomous Database Serverless Deployment in the year 2020.

- December 2020
- November 2020
- October 2020
- September 2020

- August 2020
- July 2020
- June 2020
- May 2020
- April 2020
- March 2020
- February 2020
- January 2020

December 2020

Feature	Description
Replicate Data and Capture Data with Oracle GoldenGate	Using Oracle GoldenGate you can capture changes from an Oracle Autonomous Database and replicate to any target database or platform that Oracle GoldenGate supports, including another Oracle Autonomous Database. See Oracle GoldenGate Capture for Oracle Autonomous Database for more information.
Oracle APEX Application Development	Oracle APEX Application Development (APEX Service) is a low cost Oracle Cloud service offering convenient access to the Oracle Application Express platform for rapidly building and deploying low-code applications. See Oracle APEX Application Development for more information.
Oracle Database 21c with Always Free Autonomous Database	Oracle Database 21c is available with Always Free Autonomous Database. See the following for more information: • Always Free Autonomous Database • Always Free Autonomous Database Oracle Database 21c Features
Oracle APEX 20.2	Autonomous Database uses Oracle APEX 20.2. See Creating Applications with Oracle Application Express in Autonomous Database for more information.

November 2020

Feature	Description
Amazon S3 Presigned URLs	You can use a presigned URL in any `DBMS_CLOUD` procedure that takes a URL to access files in Amazon Simple Storage Service, without the need to create a credential. See Amazon S3 Compatible URI Format for more information.
Azure Blob Storage Shared Access Signatures (SAS) URLs	You can use Shared Access Signatures (SAS) URL in any `DBMS_CLOUD` procedure that takes a URL to access files in Azure Blob Storage, without the need to create a credential. See Azure Blob Storage URI Format for more information.

October 2020

Feature	Description
User-Defined Profiles Minimum Password Length	If you create a user-defined profile with a Password Verification Function (PVF), then the minimum password length can be 8 characters. See Manage Password Complexity on Autonomous Database for more information.
PL/SQL SDK for Oracle Cloud Infrastructure Resources	PL/SQL SDK for Oracle Cloud Infrastructure enables you to write code to manage Oracle Cloud Infrastructure resources. The PL/SQL SDK is on all Autonomous Database Serverless products. See PL/SQL SDK for more information.
Documentation Addition: Oracle Database Actions User Management	You can quickly create users and modify account settings for users with Oracle Database Actions. See Create Users on Autonomous Database for more information.
Autonomous Database Events	Autonomous Database generates events that you can subscribe to with Oracle Cloud Infrastructure Events. Subscribing to Autonomous Database events allows you to create automation and to receive notifications when the events occur. See Use Autonomous Database Events for more information.

September 2020

Feature	Description
Private Endpoint Support for Tools	You can access Oracle APEX, Oracle SQL Developer Web, and Oracle REST Data Services with Private Endpoints. See Configuring Network Access with Private Endpoints for more information.
Read-Only and Restricted Modes	You can select an Autonomous Database operation mode. The default mode is Read/Write. If you select Read-Only mode users can only run queries. In addition, for either mode you can restrict access to only allow privileged users to connect to the database. See Change Autonomous Database Mode to Read/Write or Read-Only for more information.
Refreshable Clones	Autonomous Database provides cloning where you can choose to create a full clone of the active instance, create a metadata clone, or create a refreshable clone. With a refreshable clone the system creates a clone that can be easily updated with changes from the source database. See Using Refreshable Clones with Autonomous Database for more information.
Documentation Addition: Auditing Autonomous Database	This new chapter describes the auditing features that allow you to track, monitor, and record activity on your database. Auditing can help you detect security risks and improve regulatory compliance for your database. See Auditing Autonomous Database for more information.

August 2020

Feature	Description
Autonomous JSON Database	If your database version is Oracle Database 19c or higher, you can provision an Autonomous JSON Database. Autonomous JSON Database is an Oracle Cloud service that is specialized for developing NoSQL-style applications that use JavaScript Object Notation (JSON) documents. See About Autonomous JSON Database for more information.
Amazon S3-Compatible Object Stores	Autonomous Database supports Amazon S3-Compatible object stores, such as Oracle Cloud Infrastructure Object Storage, Google Cloud Storage, and Wasabi Hot Cloud Storage. See Load Data from Files in the Cloud for more information.
Asynchronous Requests using `SEND_REQUEST`	The `DBMS_CLOUD SEND_REQUEST` function supports long running requests with the additional parameters: `async_request_url`, `wait_for_states`, and `timeout`. See SEND_REQUEST Function for more information.
Rename Database	If your database version is Oracle Database 19c or higher, you can rename your Autonomous Database. See Rename Autonomous Database for more information.

July 2020

Feature	Description
Oracle APEX Release 20.1	Autonomous Database supports Oracle APEX Release 20.1. See Creating Applications with Oracle Application Express in Autonomous Database for more information. If your Autonomous Database instance is on Oracle Database 18c, to use APEX 20.1 you must upgrade to Oracle Database 19c. After you upgrade to Oracle Database 19c, APEX is automatically upgraded to APEX 20.1.
Use a Standby Database	Autonomous Database provides the Autonomous Data Guard feature to enable a standby (peer) database to provide data protection and disaster recovery for your Autonomous Database instance. See Using a Standby Database with Autonomous Database for more information.
Create and Alter User Profiles	You can create and alter database user profiles. See Manage User Profiles with Autonomous Database for more information.
Manage Time Zone File Versions	The time zone files are periodically updated to reflect the latest time zone specific changes. Autonomous Database automatically picks up the updated time zone files. See Manage Time Zone File Version on Autonomous Database for more information.

Feature	Description
Network Access Changes	If your Autonomous Database network access is configured to use a private endpoint you can change the configuration to use a public endpoint. Likewise, if your Autonomous Database instance is configured to use a public endpoint, you can change the configuration to use a private endpoint. Autonomous Database network access changes from a public to a private endpoint, or from a private to a public endpoint are only supported with database versions Oracle Database 19c onwards. See Configuring Network Access with Access Control Rules (ACLs) and Private Endpoints for more information.
Use Database Links to Access Non-Oracle Databases	You can create database links from an Autonomous Database to an Oracle Database Gateway to access Non-Oracle databases. See Create Database Links to an Oracle Database Gateway to Access Non-Oracle Databases for more information.

June 2020

Feature	Description
ORC Format and Complex Types	Autonomous Database supports loading and querying data in object store in ORC format, in addition to Avro and Parquet. Also, for ORC, Avro, and Parquet structured file types you can load and query complex data types. Support for ORC format and complex types requires Oracle Database 19c. See Load Data from Files in the Cloud and Query External Data with ORC, Parquet, or Avro Source Files for more information.
`PUT_OBJECT` Maximum File Transfer Size Increase	The `DBMS_CLOUD.PUT_OBJECT` procedure maximum size limit for file transfers to Oracle Cloud Infrastructure Object Storage is increased to 50 GB. See PUT_OBJECT Procedure for more information.
Customer Managed Oracle REST Data Services	You can use a customer managed environment to run Oracle REST Data Services (ORDS) if you want to control the configuration and management of ORDS. Installing and configuring a customer managed environment for ORDS allows you to run ORDS with configuration options that are not possible using the default Oracle managed ORDS available with Autonomous Database. See About Customer Managed Oracle REST Data Services on Autonomous Database for more information.

May 2020

Feature	Description
Move Selective Data to Object Store	You can export data to Oracle Data Pump dump files on Object Store by specifying a query to select the data to export. See Move Data to Object Store as Oracle Data Pump Files Using EXPORT_DATA for more information.

Feature	Description
Obtain Tenancy Details for a Service Request	When you file a service request for Autonomous Database, you need to provide the tenancy details for your instance. Tenancy details for the instance are available on the Oracle Cloud Infrastructure console or you can obtain these details by querying the database. See Obtain Tenancy Details to File a Service Request for more information.
Upgrade Autonomous Database to Oracle Database 19c	If your Autonomous Database instance is on Oracle Database 18c, you can upgrade to Oracle Database 19c by clicking **Upgrade to 19c**.

April 2020

Feature	Description
SODA Documents and Collections	Autonomous Database supports loading and using Simple Oracle Document Architecture (SODA) documents and collections. SODA allows you to store, search, and retrieve document collections, typically JSON documents, without using SQL. You can also access document collections with SQL/JSON operators, providing you with the simplicity of a document database but with the power of a relational database. See Using JSON Documents with Autonomous Database for more information.
Load Data with Oracle Data Pump Access Driver Dump Files	You can use Oracle Data Pump dump files in the Cloud as source files to load your data. The files for this load type must be exported from the source system using the `ORACLE_DATAPUMP` access driver in External Tables. See Create Credentials and Load Data Pump Dump Files into an Existing Table for more information.
Create External Tables and Query Data with Oracle Data Pump Access Driver Dump Files	You can query Oracle Data Pump dump files in the Cloud by creating an external table. The source files to create this type of external table must be exported from the source system using the `ORACLE_DATAPUMP` access driver in External Tables. See Query External Data Pump Dump Files for more information.
Oracle Extensions for IDEs	You can use Oracle extensions for IDEs to develop applications on Autonomous Database. Extensions are available for Eclipse, Microsoft Visual Studio, and Microsoft Visual Studio Code (VS Code). These extensions enable you to connect to, browse, and manage Autonomous Databases in Oracle Cloud directly from the IDE. See Use Oracle Extensions for IDEs to Develop Applications for more information.
Per-Second Billing	Autonomous Database instance CPU and storage usage is billed by the second, with a minimum usage period of 1 minute. Previously Autonomous Database billed in one-hour minimum increments and partial usage of an hour was rounded up to a full hour.
Stateful Rule Support in Private Endpoints	When you configure a private endpoint you define a security rule(s) in a Network Security Group (NSG); this creates a virtual firewall for your Autonomous Database and allows connections to the Autonomous Database instance. The NSG can now be configured with stateful rules. See Configure Private Endpoints with Autonomous Database for more information.

Feature	Description
Wallet zip File Contains `README` File	The wallet file that contains client credentials (wallet files) for your Autonomous Database instance now also has a `README` file with wallet expiration information. Wallet files that were downloaded prior to April 2020 do not contain this file. See Download Client Credentials (Wallets) for more information.

March 2020

Feature	Description
Restart Autonomous Database instance	Use **Restart** to restart an Autonomous Database instance. See Restart Autonomous Database for more information.
Oracle Data Pump Direct Export to Object Store	Depending on your cloud object store, you can use Oracle Data Pump to directly export to your object store to move data between Autonomous Database and other Oracle databases. See Move Data with Data Pump Export to Object Store for more information.
Oracle Data Pump Pre-authenticated URLs for Dump Files	If your source files reside on Oracle Cloud Infrastructure Object Storage, you can use pre-authenticated URLs with Oracle Data Pump. See Import Data Using Oracle Data Pump on Autonomous Database for more information.
Documentation Addition: Upgrade Autonomous Database to Oracle Database 19c	Provides information on upgrading to Oracle Database 19c if your Autonomous Database instance is on Oracle Database 18c.
Oracle Database Versions by Region	Depending on the region where you provision or clone your database, Autonomous Database supports one or more Oracle Database versions. Oracle Database 19c is available in all commercial regions.
Oracle Application Express Release 19.2	Autonomous Database supports Oracle Application Express Release 19.2. See Creating Applications with Oracle Application Express in Autonomous Database for more information.

February 2020

Feature	Description
Private Endpoints	This configuration option assigns a private endpoint, private IP and hostname, to a database and allows traffic only from the VCN you specify; access to the database from all public IPs or VCNs is blocked. This allows you to define security rules at the Network Security Group (NSG) level and to control traffic to your Autonomous Database. See About Network Access Options for more information.
Oracle Database Vault	Oracle Database Vault implements powerful security controls for your database. These unique security controls restrict access to application data by privileged database users, reducing the risk of insider and outside threats and addressing common compliance requirements. See Using Oracle Database Vault with Autonomous Database for more information.

Feature	Description
Oracle Application Express Release 19.2	If your database instance is an Always Free Autonomous Database or is a database with Oracle Database release 19c, then your database has Oracle Application Express Release 19.2. See Creating Applications with Oracle Application Express in Autonomous Database for more information.
Database Resident Connection Pool (DRCP)	Using DRCP provides you with access to a connection pool in your database that enables a significant reduction in key database resources required to support many client connections. See Use Database Resident Connection Pooling with Autonomous Database for more information.

January 2020

Feature	Description
Clone from a backup	Clone from a backup lets you create a clone by selecting a backup from a list of backups, or with a timestamp for a point-in-time to clone. See Clone Autonomous Database from a Backup for more information.
Check and set `MAX_STRING_SIZE` value	By default the database uses extended data types and the value of `MAX_STRING_SIZE` is set to the value `EXTENDED`. To support migration from older Oracle Databases or applications you can set `MAX_STRING_SIZE` to the value `STANDARD`. See Checking and Setting MAX_STRING_SIZE for more information.
Number of concurrent statements increased depending on the service you use	The maximum number of concurrent statements is increased, depending on the connection service you are using. See Predefined Database Service Names for Autonomous Database for more information.
Documentation addition showing available Oracle Database versions	The documentation shows the Oracle Database version available with Autonomous Database by region. See Oracle Database Version Availability by Region for more information.
`DBMS_CLOUD` REST API functions	The `DBMS_CLOUD` REST API functions allow you to make HTTP requests using `DBMS_CLOUD.SEND_REQUEST`. These functions provide a generic API that lets you call REST APIs from supported cloud services. See DBMS_CLOUD REST APIs for more information.
Documentation addition to describe sending mail with `UTL_SMTP`	The documentation is updated to provide the steps and sample code for sending email using `UTL_SMTP`. See Sending Mail with Email Delivery on Autonomous Database for more information.
Use ACLs to control access for Autonomous Database tools	You can now use Virtual Cloud Network, Virtual Cloud Network (OCID), IP address, or CIDR block ACLs to control access to Oracle APEX (APEX), RESTful services, and SQL Developer Web. See Use an Access Control List with Autonomous Database for more information.

2019 What's New

Announcements for the noteworthy changes made to Oracle Autonomous Database Serverless Deployment in the year 2019.

- December 2019
- November 2019
- October 2019
- September 2019
- August 2019
- July 2019
- June 2019
- May 2019
- April 2019
- March 2019
- February 2019
- January 2019

December 2019

Feature	Description
Support for `UTL_HTTP` and `UTL_SMTP` PL/SQL Packages with restrictions	Autonomous Database supports the Oracle Database PL/SQL packages `UTL_HTTP`, `UTL_SMTP`, and `DBMS_NETWORK_ACL_ADMIN` with restrictions. See Restrictions for Database PL/SQL Packages for more information.
Microsoft Active Directory Users	You can configure Autonomous Database to authenticate and authorize Microsoft Active Directory users. This allows Active Directory users to access a database using their Active Directory credentials. See Use Microsoft Active Directory with Autonomous Database for more information.
Validate external partitioned tables and hybrid partitioned tables	You can validate individual partitions of an external partitioned table or a hybrid partitioned table with the procedures `DBMS_CLOUD.VALIDATE_EXTERNAL_PART_TABLE` and `DBMS_CLOUD.VALIDATE_HYBRID_PART_TABLE`. See Validate External Partitioned Data and Validate Hybrid Partitioned Data for more information.
Use ACL IP address or CIDR block settings to control access for Autonomous Database tools	You can now use IP address and CIDR based ACLs to control access to Oracle APEX (APEX), RESTful services, and SQL Developer Web. See Use an Access Control List with Autonomous Database for more information.
Use Oracle Data Safe	Oracle Data Safe provides features that help you protect sensitive and regulated data in your database. See Safeguard Your Data with Data Safe on Autonomous Database for more information.

November 2019

Feature	Description
Use ACLs to specify Oracle Cloud Infrastructure VCNs	Use ACLs to specify Oracle Cloud Infrastructure VCNs that can connect to your Autonomous Database. See Use an Access Control List with Autonomous Database for more information.
SQL Developer Web Data Loading	In SQL Developer Web in the Worksheet page, you can upload data from local files into an existing table. See Load Data from Local Files with SQL Developer Web for more information.
Access APEX from Oracle Cloud Infrastructure Console	The Tools Tab provides access to APEX from Oracle Cloud Infrastructure Console. See Access Oracle Application Express Administration Services for more information.
Access SQL Developer Web from Oracle Cloud Infrastructure Console	The Tools Tab provides access to SQL Developer Web from Oracle Cloud Infrastructure Console. See Access SQL Developer Web as ADMIN for more information.
Access Oracle Machine Learning from Oracle Cloud Infrastructure Console	The Tools Tab provides access to Oracle Machine Learning User Administration from Oracle Cloud Infrastructure Console. See Work with Oracle Machine Learning for Data Access, Analysis, and Discovery for more information.
View Maintenance Schedule from Oracle Cloud Infrastructure Console	The Autonomous Database Information tab shows the schedule for upcoming maintenance. See About Autonomous Database Maintenance Windows for more information.
New fields with `LIST_FILES` and `LIST_OBJECTS`	The functions `DBMS_CLOUD.LIST_FILES` and `DBMS_CLOUD.LIST_OBJECTS` produce additional metadata for files and objects. See LIST_FILES Function and LIST_OBJECTS Function for more information.

October 2019

Feature	Description
Preview Version for Autonomous Database	Oracle periodically provides a preview version of Autonomous Database that allows you to test your applications and to become familiar with features in the next release of Autonomous Database. See Preview Versions for Autonomous Database for more information.
Download database specific instance wallets or regional wallets	You can download database specific instance wallets or regional wallets. See Download Client Credentials (Wallets) for more information.
Rotate database specific instance wallets or regional wallets	You can rotate database specific instance wallets or regional wallets. See Rotate Wallets for Autonomous Database for more information.
Create partitioned external tables with `DBMS_CLOUD`	You can create and validate external partitioned tables with the `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` and `DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE` statements. See Query External Partitioned Data for more information.

Feature	Description
Numeric Formats	With the `numberformat` and `numericcharacters` format options Autonomous Database supports formats to interpret numeric strings correctly. See DBMS_CLOUD Package Format Options for more information.

September 2019

Feature	Description
Always Free Autonomous Database	You have the option to create a limited number of Always Free Autonomous Databases that do not consume cloud credits. Always Free databases can be created in Oracle Cloud Infrastructure accounts that are in a trial period, have paying status, or are always free. See Always Free Autonomous Database for more information.
Autonomous Database Metrics	You can monitor the health, capacity, and performance of your databases with metrics, alarms, and, notifications. You can use Oracle Cloud Infrastructure console or Monitoring APIs to view metrics. See Monitor Performance with Autonomous Database Metrics for more information.
Object Store Using Public URL	If your source files reside on an Object Store that provides public URLs, you can use public URLs with `DBMS_CLOUD` procedures. Public means the Object Storage service supports anonymous, unauthenticated access to the Object Store files. See URI Format Using Public URL for more information.
Work Requests	Work Requests are available in the Oracle Cloud Infrastructure console. Work Requests let you track progress of database lifecycle management operations such as creating, terminating, backing up (manual), restoring, scaling, and cloning an Autonomous Database. Work requests allow you to track the progress and steps completed in a database operation. See About Work Requests for more information.

August 2019

Feature	Description
Create Directory and Drop Directory commands	The `data_pump_dir` directory is available in a database. You can use `CREATE DIRECTORY` to create additional directories. To drop a directory, you can use the `DROP DIRECTORY` command. See Creating and Managing Directories for more information.

July 2019

Feature	Description
Performance Hub	Oracle Cloud Infrastructure console includes Performance Hub for Autonomous Database. You can view real-time and historical performance data from the Performance Hub. See Monitor Autonomous Database with Performance Hub for more information.

Feature	Description
Oracle Cloud Infrastructure Native Object Store Authentication	Autonomous Database supports native authentication with Oracle Cloud Infrastructure Object Store. Using native authentication you can use key based authentication to access the Object Store (instead of using a username and password). See CREATE_CREDENTIAL Procedure for more information.
Move to a different Compartment	You can move a database to a different Oracle Cloud Infrastructure compartment. See Move an Autonomous Database to a Different Compartment for more information.

June 2019

Feature	Description
Application Express (APEX)	Oracle APEX (APEX) is a low-code development platform that enables you to build scalable, secure enterprise applications with world-class features that can be deployed anywhere. Each Autonomous Database instance includes a dedicated instance of Oracle APEX; you can use this instance to create multiple workspaces. A workspace is a shared work area where you can build applications. See About Oracle Application Express for more information.
SQL Developer Web	Oracle SQL Developer Web provides a provides a development environment and a data modeler interface for Autonomous Database. See About SQL Developer Web for more information.
Oracle REST Data Services (ORDS)	You can develop and deploy RESTful Services with native Oracle REST Data Services (ORDS) support on a database See Developing RESTful Services in Autonomous Database for more information.
Auto Scaling	Enabling auto scaling allows a database to use up to three times more CPU and IO resources than the currently specified CPU Core Count. When auto scaling is enabled, if your workload requires additional CPU and IO resources the database automatically uses the resources without any manual intervention required. See Use Auto Scaling for more information.
Azure object store with Data Pump	Autonomous Database now supports Microsoft Azure cloud storage with Oracle Data Pump. See Import Data Using Oracle Data Pump on Autonomous Database for more information.

May 2019

Feature	Description
Support for creating database links	You can create database links with `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` to access objects on another database. See Use Database Links with Autonomous Database for more information.
Oracle Spatial and Graph with limitations	

Feature	Description
Oracle Text with limitations	See Restrictions for Oracle Text for more information.
Oracle Cloud Infrastructure Native URI Format Supported with Oracle Data Pump	For importing data from Data Pump files using `impdp`, if your source files reside on the Oracle Cloud Infrastructure Object Storage you can use the Oracle Cloud Infrastructure native URIs in addition to the Swift URIs. See DBMS_CLOUD Package File URI Formats for more information.
Oracle Cloud Infrastructure Pre-Authenticated URI Format Supported with `DBMS_CLOUD`	If your source files reside on the Oracle Cloud Infrastructure Object Storage you can use the Oracle Cloud Infrastructure Pre-Authenticated URIs. When you create a pre-authenticated request, a unique URL is generated. You can use a pre-authenticated URL in any `DBMS_CLOUD` procedure that takes a URL to access files in the Oracle Cloud Infrastructure object store, without the need to create a credential. See DBMS_CLOUD Package File URI Formats for more information.
Oracle XML DB with limitations	Oracle XML DB is a high-performance, native XML storage and retrieval technology that is delivered as a part of Oracle Database. Limitations apply to Oracle XML DB with Autonomous Database database. See Restrictions for Oracle XML DB for more information.
Modify CPU/IO Shares from Service Console	Autonomous Database comes with predefined CPU/IO shares assigned to different consumer groups. You can modify predefined CPU/IO shares if your workload requires different CPU/IO resource allocations.You can change CPU/IO shares from the Service Console or using a PL/SQL package. See Manage CPU/IO Shares on Autonomous Database for more information.

April 2019

Feature	Description
Support for Network Access Control Lists	Autonomous Database now supports setting a network Access Control List (ACL) to restrict access to a specific Autonomous Database database. When you set a network ACL the database only accepts connections from addresses specified on the ACL and rejects all other client connections. See Configuring Network Access with Access Control Rules (ACLs) for more information.
Support for Updating License Type	You can now update your license type from the Oracle Cloud Infrastructure console **Actions** list. See Update License Type on Autonomous Database for more information.
Access Avro files in Object Stores	Autonomous Database allows you to directly query and load data stored as Apache Avro format files (Apache Parquet format files are also supported). You can create external tables for Parquet or Avro format data files. See DBMS_CLOUD Package Format Options for Parquet and Avro and CREATE_EXTERNAL_TABLE Procedure for Parquet or Avro Files for more information.

Feature	Description
Enhanced newline handling for file record delimiter with `DBMS_CLOUD`	Now, when reading a file, by default `DBMS_CLOUD` tries to automatically find the correct newline character to use as the record delimiter, either for Windows, newline character "`\r\n`", or for UNIX/Linux newline character "`\n`". If `DBMS_CLOUD` finds one of these it sets the record delimiter for the file. You can specify the record delimiter explicitly if you want to override the default behavior. See DBMS_CLOUD Package Format Options for more information.
Modify CPU/IO Shares	Autonomous Database comes with predefined CPU/IO shares assigned to different consumer groups. You can modify these predefined CPU/IO shares if your workload requires different CPU/IO resource allocations. See Manage CPU/IO Shares on Autonomous Database for more information.

March 2019

Feature	Description
Clone Your Database	Autonomous Database provides cloning where you can choose to clone either the full database or only the database metadata. See Cloning a Database with Autonomous Database for more information.
Oracle Cloud Infrastructure Native URI Format Supported with `DBMS_CLOUD`	If your source files reside on the Oracle Cloud Infrastructure Object Storage you can use the Oracle Cloud Infrastructure native URIs in addition to the Swift URIs. See DBMS_CLOUD Package File URI Formats for more information.
UI/API Unification and Workload Type field	Oracle Cloud Infrastructure Console and the Oracle Cloud Infrastructure APIs for Autonomous Data Warehouse and Autonomous Transaction Processing are converged to a single, unified framework. These changes allow you to more easily manage both types of Autonomous Databases. There is a new Workload Type field on the console showing values of either “Transaction Processing” or “Data Warehouse”, depending on the type of database you are viewing. You can also select a Workload Type when you provision an Autonomous Database. See Workload Types for more information.
Service Gateway	You can now set up a Service Gateway with Autonomous Database. A service gateway allows connectivity to Autonomous Database from private IP addresses in private subnets without requiring an Internet Gateway in your VCN. See Access Autonomous Database with Service Gateway and Access to Oracle Services: Service Gateway for more information.
Documentation changes	The procedure `DBMS_CLOUD.DELETE_ALL_OPERATIONS` is now documented. See DELETE_ALL_OPERATIONS Procedure for more information.

February 2019

Feature	Description
Application Continuity	You can now enable and disable Application Continuity. Application Continuity masks outages from end users and applications by recovering the in-flight work for impacted database sessions following outages. Application Continuity performs this recovery beneath the application so that the outage appears to the application as a slightly delayed execution. See Using Application Continuity on Autonomous Database for more information.
Documentation changes	Autonomous Database documentation includes a new chapter that describes moving data to other Oracle databases. See Moving Data from Autonomous Database to Other Oracle Databases for more information.

January 2019

Feature	Description
Access Parquet files in Object Stores	Autonomous Database allows you to directly query and load data stored as parquet files in object stores. You can also create external parquet format tables in object stores. See COPY_DATA Procedure for Parquet Files and CREATE_EXTERNAL_TABLE Procedure for Parquet Files for more information.

2018 What's New

Announcements for the noteworthy changes made to Oracle Autonomous Database Serverless Deployment in the year 2018.

- October 2018
- September 2018
- August 2018
- July 2018
- June 2018
- May 2018
- April 2018

October 2018

Feature	Description
Oracle Cloud Infrastructure Console changes	The Autonomous Database console has a new layout and provides new buttons that improve the usability of Autonomous Database. These changes include a new **DB Connection** button that makes it easier to download client credentials.

September 2018

Feature	Description
Table compression methods	In addition to Hybrid Columnar Compression all table compression types are now available in Autonomous Database. For more information, see Managing DML Performance and Compression.
Idle timeout changes	The idle timeout of 60 minutes is now lifted. Idle sessions that do not hold resources required by other sessions will not be terminated after 60 minutes now. For more information, see Manage Concurrency and Priorities on Autonomous Database.

August 2018

Feature	Description
The Oracle Cloud Infrastructure page has a new option Autonomous Transaction Processing	The sign in for Oracle Cloud Infrastructure now lists the new product, Oracle Autonomous Transaction Processing, in addition to Oracle Autonomous Database.

July 2018

Feature	Description
SQL Developer 18.2.0 and later without Keystore Password field for connections	When you connect to Autonomous Database with SQL Developer 18.2.0 or later, you do not need to supply a keystore password. The keystore password was required in previous SQL Developer versions.

June 2018

Feature	Description
New management interfaces	Autonomous Database is now provisioned and managed using the native Oracle Cloud Infrastructure. This provides a more intuitive user interface to make managing your Autonomous Database instances easier with additional capabilities including sorting and filtering. For more information, see Oracle Help Center.
Better authorization management	With OCI Identity and Access Management (IAM) you can better organize and isolate your Autonomous Database instances using Compartments and control what type of access a user or group of users have.
Built-in auditing	The OCI Audit service records use of Autonomous Database application programming interface (API) endpoints as log events in JSON format. You can view, copy, and analyze audit events with standard log analysis tools or using the Audit service Console, the Audit APIs, or the Java SDK.
Availability of the Phoenix region	Autonomous Database is now available in the Phoenix region in addition to Ashburn and Frankfurt regions.
User Assistance changes	The documentation, videos, and examples are updated on the Oracle Help Center to include the procedures to create and control Autonomous Database instances with the Oracle Cloud Infrastructure Console. A new Related Resources page shows related resources including the Autonomous Data Warehouse Forum .
Idle timeouts in database services	The idle timeouts for three database services, high, medium, and low, have now been relaxed. The previous idle timeout setting of 5 minutes will now apply to only idle sessions that hold resources needed by other active users. See Manage Concurrency and Priorities on Autonomous Database.

May 2018

Feature	Description
Oracle Cloud Infrastructure Object Storage Credentials	The name Swift Password is now Auth token. Use of Swift password in the documentation is now replaced with Auth token.

April 2018

Feature	Description
Microsoft Azure Blob Storage or Azure Data Lake Storage integration	Autonomous Database now supports Microsoft Azure Blob Storage or Azure Data Lake Storage for data loading and querying external data. You can load data from or run queries on files residing on Azure Blob Storage. See Loading Data from Files in the Cloud.

Feature	Description
Manage files on the local file system	You can now list files residing on the local file system on Autonomous Database, see LIST_FILES Function. You can also remove files from the local file system. See DELETE_FILE Procedure.

Sample Star Schema Benchmark (SSB) Queries and Analytic Views

The SSB schema contains the tables: lineorder, customer, supplier, part, and dwdate. The following is a list of sample queries and analytic views you can use against the SSB schema. Note that you need to prefix the table names with the schema name SSB in your queries.

> **Note:**
>
> Both SH and SSB are provided as schema-only users, so you cannot unlock or drop those users or set a password. And the storage of the sample data sets does not count towards your database storage.

- Star Schema Benchmark Queries
- Star Schema Benchmark Analytic Views

Star Schema Benchmark Queries

Star Schema Benchmark Queries

```
select sum(lo_extendedprice*lo_discount) as revenue
from ssb.lineorder, ssb.dwdate
where lo_orderdate = d_datekey
and d_yearmonthnum = 199401
and lo_discount between 4 and 6
and lo_quantity between 26 and 35;

select sum(lo_extendedprice*lo_discount) as revenue
from ssb.lineorder, ssb.dwdate
where lo_orderdate = d_datekey
and d_year = 1993
and lo_discount between 1 and 3
and lo_quantity < 25;

select sum(lo_extendedprice*lo_discount) as revenue
from ssb.lineorder, ssb.dwdate
where lo_orderdate = d_datekey
and d_yearmonthnum = 199401
and lo_discount between 4 and 6
```

```
and lo_quantity between 26 and 35;

select sum(lo_extendedprice*lo_discount) as revenue
from ssb.lineorder, ssb.dwdate
where lo_orderdate = d_datekey
and d_weeknuminyear = 6
and d_year = 1994
and lo_discount between 5 and 7
and lo_quantity between 26 and 35;

select sum(lo_revenue), d_year, p_brand1
from ssb.lineorder, ssb.dwdate, ssb.part, ssb.supplier
where lo_orderdate = d_datekey
and lo_partkey = p_partkey
and lo_suppkey = s_suppkey
and p_category = 'MFGR#12'
and s_region = 'AMERICA'
group by d_year, p_brand1
order by d_year, p_brand1;

select sum(lo_revenue), d_year, p_brand1
from ssb.lineorder, ssb.dwdate, ssb.part, ssb.supplier
where lo_orderdate = d_datekey
and lo_partkey = p_partkey
and lo_suppkey = s_suppkey
and p_brand1 between 'MFGR#2221' and 'MFGR#2228'
and s_region = 'ASIA'
group by d_year, p_brand1
order by d_year, p_brand1;

select sum(lo_revenue), d_year, p_brand1
from ssb.lineorder, ssb.dwdate, ssb.part, ssb.supplier
where lo_orderdate = d_datekey
and lo_partkey = p_partkey
and lo_suppkey = s_suppkey
and p_brand1 = 'MFGR#2221'
and s_region = 'EUROPE'
group by d_year, p_brand1
order by d_year, p_brand1;

select c_nation, s_nation, d_year, sum(lo_revenue) as revenue
from ssb.customer, ssb.lineorder, ssb.supplier, ssb.dwdate
where lo_custkey = c_custkey
and lo_suppkey = s_suppkey
and lo_orderdate = d_datekey
and c_region = 'ASIA' and s_region = 'ASIA'
and d_year >= 1992 and d_year <= 1997
group by c_nation, s_nation, d_year
order by d_year asc, revenue desc;

select c_city, s_city, d_year, sum(lo_revenue) as revenue
from ssb.customer, ssb.lineorder, ssb.supplier, ssb.dwdate
where lo_custkey = c_custkey
and lo_suppkey = s_suppkey
and lo_orderdate = d_datekey
and c_nation = 'UNITED STATES'
```

```
and s_nation = 'UNITED STATES'
and d_year >= 1992 and d_year <= 1997
group by c_city, s_city, d_year
order by d_year asc, revenue desc;

select c_city, s_city, d_year, sum(lo_revenue) as revenue
from ssb.customer, ssb.lineorder, ssb.supplier, ssb.dwdate
where lo_custkey = c_custkey
and lo_suppkey = s_suppkey
and lo_orderdate = d_datekey
and (c_city='UNITED KI1' or c_city='UNITED KI5')
and (s_city='UNITED KI1' or s_city='UNITED KI5')
and d_year >= 1992 and d_year <= 1997
group by c_city, s_city, d_year
order by d_year asc, revenue desc;

select c_city, s_city, d_year, sum(lo_revenue) as revenue
from ssb.customer, ssb.lineorder, ssb.supplier, ssb.dwdate
where lo_custkey = c_custkey
and lo_suppkey = s_suppkey
and lo_orderdate = d_datekey
and (c_city='UNITED KI1' or c_city='UNITED KI5')
and (s_city='UNITED KI1' or s_city='UNITED KI5')
and d_yearmonth = 'Dec1997'
group by c_city, s_city, d_year
order by d_year asc, revenue desc;

select d_year, c_nation, sum(lo_revenue - lo_supplycost) as profit
from ssb.dwdate, ssb.customer, ssb.supplier, ssb.part, ssb.lineorder
where lo_custkey = c_custkey
 and lo_suppkey = s_suppkey
 and lo_partkey = p_partkey
 and lo_orderdate = d_datekey
 and c_region = 'AMERICA'
 and s_region = 'AMERICA'
 and (p_mfgr = 'MFGR#1' or p_mfgr = 'MFGR#2')
group by d_year, c_nation
order by d_year, c_nation;

select d_year, s_nation, p_category, sum(lo_revenue - lo_supplycost) as profit
from ssb.dwdate, ssb.customer, ssb.supplier, ssb.part, ssb.lineorder
where lo_custkey = c_custkey
and lo_suppkey = s_suppkey
and lo_partkey = p_partkey
and lo_orderdate = d_datekey
and c_region = 'AMERICA'
and s_region = 'AMERICA'
and (d_year = 1997 or d_year = 1998)
and (p_mfgr = 'MFGR#1'
or p_mfgr = 'MFGR#2')
group by d_year, s_nation, p_category order by d_year, s_nation, p_category;

select d_year, s_city, p_brand1, sum(lo_revenue - lo_supplycost) as profit
from ssb.dwdate, ssb.customer, ssb.supplier, ssb.part, ssb.lineorder
where lo_custkey = c_custkey
and lo_suppkey = s_suppkey
```

```
and lo_partkey = p_partkey
and lo_orderdate = d_datekey
and c_region = 'AMERICA'
and s_nation = 'UNITED STATES'
and (d_year = 1997 or d_year = 1998)
and p_category = 'MFGR#14'
group by d_year, s_city, p_brand1 order by d_year, s_city, p_brand1;
```

Star Schema Benchmark Analytic Views

SSB Analytic Views

Analytic views make it easy to extend a star schema with a hierarchical business model, aggregation and measure calculation rules, presentation and application-specific metadata that can be used to enhance the content of a data set and to simplify the development of business intelligence applications. The SSB schema includes an analytic view and four hierarchies that use the tables of the star schema. Use the following queries to query the analytic Sample SSB view. Note that the analytic view is in the SSB schema.

```
SELECT
  dwdate_hier.member_name as year,
  part_hier.member_name as part,
  customer_hier.c_region,
  customer_hier.member_name as customer,
  lo_quantity,
  lo_revenue
FROM  ssb.ssb_av
  HIERARCHIES (
    dwdate_hier,
    part_hier,
    customer_hier)
WHERE
  dwdate_hier.d_year = '1998'
  AND dwdate_hier.level_name = 'MONTH'
  AND part_hier.level_name = 'MANUFACTURER'
  AND customer_hier.c_region = 'AMERICA'
  AND customer_hier.level_name = 'NATION'
ORDER BY
  dwdate_hier.hier_order,
  part_hier.hier_order,
  customer_hier.hier_order;

SELECT
  dwdate_hier.member_name as time,
  part_hier.member_name as part,
  customer_hier.member_name as customer,
  supplier_hier.member_name as supplier,
  lo_quantity,
  lo_supplycost
FROM  ssb.ssb_av
  HIERARCHIES (
    dwdate_hier,
    part_hier,
    customer_hier,
```

```
    supplier_hier)
WHERE
  dwdate_hier.d_year = '1998'
  AND dwdate_hier.level_name = 'MONTH'
  AND part_hier.level_name = 'MANUFACTURER'
  AND customer_hier.c_region = 'AMERICA'
  AND customer_hier.c_nation = 'CANADA'
  AND customer_hier.level_name = 'CITY'
  AND supplier_hier.s_region = 'ASIA'
  AND supplier_hier.level_name = 'REGION'
ORDER BY
  dwdate_hier.hier_order,
  part_hier.hier_order,
  customer_hier.hier_order,
  supplier_hier.hier_order;

SELECT
  dwdate_hier.member_name as year,
  part_hier.member_name as part,
  customer_hier.member_name as customer,
  supplier_hier.member_name as supplier,
  lo_quantity,
  lo_revenue,
  lo_supplycost
FROM  ssb.ssb_av
  HIERARCHIES (
    dwdate_hier,
    part_hier,
    customer_hier,
    supplier_hier)
WHERE
  dwdate_hier.d_yearmonth = 'Apr1998'
  AND dwdate_hier.level_name = 'DAY'
  AND part_hier.level_name = 'MANUFACTURER'
  AND customer_hier.c_region = 'AMERICA'
  AND customer_hier.c_nation = 'CANADA'
  AND customer_hier.level_name = 'CITY'
  AND supplier_hier.level_name = 'REGION'
ORDER BY
  dwdate_hier.hier_order,
  part_hier.hier_order,
  customer_hier.hier_order,
  supplier_hier.hier_order;

SELECT
  dwdate_hier.member_name as year,
  part_hier.member_name as part,
  supplier_hier.member_name as supplier,
  lo_quantity,
  lo_extendedprice,
  lo_ordtotalprice,
  lo_revenue,
  lo_supplycost
FROM  ssb.ssb_av
  HIERARCHIES (
    dwdate_hier,
```

```
    part_hier,
    supplier_hier)
WHERE
  dwdate_hier.level_name = 'YEAR'
  AND part_hier.level_name = 'MANUFACTURER'
  AND supplier_hier.level_name = 'SUPPLIER'
  AND supplier_hier.s_suppkey = '23997';

SELECT
  dwdate_hier.member_name as time,
  part_hier.p_container,
  part_hier.member_name as part,
  lo_quantity,
  lo_extendedprice,
  lo_ordtotalprice,
  lo_revenue,
  lo_supplycost
FROM  ssb.ssb_av
  HIERARCHIES (
    dwdate_hier,
    part_hier)
WHERE
  dwdate_hier.member_name = 'June 10, 1998     '
  AND dwdate_hier.level_name = 'DAY'
  AND part_hier.level_name = 'PART'
  AND part_hier.p_size = 32;

SELECT
  dwdate_hier.member_name as time,
  part_hier.member_name as part,
  part_hier.p_name,
  part_hier.p_color,
  lo_quantity,
  lo_revenue,
  lo_supplycost,
  lo_revenue - lo_supplycost as profit
FROM  ssb.ssb_av
  HIERARCHIES (
    dwdate_hier,
    part_hier)
WHERE
  dwdate_hier.d_yearmonth = 'Aug1996'
  AND dwdate_hier.d_dayofweek = 'Friday   '
  AND dwdate_hier.level_name = 'DAY'
  AND part_hier.level_name = 'PART'
  AND part_hier.p_color in ('ivory','coral')
ORDER BY
  dwdate_hier.hier_order,
  part_hier.hier_order;
```

Autonomous Database Supplied Package Reference

This appendix provides information about the packages you use with Autonomous Database. The `DBMS_CLOUD` topic also covers the `DBMS_CLOUD REST APIs`.

- DBMS_CLOUD Package
  The DBMS_CLOUD package provides database routines for working with cloud resources.
- DBMS_CLOUD_ADMIN Package
  The DBMS_CLOUD_ADMIN package provides administrative routines for configuring a database.
- DBMS_CLOUD_AI Package
  The DBMS_CLOUD_AI package, with Select AI, facilitates and configures the translation of natural language prompts to generate, run, explain SQL statements. Also, enables retrieval augmented generation and natural language-based interactions, including chatting with LLMs.
- DBMS_CLOUD_FUNCTION Package
  The DBMS_CLOUD_FUNCTION package allows you invoke OCI and AWS Lambda remote functions in your Autonomous Database as SQL functions.
- DBMS_CLOUD_LINK Package
  The DBMS_CLOUD_LINK package allows a user to register a table or a view as a data set for read only access with Cloud Links.
- DBMS_CLOUD_LINK_ADMIN Package
- DBMS_CLOUD_MACADM Package
  This section covers the DBMS_CLOUD_MACADM subprograms provided with Autonomous Database.
- DBMS_CLOUD_MIGRATION Package
  The DBMS_CLOUD_MIGRATION package facilitates translation of the SQL code from a non-Oracle SQL flavor to Oracle SQL.
- DBMS_CLOUD_NOTIFICATION Package
  The DBMS_CLOUD_NOTIFICATION package allows you to send messages or the output of a SQL query to a provider.
- DBMS_CLOUD_PIPELINE Package
  The DBMS_CLOUD_PIPELINE package allows you to create data pipelines for loading and exporting data in the cloud. This package supports continuous incremental data load of files in object store into the database. DBMS_CLOUD_PIPELINE also supports continuous incremental export of table data or query results from the database to object store based on a timestamp column.
- DBMS_CLOUD_REPO Package
  The DBMS_CLOUD_REPO package provides for use of and management of cloud hosted code repositories from Oracle Database. Supported cloud code repositories include GitHub, AWS CodeCommit , and Azure Repos.
- DBMS_AUTO_PARTITION Package
  The DBMS_AUTO_PARTITION package provides administrative routines for managing automatic partitioning of schemas and tables.
- DBMS_DATA_ACCESS Package
  The DBMS_DATA_ACCESS package provides routines to generate and manage Table Hyperlinks for data sets.
- DBMS_DCAT Package
  The DBMS_DCAT package provides functions and procedures to help Autonomous Database users leverage the data discovery and centralized metadata management system of OCI Data Catalog.

- DBMS_INMEMORY_ADMIN Package
  The DBMS_INMEMORY_ADMIN package enables you to manage the In-Memory area and In-Memory FastStart.
- DBMS_MAX_STRING_SIZE Package
  The DBMS_MAX_STRING_SIZE package provides an interface for checking and changing the value of the DBMS_MAX_STRING_SIZE initialization parameter.
- DBMS_PIPE Package
- DBMS_PROXY_SQL
  Describes the procedures to enable and disable query offload for an elastic pool leader.
- DBMS_SHARE Package
  The DBMS_SHARE subprograms and views are used to share local data with external systems.
- CS_RESOURCE_MANAGER Package
  The CS_RESOURCE_MANAGER package provides an interface to list and update consumer group parameters, and to revert parameters to default values.
- CS_SESSION Package
  The CS_SESSION package provides an interface to switch the database service and consumer group of the existing session.

DBMS_CLOUD Package

The DBMS_CLOUD package provides database routines for working with cloud resources.

- DBMS_CLOUD Endpoint Management
- DBMS_CLOUD Subprograms and REST APIs
  This section covers the DBMS_CLOUD subprograms and REST APIs provided with Autonomous Database.
- DBMS_CLOUD URI Formats
  Describes the format of the source file URIs in operations with DBMS_CLOUD. The format depends on the object storage service you are using.
- DBMS_CLOUD Package Format Options
  The format argument in DBMS_CLOUD specifies the format of source files.
- DBMS_CLOUD Package Format Options for EXPORT_DATA
  Describes the valid format parameter options for DBMS_CLOUD.EXPORT_DATA with text file formats, CSV, JSON, Parquet, or XML, and for Oracle Data Pump.
- DBMS_CLOUD Avro, ORC, and Parquet Support
  This section covers the DBMS_CLOUD Avro, ORC, and Parquet support provided with Autonomous Database.
- DBMS_CLOUD Exceptions
  The following table describes exceptions for DBMS_CLOUD.

DBMS_CLOUD Endpoint Management

Describes the format of pre-configured endpoint URIs in operations with DBMS_CLOUD.
The DBMS_CLOUD package supports a number of object store and REST endpoints that are all pre-configured for the service. This package also allows you to access additional non-preconfigured endpoints after properly enabling such endpoints through network access control lists (**ACLs**).

The authentication for pre-configured endpoints is automatically understood and derived in `DBMS_CLOUD`. In order to know the proper authentication type for additional, customer-managed endpoints, `DBMS_CLOUD` supports `URI` schemes to indicate the authentication type for the URI endpoint. The `URI` endpoint must support HTTPS on port 443 for secure HTTP requests.

- For more information on pre-configured URIs, see DBMS_CLOUD URI Formats.
- For more information on customer-managed URIs, see Additional Customer-Managed URI Formats

DBMS_CLOUD Subprograms and REST APIs

This section covers the `DBMS_CLOUD` subprograms and REST APIs provided with Autonomous Database.

> **Note:**
>
> To run `DBMS_CLOUD` subprograms with a user other than ADMIN you need to grant `EXECUTE` privileges to that user. For example, run the following command as ADMIN to grant privileges to `adb_user`:
>
> ```
> GRANT EXECUTE ON DBMS_CLOUD TO adb_user;
> ```

The `DBMS_CLOUD` package is made up of the following:

- DBMS_CLOUD for Access Management
  The subprograms for credential management within the DBMS_CLOUD package, including creating, deleting, and updating credentials.
- DBMS_CLOUD for Objects and Files
  The subprograms for object and file management within the DBMS_CLOUD package.
- DBMS_CLOUD for Bulk File Management
  The subprograms for bulk file operations within the DBMS_CLOUD package.
- DBMS_CLOUD REST APIs
  This section covers the `DBMS_CLOUD` REST APIs provided with Autonomous Database.

DBMS_CLOUD for Access Management

The subprograms for credential management within the DBMS_CLOUD package, including creating, deleting, and updating credentials.

Subprogram	Description
CREATE_CREDENTIAL Procedure	This procedure stores cloud service credentials in Autonomous Database.
DROP_CREDENTIAL Procedure	This procedure removes an existing credential from Autonomous Database.
REFRESH_VAULT_CREDENTIAL Procedure	This procedure immediately refreshes the vault secret of a vault secret credential to get the latest version of the vault secret for the specified `credential_name` in Autonomous Database.

Subprogram	Description
UPDATE_CREDENTIAL Procedure	This procedure updates cloud service credential attributes in Autonomous Database.

- CREATE_CREDENTIAL Procedure
  This procedure stores cloud service credentials in Autonomous Database.
- DROP_CREDENTIAL Procedure
  This procedure removes an existing credential from Autonomous Database.
- REFRESH_VAULT_CREDENTIAL Procedure
  This procedure refreshes the vault secret of a vault secret credential.
- UPDATE_CREDENTIAL Procedure
  This procedure updates an attribute with a new value for a specified `credential_name`.

CREATE_CREDENTIAL Procedure

This procedure stores cloud service credentials in Autonomous Database.

Use stored cloud service credentials to access the cloud service for data loading, for querying external data residing in the cloud, or for other cases when you use `DBMS_CLOUD` procedures with a `credential_name` parameter. This procedure is overloaded:

- Use the Oracle Cloud Infrastructure-related parameters, including: `user_ocid`, `tenancy_ocid`, `private_key`, and `fingerprint` only when you are using Oracle Cloud Infrastructure Signing Keys authentication.
- Use the `params` parameter for one of the following:
  - Amazon Resource Names (ARNs) credentials
  - Google Analytics or Google BigQuery credentials
  - Vault secret credentials for use with a supported vault:
    - Oracle Cloud Infrastructure Vault
    - Azure Key Vault
    - AWS Secrets Manager
    - GCP Secret Manager

Syntax

```
DBMS_CLOUD.CREATE_CREDENTIAL (
      credential_name   IN VARCHAR2,
      username          IN VARCHAR2,
      password          IN VARCHAR2 DEFAULT NULL,
      comments          IN VARCHAR2 DEFAULT NULL);


DBMS_CLOUD.CREATE_CREDENTIAL (
      credential_name   IN VARCHAR2,
      user_ocid         IN VARCHAR2,
      tenancy_ocid      IN VARCHAR2,
      private_key       IN VARCHAR2,
      fingerprint       IN VARCHAR2,
      comments          IN VARCHAR2 DEFAULT NULL);
```

```
DBMS_CLOUD.CREATE_CREDENTIAL (
      credential_name  IN VARCHAR2,
      params           IN CLOB DEFAULT);
```

Parameters

Parameter	Description
`credential_name`	The name of the credential to be stored. The `credential_name` parameter must conform to Oracle object naming conventions. The `credential_name` parameter can be defined as a nonquoted identifier or as a quoted identifier. See Database Object Naming Rules for more information.
`username`	The `username` and `password` arguments together specify your cloud service credentials. See the usage notes for what to specify for the `username` and `password` for different cloud services.
`password`	The `username` and `password` arguments together specify your cloud service credentials.
`comments`	Specifies comment text for the credential. This field can be useful to describe and distinguish between credentials. The maximum comment length is 3800 characters.
`user_ocid`	Specifies the user's OCID. See Where to Get the Tenancy's OCID and User's OCID for details on obtaining the User's OCID.
`tenancy_ocid`	Specifies the tenancy's OCID. See Where to Get the Tenancy's OCID and User's OCID for details on obtaining the Tenancy's OCID.
`private_key`	Specifies the generated private key. Private keys generated with a passphrase are not supported. You need to generate the private key without a passphrase. See How to Generate an API Signing Key for details on generating a key pair in PEM format.
`fingerprint`	Specifies a fingerprint. After a generated public key is uploaded to the user's account the fingerprint is displayed in the console. Use the displayed fingerprint for this argument. See How to Get the Key's Fingerprint and How to Generate an API Signing Key for more details.
`params`	Specifies credential parameters for one of the following: • Amazon Resource Names (ARNs) credentials • Google Analytics or Google BigQuery credentials • Vault secret credentials for use with username/password type credentials where you store the password in a supported vault: – Oracle Cloud Infrastructure Vault – Azure Key Vault – AWS Secrets Manager – GCP Secret Manager To create a vault secret credential you must have `EXECUTE` privilege on the `DBMS_CLOUD` package.

Usage Notes

- This operation stores the credentials in the database in an encrypted format.
- You can see the credentials in your schema by querying the `user_credentials` table.

- You can see the credential comments in your schema by querying the `user_credentials` table. For example:

```
SELECT comments FROM user_credentials
        WHERE credential_name = 'MY_CRED';
```

- The `ADMIN` user can see all the credentials by querying the `dba_credentials` table.
- You only need to create credentials once unless your cloud service credentials change. Once you store the credentials you can then use the same credential name for `DBMS_CLOUD` procedures that require a `credential_name` parameter.
- This procedure is overloaded. If you provide one of the key based authentication attributes, `user_ocid`, `tenancy_ocid`, `private_key`, or `fingerprint`, the call is assumed to be an Oracle Cloud Infrastructure Signing Key based credential.
- You can list credentials from the view `ALL_CREDENTIALS`. For example, run the following command to list credentials:

```
SELECT credential_name, username, comments FROM all_credentials;
```

Oracle Cloud Infrastructure Credentials (Auth Tokens)

For Oracle Cloud Infrastructure the `username` is your Oracle Cloud Infrastructure user name. The `password` is your Oracle Cloud Infrastructure auth token. See Working with Auth Tokens.

For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'adb_user@example.com',
    password => 'password' );
END;
/
```

Use Auth Token based credentials when you are authenticating calls to OCI Object Storage. For calls to any other type of Oracle Cloud Infrastructure cloud service, use Oracle Cloud Infrastructure Signing Key Based Credentials.

For OCI Object Storage, `username` parameter value must include the Identity domain and the user name from your profile. You can find the Identity domain associated with a user in the Oracle Cloud Infrastructure Console.

For example:

```
oracleidentitycloudservice/adb_user@example.com
```

With the default Identity domain you are not required to include the domain name `Default`. For example:

```
adb_user@example.com
```

Oracle Cloud Infrastructure Signing Key Based Credentials

Use the Oracle Cloud Infrastructure signing key related parameters, including: `user_ocid`, `tenancy_ocid`, `private_key`, and `fingerprint` with Oracle Cloud Infrastructure Signing Keys authentication.

For example:

```
BEGIN
   DBMS_CLOUD.CREATE_CREDENTIAL (
       credential_name => 'OCI_KEY_CRED',
       user_ocid       =>
'ocid1.user.oc1..aaaaaaaauq54mi7zdyfhw33ozkwuontjceel7fok5nq3bf2vwetkpqsoa',
       tenancy_ocid    =>
'ocid1.tenancy.oc1..aabbbbbbaafcue47pqmrf4vigneebgbcmmoy5r7xvoypicjqqge32ewnrc
yx2a',
       private_key     =>
'MIIEogIBAAKCAQEAtUnxbmrekwgVac6FdWeRzoXvIpA9+0r1.....wtnNpESQQQ0QLGPD8NM//
JEBg=',
       fingerprint     => 'f2:db:f9:18:a4:aa:fc:94:f4:f6:6c:39:96:16:aa:27');
END;
/
```

Private keys generated with a passphrase are not supported. You need to generate the private key without a passphrase. See How to Generate an API Signing Key for more information.

Oracle Cloud Infrastructure Object Storage Classic Credentials

If your source files reside in Oracle Cloud Infrastructure Object Storage Classic, the `username` is your Oracle Cloud Infrastructure Classic user name and the `password` is your Oracle Cloud Infrastructure Classic password.

Amazon Web Services (AWS) Credentials

If your source files reside in Amazon S3 or you are calling an AWS API, the `username` is your AWS access key ID and the `password` is your AWS secret access key. See AWS Identity and Access Management.

Microsoft Azure Credentials

If your source files reside in Azure Blob Storage or Azure Data Lake Storage or you are calling an Azure API, the `username` is your Azure storage account name and the `password` is an Azure storage account access key. See About Azure storage accounts.

Amazon S3-Compatible Credentials

Service	Credentials Information
Oracle Cloud Infrastructure (Customer Secret Keys)	If your source files reside in Oracle Cloud Infrastructure, then you need to use Customer Secret Keys with S3-compatible URLs. See Working with Customer Secret Keys for more information.

Service	Credentials Information
Google Cloud Storage	If your source files reside in Google Cloud Storage or you are calling Google Cloud Storage APIs, then you need to set a default Google project and obtain an HMAC key to create credentials to supply with Google Cloud Storage S3-compatible URLs. Use the HMAC key id as the username, and the HMAC secret as the password. See Projects and HMAC Keys for more information.
Wasabi Hot Cloud Storage	If your source files reside in Wasabi Hot Cloud Storage or you are calling Wasabi Hot Cloud Storage APIs, then you need Access Keys to create credentials to supply with S3-compatible URLs. Use the Wasabi Hot Cloud Storage Access Key as the username, and the Wasabi Hot Cloud Storage Secret Key as the password. See Creating a Wasabi API Access Key Set for more information.

AWS Amazon Resource Names (ARN) Credentials

If your source files reside in Amazon S3 or you are calling an AWS API, use `params` to specify the parameters for the Amazon Resource Names (ARN).

Parameter	Value
`aws_role_arn`	Specifies the Amazon Resource Name (ARN) that identifies the AWS role. If this parameter is not supplied when creating the credential, `ORA-20041` is raised.
`external_id_type`	Optionally set the `external_id_type` to use the Autonomous Database compartment OCID, database OCID, or tenancy OCID by supplying one of: `compartment_ocid`, `database_ocid`, or `tenant_ocid`. If this parameter is not given when creating the credential, the default value is `database_ocid`.

For example:

```
BEGIN DBMS_CLOUD.CREATE_CREDENTIAL(
        credential_name      => 'MY_CRED',
        params               => JSON_OBJECT(
               'aws_role_arn'      value 'arn:aws:iam::123456:role/
AWS_ROLE_ARN',
               'external_id_type'  value 'database_ocid'));
END;
/
```

GitHub Personal Access Token

If your source files reside in a GitHub repository or you are calling a GitHub API, the `username` is your GitHub email and the `password` is your GitHub personal access token. See Creating a personal access token for more information.

For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'MY_GITHUB_CRED',
```

```
    username => 'user@example.com',
    password => 'your_personal_access_token' );
END;
/
```

Google Analytics or Google BigQuery Credentials

If you are accessing Google Analytics or Google BigQuery, use the `params` parameter to specify the Google OAuth 2.0 credential parameters.

Parameter	Value
`gcp_oauth2`	Specifies OAuth 2.0 access for Google Analytics or Google BigQuery with a JSON object that includes the following parameters and their values: • `client_id`: See the Google API Console to obtain the client ID. • `client_secret`: See the Google API Console to obtain the client secret. • `refresh_token`: A refresh token allows your application to obtain new access tokens. See Using OAuth 2.0 to Access Google APIs for more information on Google OAuth credentials.

For example:

```
BEGIN DBMS_CLOUD.CREATE_CREDENTIAL(
      credential_name => 'GOOGLE_BIGQUERY_CRED',
      params => JSON_OBJECT('gcp_oauth2' value
                                  JSON_OBJECT(
                                       'client_id'     value 'client_id',
                                       'client_secret' value 'client_secret',
                                       'refresh_token' value
'refresh_token' )));
END;
/
```

Vault Secret Credentials with Oracle Cloud Infrastructure Vault

To create vault secret credentials with Oracle Cloud Infrastructure Vault, use the `params` parameter to specify the required parameters:

- `username`: Specifies the username of any type of username/password credential such as the username of OCI Swift password. For example, if you have a Swift credential with username as “scott” and password as “*password*”, provide “scott” as the `username` parameter.
- `secret_id`: Is the vault secret ID. Specify the `secret_id` value as the vault secret OCID. See Overview of Vault for more information.
- `region`: Is an optional parameter that specifies the oracle cloud region identifier. The region, when specified, indicates the location where Oracle Cloud Infrastructure Vault secret is located.

  By default, `CREATE_CREDENTIAL` uses the `region` mapped from the region key in the `secret_id`. An example of a region is `us-ashburn-1`.

  See Regions and Availability Domains for a complete list of regions.

For example:

```
BEGIN DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'OCI_SECRET_CRED',
    params => JSON_OBJECT(
                'username'   value 'scott',
                'region'     value 'us-ashburn-1',
                'secret_id'  value 'ocid1.vaultsecret.co1.ap-
mumbai-1.example..aaaaaaaauq5ok5nq3bf2vwetkpqsoa'));
END;
/
```

Notes for using an Oracle Cloud Infrastructure Vault secret to store vault secrets:

- When you use an Oracle Cloud Infrastructure Vault, on the Autonomous Database instance you must enable principal authentication with `DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL`.
- On Oracle Cloud Infrastructure you must specify a policy for the resource principal to access the secret.

To create a vault secret credential you must have `EXECUTE` privilege on the `DBMS_CLOUD` package.

Vault Secret Credentials with Azure Key Vault

To create Azure Key Vault credentials, use the `params` parameter to specify the required parameters:

- `username`: Specifies the username associated with the key.
- `secret_id`: Specifies the secret name.
- `azure_vault_name`: Specifies the name of the vault where the secret is located.

See Create a key vault for more information.

To create a vault secret credential you must have `EXECUTE` privilege on the `DBMS_CLOUD` package.

Vault Secret Credentials with AWS Secrets Manager

To create vault secret credentials with AWS Secrets Manager, use the `params` parameter to specify the required parameters:

- `username`: Specifies the AWS Secrets Manager access key.
- `secret_id`: Is the AWS Secrets Manager AWS ARN.
- `region`: (Optional) Specifies the AWS service region where the vault and secret are located. An example of the AWS region is "us-east-2". The default `region` is the region specified with the ARN in the `secret_id` parameter.

  See Managing AWS Regions for more information.

  To create a vault secret credential you must have `EXECUTE` privilege on the `DBMS_CLOUD` package.

Vault Secret Credentials with GCP Secret Manager

To create GCP Secret Manager credentials, use the `params` parameter to specify the required parameters:

- `username`: Specifies the username associated with the secret.
- `secret_id`: Is the secret name.
- `gcp_project_id`: Specifies the ID of the project where the secret is located.

See Secret Manager for more information.

To create a vault secret credential you must have `EXECUTE` privilege on the `DBMS_CLOUD` package.

DROP_CREDENTIAL Procedure

This procedure removes an existing credential from Autonomous Database.

Syntax

```
DBMS_CLOUD.DROP_CREDENTIAL (
   credential_name      IN VARCHAR2);
```

Parameters

Parameter	Description
`credential_name`	The name of the credential to be removed.

REFRESH_VAULT_CREDENTIAL Procedure

This procedure refreshes the vault secret of a vault secret credential.

This procedure lets you immediately refresh the vault secret of a vault secret credential to get the latest version of the vault secret for the specified `credential_name`.

Syntax

```
DBMS_CLOUD.REFRESH_VAULT_CREDENTIAL (
    credential_name    IN VARCHAR2);
```

Parameters

Parameter	Description
`credential_name`	The name of the credential to refresh.

Usage Notes

- The `ADMIN` user can see all the credentials by querying the `dba_credentials` table.

- You can list credentials from the view `ALL_CREDENTIALS`. For example, run the following command to list credentials:

```
SELECT credential_name, username, comments FROM all_credentials;
```

Example

```
BEGIN
  DBMS_CLOUD.REFRESH_VAULT_CREDENTIAL(
     credential_name => 'AZURE_SECRET_CRED');
END;
/
```

UPDATE_CREDENTIAL Procedure

This procedure updates an attribute with a new value for a specified `credential_name`.

Use stored credentials for data loading, for querying external data residing in the Cloud, or wherever you use `DBMS_CLOUD` procedures with a `credential_name` parameter.

Syntax

```
DBMS_CLOUD.UPDATE_CREDENTIAL (
    credential_name   IN VARCHAR2,
    attribute         IN VARCHAR2,
    value             IN VARCHAR2);
```

Parameters

Parameter	Description
`credential_name`	The name of the credential to be updated.
`attribute`	Name of attribute to update. For a username/password type credential, the valid `attribute` values are: `USERNAME` and `PASSWORD`. For a credential for an Amazon ARN, the valid `attribute` values are: `aws_role_arn` and `external_id_type`. For a credential for Google BigQuery or Google Analytics, the valid `attribute` values are: `client_id`, `client_secret`, and `refresh_token`. Depending on the vault you are using, for Vault Secret Credentials the valid `attribute` values are: • Oracle Cloud Infrastructure Vault: `secret_id`, `region` • Azure Key Vault: `secret_id`, `azure_vault_name` • AWS Secrets Manager: `secret_id`, `region` • GCP Secret Manager: `secret_id`, `gcp_project_id` See CREATE_CREDENTIAL Procedure for more information.
`value`	New value for the specified attribute.

Usage Notes

- The username value is case sensitive. It cannot contain double quotes or spaces.

- The `ADMIN` user can see all the credentials by querying `dba_credentials`.
- You only need to create credentials once unless your cloud service credentials change. Once you store the credentials you can then use the same credential name for `DBMS_CLOUD` procedures that require a `credential_name` parameter.
- You can list credentials from the view `ALL_CREDENTIALS`. For example, run the following command to list credentials:

```
SELECT credential_name, username, comments FROM all_credentials;
```

Examples

```
BEGIN
  DBMS_CLOUD.UPDATE_CREDENTIAL(
     credential_name => 'OBJ_STORE_CRED',
     attribute => 'PASSWORD',
     value => 'password');
END;
/


BEGIN
  DBMS_CLOUD.UPDATE_CREDENTIAL(
     credential_name => 'ARN_CRED',
     attribute => 'aws_role_arn',
     value => 'NEW_AWS_ARN');
END;
/
```

DBMS_CLOUD for Objects and Files

The subprograms for object and file management within the DBMS_CLOUD package.

Subprogram	Description
COPY_COLLECTION Procedure	This procedure loads data into existing SODA collection either from Cloud Object Storage or from files in a directory.
COPY_DATA Procedure	This procedure loads data into existing Autonomous Database tables either from Cloud Object Storage or from files in a directory.
COPY_DATA Procedure for Avro, ORC, or Parquet Files	This procedure with the `format` parameter `type` set to the value `orc`, `parquet`, or `avro` loads data into existing Autonomous Database tables from ORC, Parquet, or Avro files in the Cloud or from ORC, Parquet, or Avro files in a directory. Similar to text files, the data is copied from the source ORC, Parquet, or Avro file into the preexisting internal table.
COPY_OBJECT Procedure	This procedure copies files from one Cloud Object Storage bucket to another.
CREATE_EXTERNAL_TABLE Procedure	This procedure creates an external table on files in the Cloud or on files in a directory. This allows you to run queries on external data from Autonomous Database.
CREATE_CLOUD_TABLE Procedure	This procedure creates a cloud table where all persistent data is stored in Oracle-Managed Object Storage.

Subprogram	Description
CREATE_EXTERNAL_TABLE Procedure for Apache Iceberg	This procedure creates external tables for Apache Iceberg tables in the supported configurations.
CREATE_EXTERNAL_TABLE Procedure for Avro, ORC, or Parquet Files	This procedure with the `format` parameter `type` set to the value `parquet`, `orc`, or `avro`, creates an external table with either Parquet, ORC, or Avro format files in the Cloud or in a directory. This allows you to run queries on external data from Autonomous Database.
CREATE_EXTERNAL_PART_TABLE Procedure	This procedure creates an external partitioned table on files in the Cloud. This allows you to run queries on external data from Autonomous Database.
CREATE_EXTERNAL_TEXT_INDEX Procedure	This procedure creates text index on the object store files.
CREATE_HYBRID_PART_TABLE Procedure	This procedure creates a hybrid partitioned table. This allows you to run queries on hybrid partitioned data from Autonomous Database.
DELETE_ALL_OPERATIONS Procedure	This procedure clears either all data load operations logged in the `user_load_operations` table in your schema or clears all the data load operations of the specified type, as indicated with the `type` parameter.
DELETE_FILE Procedure	This procedure removes the specified file from the specified directory on Autonomous Database
DELETE_OBJECT Procedure	This procedure deletes the specified object on object store.
DELETE_OPERATION Procedure	This procedure accepts an `operation_id` as input and deletes the logs associated with the specified `operation_id`.
DROP_EXTERNAL_TEXT_INDEX Procedure	This procedure drops text index on the object store files.
EXPORT_DATA Procedure	This procedure exports data from Autonomous Database to files in the Cloud based on the result of a query. The overloaded form enables you to use the `operation_id` parameter. Depending on the `format` parameter `type` option specified, the procedure exports rows to the Cloud Object store as text with options of CSV, JSON, Parquet, or XML; or using the ORACLE_DATAPUMP access driver to write data to a dump file.
GET_OBJECT Procedure and Function	This procedure is overloaded. The procedure form reads an object from Cloud Object Storage and copies it to Autonomous Database. The function form reads an object from Cloud Object Storage and returns a `BLOB` to Autonomous Database.
LIST_FILES Function	This function lists the files in the specified directory. The results include the file names and additional metadata about the files such as file size in bytes, creation timestamp, and the last modification timestamp.
LIST_OBJECTS Function	This function lists objects in the specified location on object store. The results include the object names and additional metadata about the objects such as size, checksum, creation timestamp, and the last modification timestamp.
MOVE_OBJECT Procedure	This procedure moves an object from one Cloud Object Storage bucket to another one.

Subprogram	Description
PUT_OBJECT Procedure	This procedure is overloaded. In one form the procedure copies a file from Autonomous Database to the Cloud Object Storage. In another form the procedure copies a `BLOB` from Autonomous Database to the Cloud Object Storage.
SYNC_EXTERNAL_PART_TABLE Procedure	This procedure simplifies updating an external partitioned table from files in the Cloud. Run this procedure whenever new partitions are added or when partitions are removed from the Object Store source for the external partitioned table.
VALIDATE_EXTERNAL_TABLE Procedure	This procedure validates the source files for an external table, generates log information, and stores the rows that do not match the format options specified for the external table in a *badfile* table on Autonomous Database.
VALIDATE_EXTERNAL_PART_TABLE Procedure	This procedure validates the source files for an external partitioned table, generates log information, and stores the rows that do not match the format options specified for the external table in a *badfile* table on Autonomous Database.
VALIDATE_HYBRID_PART_TABLE Procedure	This procedure validates the source files for a hybrid partitioned table, generates log information, and stores the rows that do not match the format options specified for the hybrid table in a *badfile* table on Autonomous Database.

- COPY_COLLECTION Procedure
  This procedure loads data into a SODA collection from Cloud Object Storage or from a directory. If the specified SODA collection does not exist, the procedure creates it. The overloaded form enables you to use the `operation_id` parameter.
- COPY_DATA Procedure
  This procedure loads data into existing Autonomous Database tables from files in the Cloud, or from files in a directory. The overloaded form enables you to use the `operation_id` parameter.
- COPY_DATA Procedure for Avro, ORC, or Parquet Files
- COPY_OBJECT Procedure
  This procedure copies an object from one Cloud Object Storage bucket or folder to another.
- CREATE_CLOUD_TABLE Procedure
  This procedure creates a Cloud Table. All Cloud Table data is stored in Oracle managed Object Storage (Cloud Tables only store their metadata in the database).
- CREATE_EXTERNAL_PART_TABLE Procedure
  This procedure creates an external partitioned table on files in the Cloud, or from files in a directory. This allows you to run queries on external data from Autonomous Database.
- CREATE_EXTERNAL_TABLE Procedure
  This procedure creates an external table on files in the Cloud or from files in a directory. This allows you to run queries on external data from Autonomous Database.
- CREATE_EXTERNAL_TABLE Procedure for Apache Iceberg
  This procedure creates external tables for Apache Iceberg tables in the supported configurations.
- CREATE_EXTERNAL_TABLE Procedure for Avro, ORC, or Parquet Files
- CREATE_EXTERNAL_TEXT_INDEX Procedure
  This procedure creates a text index on Object Storage files.

- CREATE_HYBRID_PART_TABLE Procedure
  This procedure creates a hybrid partitioned table. This allows you to run queries on hybrid partitioned data from Autonomous Database using database objects and files in the Cloud, or database objects and files in a directory.
- DELETE_ALL_OPERATIONS Procedure
  This procedure clears either all data load operations logged in the `user_load_operations` table in your schema or clears all the data load operations of the specified type, as indicated with the `type` parameter.
- DELETE_FILE Procedure
  This procedure removes the specified file from the specified directory on Autonomous Database.
- DELETE_OBJECT Procedure
  This procedure deletes the specified object on object store.
- DELETE_OPERATION Procedure
  This procedure clears the data load entries for the specified operation ID logged in the `user_load_operations` or `dba_load_operations` tables in your schema.
- DROP_EXTERNAL_TEXT_INDEX Procedure
  This procedure drops text index on the Object Storage files.
- EXPORT_DATA Procedure
- GET_OBJECT Procedure and Function
  This procedure is overloaded. The procedure form reads an object from Cloud Object Storage and copies it to Autonomous Database. The function form reads an object from Cloud Object Storage and returns a `BLOB` to Autonomous Database.
- LIST_FILES Function
  This function lists the files in the specified directory. The results include the file names and additional metadata about the files such as file size in bytes, creation timestamp, and the last modification timestamp.
- LIST_OBJECTS Function
  This function lists objects in the specified location on object store. The results include the object names and additional metadata about the objects such as size, checksum, creation timestamp, and the last modification timestamp.
- MOVE_OBJECT Procedure
  This procedure moves an object from one Cloud Object Storage bucket or folder to another.
- PUT_OBJECT Procedure
  This procedure is overloaded. In one form the procedure copies a file from Autonomous Database to the Cloud Object Storage. In another form the procedure copies a `BLOB` from Autonomous Database to the Cloud Object Storage.
- SYNC_EXTERNAL_PART_TABLE Procedure
  This procedure simplifies updating an external partitioned table from files in the Cloud. Run this procedure whenever new partitions are added or when partitions are removed from the Object Store source for the external partitioned table.
- VALIDATE_EXTERNAL_PART_TABLE Procedure
  This procedure validates the source files for an external partitioned table, generates log information, and stores the rows that do not match the format options specified for the external table in a *badfile* table on Autonomous Database. The overloaded form enables you to use the `operation_id` parameter.

- VALIDATE_EXTERNAL_TABLE Procedure
  This procedure validates the source files for an external table, generates log information, and stores the rows that do not match the format options specified for the external table in a *badfile* table on Autonomous Database. The overloaded form enables you to use the `operation_id` parameter.
- VALIDATE_HYBRID_PART_TABLE Procedure
  This procedure validates the source files for a hybrid partitioned table, generates log information, and stores the rows that do not match the format options specified for the hybrid table in a *badfile* table on Autonomous Database. The overloaded form enables you to use the `operation_id` parameter.

COPY_COLLECTION Procedure

This procedure loads data into a SODA collection from Cloud Object Storage or from a directory. If the specified SODA collection does not exist, the procedure creates it. The overloaded form enables you to use the `operation_id` parameter.

Syntax

```
DBMS_CLOUD.COPY_COLLECTION (
    collection_name   IN VARCHAR2,
    credential_name   IN VARCHAR2 DEFAULT NULL,
    file_uri_list     IN CLOB,
    format            IN CLOB     DEFAULT NULL
);

DBMS_CLOUD.COPY_COLLECTION (
    collection_name   IN VARCHAR2,
    credential_name   IN VARCHAR2 DEFAULT NULL,
    file_uri_list     IN CLOB,
    format            IN CLOB     DEFAULT NULL,
    operation_id      OUT NOCOPY NUMBER
);
```

Parameters

Parameter	Description
`collection_name`	The name of the SODA collection into which data will be loaded. If a collection with this name already exists, the specified data will be loaded, otherwise a new collection is created.
`credential_name`	The name of the credential to access the Cloud Object Storage. You can use `'OCI$RESOURCE_PRINCIPAL'` as the `credential_name` when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. This parameter is not used when you specify a directory with `file_uri_list`.

Parameter	Description
file_uri_list	This parameter specifies either a comma-delimited list of source file URIs or one or more directories and source files. **Cloud source file URIs** You can use wildcards as well as regular expressions in the file names in Cloud source file URIs. Regular expressions can only be used when the regexuri format parameter is set to TRUE. The characters "*" and "?" are considered wildcard characters when the regexuri parameter is set to FALSE. When the regexuri parameter is set to TRUE the characters "*" and "?" are part of the specified regular expression pattern. Regular expression patterns are only supported for the file name or subfolder path in your URIs and the pattern matching is identical to that performed by the REGEXP_LIKE function. For example: `file_uri_list => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o(/)*year=[0-9]+(/)*month=[0-9]+(/)*[a-z]+[1-3]??.csv'` The format of the URIs depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD URI Formats. See REGEXP_LIKE Condition for more information on REGEXP_LIKE condition. **Directory** You can specify one directory and one or more file names or use a comma separated list of directories and file names. The format to specify a directory is: '*MY_DIR*:*filename.ext*'. By default the directory name *MY_DIR* is a database object and is case-insensitive. The file name is case sensitive. Regular expressions are not supported when specifying the file names in a directory. You can only use wildcards to specify file names in a directory. The character "*" can be used as the wildcard for multiple characters, and the character "?" can be used as the wildcard for a single character. For example: '*MY_DIR*:*" or '*MY_DIR*:test?' To specify multiple directories, use a comma separated list of directories: For example: '*MY_DIR1*:*, *MY_DIR2*:test?' Use double quotes to specify a case-sensitive directory name. For example: '"*my_dir1*":*, "*my_dir2*":Test?' To include a quote character, use two quotes. For example: '*MY_DIR*:''filename.ext'. This specifies the filename starts with a quote ('').
format	The options describing the format of the source files. These options are specified as a JSON string. Supported formats are: characterset, compression, encryption, ignoreblanklines, jsonpath, maxdocsize, recorddelimiter, rejectlimit, type, unpackarrays, keyassignment, and keypath. Apart from the mentioned formats for JSON data, Autonomous Database supports other formats too. For the list of format arguments supported by Autonomous Database, see DBMS_CLOUD Package Format Options.

Parameter	Description
operation_id	Use this parameter to track the progress and final status of the load operation as the corresponding ID in the USER_LOAD_OPERATIONS view.

Example

```
BEGIN
    DBMS_CLOUD.CREATE_CREDENTIAL(
            credential_name => 'OBJ_STORE_CRED',
            username        => 'user_name@oracle.com',
            password        => 'password'
            );

    DBMS_CLOUD.COPY_COLLECTION(
            collection_name => 'myCollection',
            credential_name => 'OBJ_STORE_CRED',
            file_uri_list   => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/adbexample/b/json/o/myCollection.json'
            );
END;
/
```

COPY_DATA Procedure

This procedure loads data into existing Autonomous Database tables from files in the Cloud, or from files in a directory. The overloaded form enables you to use the operation_id parameter.

Syntax

```
DBMS_CLOUD.COPY_DATA (
    table_name         IN VARCHAR2,
    credential_name    IN VARCHAR2 DEFAULT NULL,
    file_uri_list      IN CLOB,
    schema_name        IN VARCHAR2,
    field_list         IN CLOB,
    format             IN CLOB);

DBMS_CLOUD.COPY_DATA (
    table_name         IN VARCHAR2,
    credential_name    IN VARCHAR2 DEFAULT NULL,
    file_uri_list      IN CLOB DEFAULT NULL,
    schema_name        IN VARCHAR2 DEFAULT NULL,
    field_list         IN CLOB DEFAULT NULL,
    format             IN CLOB DEFAULT NULL
    operation_id       OUT NOCOPY NUMBER);
```

Parameters

Parameter	Description
table_name	The name of the target table on the database. The target table needs to be created before you run COPY_DATA.

Parameter	Description
`credential_name`	The name of the credential to access the Cloud Object Storage. You can use `'OCI$RESOURCE_PRINCIPAL'` as the `credential_name` when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. This parameter is not used when you specify a directory or Table Hyperlink URL with `file_uri_list`.

Parameter	Description

file_uri_list	This parameter specifies one of the following:
	• Comma-delimited list of source file URIs
	• Comma-delimited list of Table Hyperlink URLs
	• One or more directories and source files
	You can use wildcards as well as regular expressions in the file names in Cloud source file URIs.
	Cloud source file URIs
	This parameter specifies either a comma-delimited list of source file URIs or one or more directories and source files.
	Regular expressions can only be used when the regexuri format parameter is set to TRUE.
	The characters "*" and "?" are considered wildcard characters when the regexuri parameter is set to FALSE. When the regexuri parameter is set to TRUE the characters "*" and "?" are part of the specified regular expression pattern.
	Regular expression patterns are only supported for the file name or subfolder path in your URIs and the pattern matching is identical to that performed by the REGEXP_LIKE function.
	For example:
	`file_uri_list => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o(/)*year=[0-9]+(/)*month=[0-9]+(/)*[a-z]+[1-3]??.csv'`
	The format of the URIs depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD URI Formats.
	See REGEXP_LIKE Condition for more information on REGEXP_LIKE condition.
	Table Hyperlink URL
	You can use a Table Hyperlink in the following ways:
	• Specify a single Table Hyperlink URL.
	For example:
	`file_uri_list => 'https://dataaccess.adb.us-chicago-1.oraclecloudapps.com/adb/p/K6XExample/data'`
	• Specify a comma-delimited list of Table Hyperlink URLs, you must ensure that all included Table Hyperlinks have the same column names, column order, and column data types.
	For example:
	`file_uri_list => 'https://dataaccess.adb.us-chicago-1.oraclecloudapps.com/adb/p/K6XExample/data','https://dataaccess.adb.us-chicago-1.oraclecloudapps.com/adb/p/K7XExample/data'`
	See About Table Hyperlinks on Autonomous Database and GET_PREAUTHENTICATED_URL Procedure for more information.
	Directory
	You can specify one directory and one or more file names or use a comma separated list of directories and file names. The format to specify a directory is: '*MY_DIR*:*filename.ext*'. By default the directory name

Parameter	Description
	MY_DIR is a database object and is case-insensitive. The file name is case sensitive. Regular expressions are not supported when specifying the file names in a directory. You can only use wildcards to specify file names in a directory. The character "*" can be used as the wildcard for multiple characters, and the character "?" can be used as the wildcard for a single character. For example: '*MY_DIR*:*" or '*MY_DIR*:test?' To specify multiple directories, use a comma separated list of directories: For example: '*MY_DIR1*:*, *MY_DIR2*:test?' Use double quotes to specify a case-sensitive directory name. For example: '"*my_dir1*":*, "*my_dir2*":Test?' To include a quote character, use two quotes. For example: '*MY_DIR*:''filename.ext'. This specifies the `filename` starts with a quote (`'`).
`schema_name`	The name of the schema where the target table resides. The default value is NULL meaning the target table is in the same schema as the user running the procedure.
`field_list`	Identifies the fields in the source files and their data types. The default value is NULL meaning the fields and their data types are determined by the target table definition. This argument's syntax is the same as the `field_list` clause in regular Oracle external tables. For more information about `field_list` see *Oracle® Database Utilities*. When the `format` parameter `type` option value is `json`, this parameter is ignored. For an example using `field_list`, see CREATE_EXTERNAL_TABLE Procedure.
`format`	The options describing the format of the source, log, and bad files. For the list of the options and how to specify the values see DBMS_CLOUD Package Format Options. For Avro, ORC, or Parquet file format options, see DBMS_CLOUD Package Format Options for Avro, ORC, or Parquet.
`operation_id`	Use this parameter to track the progress and final status of the load operation as the corresponding ID in the `USER_LOAD_OPERATIONS` view.

Usage Note

The default record delimiter is `detected newline`. With `detected newline`, `DBMS_CLOUD` tries to automatically find the correct newline character to use as the record delimiter. `DBMS_CLOUD` first searches for the Windows newline character `\r\n`. If it finds the Windows newline character, this is used as the record delimiter for all files in the procedure. If a Windows newline character is not found, `DBMS_CLOUD` searches for the UNIX/Linux newline character `\n`, and if it finds one it uses `\n` as the record delimiter for all files in the procedure. If the source files use a combination of different record delimiters, you may encounter an error such as, "`KUP-04020: found record longer than buffer size supported`". In this case, you need to either modify the source files to use the same record delimiter or only specify the source files that use the same record delimiter.

See DBMS_CLOUD Package Format Options for information on the `recorddelmiter` format option.

Examples

```
BEGIN
    DBMS_CLOUD.CREATE_CREDENTIAL(
            credential_name => 'DEF_CRED_NAME',
            username        => 'user_name@oracle.com',
            password        => 'password'
            );
END;
/
```

```
BEGIN
 DBMS_CLOUD.COPY_DATA(
    table_name =>'CHANNELS',
    credential_name =>'DEF_CRED_NAME',
    file_uri_list =>'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/channels.txt',
    format => json_object('delimiter' value ',')
 );
END;
/
```

```
BEGIN
    DBMS_CLOUD.COPY_DATA(
            table_name      => 'ORDERS',
            schema_name     => 'TEST_SCHEMA',
            credential_name => 'DEF_CRED_NAME',
         file_uri_list   => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/adbexample/b/json/o/orde[r]s.tbl.1'
            format          =>  json_object('ignoreblanklines' value TRUE,
                                            'rejectlimit' value '0',
                                            'dateformat' value 'yyyy-mm-dd',
                                            'regexuri' value TRUE)
            );
END;
/
```

COPY_DATA Procedure for Avro, ORC, or Parquet Files

This procedure with the `format` parameter `type` set to the value `avro`, `orc`, or `parquet` loads data into existing Autonomous Database tables from Avro, ORC, or Parquet files in the Cloud or from files in a directory.

Similar to text files, the data is copied from the source Avro, ORC, or Parquet file into the preexisting internal table.

Syntax

```
DBMS_CLOUD.COPY_DATA (
    table_name        IN VARCHAR2,
    credential_name   IN VARCHAR2 DEFAULT NULL,
    file_uri_list     IN CLOB,
    schema_name       IN VARCHAR2 DEFAULT,
```

```
    field_list          IN CLOB DEFAULT,
    format              IN CLOB DEFAULT);
```

Parameters

Parameter	Description
`table_name`	The name of the target table on the database. The target table needs to be created before you run `COPY_DATA`.
`credential_name`	The name of the credential to access the Cloud Object Storage. You can use `'OCI$RESOURCE_PRINCIPAL'` as the `credential_name` when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. This parameter is not used when you specify a directory with `file_uri_list`.

Parameter	Description
file_uri_list	This parameter specifies either a comma-delimited list of source file URIs or one or more directories and source files. **Cloud source file URIs** You can use wildcards as well as regular expressions in the file names in Cloud source file URIs. Regular expressions can only be used when the regexuri format parameter is set to TRUE. The characters "*" and "?" are considered wildcard characters when the regexuri parameter is set to FALSE. When the regexuri parameter is set to TRUE the characters "*" and "?" are part of the specified regular expression pattern. Regular expression patterns are only supported for the file name or subfolder path in your URIs and the pattern matching is identical to that performed by the REGEXP_LIKE function. For example: `file_uri_list => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o(/)*year=[0-9]+(/)*month=[0-9]+(/)*[a-z]+[1-3]??.csv'` The format of the URIs depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD URI Formats. See REGEXP_LIKE Condition for more information on REGEXP_LIKE condition. **Directory** You can specify one directory and one or more file names or use a comma separated list of directories and file names. The format to specify a directory is: '*MY_DIR*:*filename.ext*'. By default the directory name *MY_DIR* is a database object and is case-insensitive. The file name is case sensitive. Regular expressions are not supported when specifying the file names in a directory. You can only use wildcards to specify file names in a directory. The character "*" can be used as the wildcard for multiple characters, and the character "?" can be used as the wildcard for a single character. For example: '*MY_DIR*:*" or '*MY_DIR*:test?' To specify multiple directories, use a comma separated list of directories: For example: '*MY_DIR1*:*, *MY_DIR2*:test?' Use double quotes to specify a case-sensitive directory name. For example: '"*my_dir1*":*, "*my_dir2*":Test?' To include a quote character, use two quotes. For example: '*MY_DIR*:''filename.ext'. This specifies the filename starts with a quote (') .
schema_name	The name of the schema where the target table resides. The default value is NULL meaning the target table is in the same schema as the user running the procedure.

Parameter	Description
field_list	Ignored for Avro, ORC, or Parquet files. The fields in the source match the external table columns by name. Source data types are converted to the external table column data type. For ORC files, see DBMS_CLOUD Package ORC to Oracle Data Type Mapping. For Parquet files, see DBMS_CLOUD Package Parquet to Oracle Data Type Mapping for details on mapping. For Avro files, see DBMS_CLOUD Package Avro to Oracle Data Type Mapping for details on mapping.
format	The options describing the format of the source files. For Avro, ORC, or Parquet files, only two options are supported: see DBMS_CLOUD Package Format Options for Avro, ORC, or Parquet.

Usage Notes

- As with other data files, Avro, ORC, and Parquet data loads generate logs that are viewable in the tables dba_load_operations and user_load_operations. Each load operation adds a record to dba[*user*]_load_operations that indicates the table containing the logs.

  The log table provides summary information about the load.

- For Avro, ORC, or Parquet, when the format parameter type is set to the value avro, orc, or parquet, the BADFILE_TABLE table is always empty.
  - For Parquet files, PRIMARY KEY constraint errors throw an ORA error.
  - If data for a column encounters a conversion error, for example, the target column is not large enough to hold the converted value, the value for the column is set to NULL. This does not produce a rejected record.

COPY_OBJECT Procedure

This procedure copies an object from one Cloud Object Storage bucket or folder to another.

The source and target bucket or folder can be in the same or different Cloud Object store provider.

When the source and target are in distinct Object Stores or have different accounts with the same cloud provider, you can give separate credential names for the source and target locations.

The source credential name is by default also used by the target location when target credential name is not provided.

Syntax

```
DBMS_CLOUD.COPY_OBJECT (
    source_credential_name  IN  VARCHAR2 DEFAULT NULL,
    source_object_uri       IN  VARCHAR2,
    target_object_uri       IN  VARCHAR2,
    target_credential_name  IN  VARCHAR2 DEFAULT NULL
);
```

Parameters

Parameter	Description
source_credential_name	The name of the credential to access the source Cloud Object Storage. You can use 'OCI$RESOURCE_PRINCIPAL' as the credential_name when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. If you do not supply a source_credential_name value, the credential_name is set to NULL.
source_object_uri	Specifies URI, that point to the source Object Storage bucket or folder location. This parameter is mandatory. The format of the URIs depends on the Cloud Object Storage service. See DBMS_CLOUD URI Formats for more information.
target_object_uri	Specifies the URI for the target Object Store. This parameter is mandatory. The format of the URIs depends on the Cloud Object Storage service. See DBMS_CLOUD URI Formats for more information.
target_credential_name	The name of the credential to access the target Cloud Object Storage location. You can use 'OCI$RESOURCE_PRINCIPAL' as the credential_name when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. If you do not supply a target_credential_name value, the target_object_uri is set to the source_credential_name value.

Example

```
BEGIN
DBMS_CLOUD.COPY_OBJECT (
    source_credential_name => 'OCI_CRED',
    source_object_uri    => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname1/bgfile.csv',
    target_object_uri    => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname2/myfile.csv'
);
END;
/
```

CREATE_CLOUD_TABLE Procedure

This procedure creates a Cloud Table. All Cloud Table data is stored in Oracle managed Object Storage (Cloud Tables only store their metadata in the database).

Syntax

```
DBMS_CLOUD.CREATE_CLOUD_TABLE (
    table_name       IN VARCHAR2,
    column_list      IN CLOB,
    params           IN CLOB);
```

Parameters

Parameter	Description
table_name	The name of the Cloud Table.
column_list	Comma-delimited list of column names and data types for the Cloud Table.

Usage Notes

- DEFAULT attributes: The column_list can include DEFAULT clause, which functions like the DEFAULT clause in an ordinary CREATE TABLE. See CREATE TABLE for information on the behavior of the DEFAULT clause.
- Use DROP TABLE to drop a Cloud Table. Cloud Tables do not support the recycle bin.

  For example:

  ```
  DROP TABLE CLOUD_TAB1;
  ```

- You can grant SELECT, INSERT, and UPDATE privileges for a Cloud Table. No other privileges can be granted to a Cloud Table.

  See Configuring Privilege and Role Authorization for more information.

Examples

```
EXEC DBMS_CLOUD.CREATE_CLOUD_TABLE( 'CLOUD_TAB1', 'I INTEGER, J INTEGER' );


BEGIN
  DBMS_CLOUD.CREATE_CLOUD_TABLE(
   table_name  => 'CLOUD_TABLE_WITH_DEFAULT',
   column_list => 'I INTEGER,
                   A VARCHAR2(32) DEFAULT ''ABC''' );
END;
/
```

CREATE_EXTERNAL_PART_TABLE Procedure

This procedure creates an external partitioned table on files in the Cloud, or from files in a directory. This allows you to run queries on external data from Autonomous Database.

Syntax

```
DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE (
    table_name           IN VARCHAR2,
    credential_name      IN VARCHAR2,
    partitioning_clause  IN CLOB,
    column_list          IN CLOB,
    field_list           IN CLOB DEFAULT,
    format               IN CLOB DEFAULT);


DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE (
```

```
    table_name           IN VARCHAR2,
    credential_name      IN VARCHAR2,
    file_uri_list        IN VARCHAR2,
    column_list          IN CLOB,
    field_list           IN CLOB DEFAULT,
    format               IN CLOB DEFAULT);
```

Parameters

Parameter	Description
`table_name`	The name of the external table.
`credential_name`	The name of the credential to access the Cloud Object Storage. You can use `'OCI$RESOURCE_PRINCIPAL'` as the `credential_name` when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information.
`partitioning_clause`	Specifies the complete partitioning clause, including the location information for individual partitions. If you use the `partitioning_clause` parameter, the `file_uri_list` parameter is not allowed.
`file_uri_list`	This parameter specifies either a comma-delimited list of source file URIs or one or more directories and source files. **Cloud source file URIs** You can use wildcards as well as regular expressions in the file names in Cloud source file URIs. Regular expressions can only be used when the `regexuri format` parameter is set to `TRUE`. The characters "*" and "?" are considered wildcard characters when the `regexuri` parameter is set to `FALSE`. When the `regexuri` parameter is set to `TRUE` the characters "*" and "?" are part of the specified regular expression pattern. Regular expression patterns are only supported for the file name or subfolder path in your URIs and the pattern matching is identical to that performed by the `REGEXP_LIKE` function. This option is only supported with external tables that are created on a file in the Object Storage. For example: `file_uri_list => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o(/)*year=[0-9]+(/)*month=[0-9]+(/)*[a-z]+[1-3]??.csv'` If you use the parameter `file_uri_list`, the `partitioning_clause` parameter is not allowed. The format of the URIs depends on the Cloud Object Storage service. See DBMS_CLOUD URI Formats for more information. See REGEXP_LIKE Condition for more information on `REGEXP_LIKE` condition.

Parameter	Description
column_list	Comma-delimited list of column names and data types for the external table. This parameter has the following requirements, depending on the type of the data files specified with the file_uri_list parameter: • The column_list parameter is required with unstructured files. Using unstructured files, for example with CSV text files, the column_list parameter must specify all the column names and data types inside the data file as well as the partition columns derived from the object name. • The column_list parameter is optional with structured files. For example, with Avro, ORC, or Parquet data files, the column_list is not required. When the column_list is not included, the format parameter partition_columns option must include specifications for both column names (name) and data types (type).
field_list	Identifies the fields in the source files and their data types. The default value is NULL meaning the fields and their data types are determined by the *column_list* parameter. This argument's syntax is the same as the field_list clause in regular Oracle external tables. For more information about field_list see *Oracle® Database Utilities*.

Parameter	Description
format	The format option partition_columns specifies the DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE column names and data types of partition columns when the partition columns are derived from the file path, depending on the type of data file, structured or unstructured: • When the DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE includes the column_list parameter and the data files are unstructured, such as with CSV text files, partition_columns does not include the data type. For example, use a format such as the following for this type of partition_columns specification: '"partition_columns":["state","zipcode"]' The data type is not required because it is specified in the DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE column_list parameter. • When the DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE does not include the column_list parameter and the data files are structured, such as Avro, ORC, or Parquet files, the partition_columns option includes both the column name, name sub-clause, and the data type, type sub-clause. For example, the following shows a partition_columns specification: '"partition_columns":[{"name":"country", "type":"varchar2(10)"}, {"name":"year", "type":"number"}, {"name":"month", "type":"varchar2(10)"}]' If the data files are unstructured and the type sub-clause is specified with partition_columns, the type sub-clause is ignored. For object names that are not based on hive format, the order of the partition_columns specified columns must match the order as they appear in the object name in the file path specified in the file_uri_list parameter. To see all the format parameter options describing the format of the source files, see DBMS_CLOUD Package Format Options.

Usage Notes

- You cannot call this procedure with both partitioning_clause and file_uri_list parameters.
- Specifying the column_list parameter is optional with structured data files, including Avro, Parquet, or ORC data files. If column_list is not specified, the format parameter partition_columns option must include both name and type.
- The column_list parameter is required with unstructured data files, such as CSV text files.
- The procedure DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE supports external partitioned files in the supported cloud object storage services, including:
  - Oracle Cloud Infrastructure Object Storage

  - Azure Blob Storage or Azure Data Lake Storage
  - Amazon S3
  - Amazon S3-Compatible, including: Oracle Cloud Infrastructure Object Storage, Google Cloud Storage, and Wasabi Hot Cloud Storage.
  - GitHub Repository

  See DBMS_CLOUD URI Formats for more information.

- The procedure `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` supports external partitioned files in directories, either in a local file system or in a network file system.
- When you call `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` with the `file_uri_list` parameter, the types for columns specified in the Cloud Object Store file name must be one of the following types:

  ```
  VARCHAR2(n)
  NUMBER(n)
  NUMBER(p,s)
  NUMBER
  DATE
  TIMESTAMP(9)
  ```

- The default record delimiter is `detected newline`. With `detected newline`, `DBMS_CLOUD` tries to automatically find the correct newline character to use as the record delimiter. `DBMS_CLOUD` first searches for the Windows newline character `\r\n`. If it finds the Windows newline character, this is used as the record delimiter for all files in the procedure. If a Windows newline character is not found, `DBMS_CLOUD` searches for the UNIX/Linux newline character `\n`, and if it finds one it uses `\n` as the record delimiter for all files in the procedure. If the source files use a combination of different record delimiters, you may encounter an error such as, "`KUP-04020: found record longer than buffer size supported`". In this case, you need to either modify the source files to use the same record delimiter or only specify the source files that use the same record delimiter.

  See DBMS_CLOUD Package Format Options for information on the `recorddelmiter` format option.
- The external partitioned tables you create with `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` include two invisible columns `file$path` and `file$name`. These columns help identify which file a record is coming from.
  - `file$path`: Specifies the file path text up to the beginning of the object name.
  - `file$name`: Specifies the object name, including all the text that follows the bucket name.

Examples

Example using the `partitioning_clause` parameter:

```
BEGIN
   DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE(
      table_name =>'PET1',
      credential_name =>'OBJ_STORE_CRED',
      format => json_object('delimiter' value ',', 'recorddelimiter' value
'newline', 'characterset' value 'us7ascii'),
      column_list => 'col1 number, col2 number, col3 number',
      partitioning_clause => 'partition by range (col1)
```

```
                                        (partition p1 values less than (1000) location
                                            ( ''&base_URL//file_11.txt'')
                                         ,
                                         partition p2 values less than (2000) location
                                            ( ''&base_URL/file_21.txt'')
                                         ,
                                         partition p3 values less than (3000)
location
                                            ( ''&base_URL/file_31.txt'')
                                         )'

     );
   END;
/


BEGIN
    DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE(
       table_name          => 'PET',
       format              => json_object('delimiter'value ','),
       column_list         => 'name varchar2(20), gender varchar2(10), salary
number',
       partitioning_clause => 'partition by range (salary)
              (   -- Use test1.csv in the DEFAULT DIRECTORY DATA_PUMP_DIR
                  partition p1 values less than (100) LOCATION
(''test1.csv''),
                   -- Use test2.csv in a specified directory MY_DIR
                  partition p2 values less than (300) DEFAULT DIRECTORY
MY_DIR LOCATION (''test2.csv'')        )'   );
END;
/
```

Example using the `file_uri_list` and `column_list` parameters with unstructured data files:

```
BEGIN
  DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE(
   table_name => 'MYSALES',
   credential_name => 'DEF_CRED_NAME',
   file_uri_list     => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-
string/b/bucketname/o/*.csv',
   column_list       => 'product varchar2(100), units number, country varchar2(100), year
number, month varchar2(2)',
   field_list        => 'product, units', --[Because country, year and month are not in the
file, they are not listed in the field list]
   format            => '{"type":"csv", "partition_columns":["country","year","month"]}');
END;
/
```

Example using the `file_uri_list` without the `column_list` parameter with structured data files:

```
BEGIN
  DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE(
  table_name => 'MYSALES',
  credential_name => 'DEF_CRED_NAME',
  DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE(
```

```
    table_name      => 'MYSALES',
    credential_name => 'DEF_CRED_NAME',
    file_uri_list   => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-
string/b/bucketname/o/*.parquet',
    format          =>
        json_object('type' value 'parquet', 'schema' value 'first',
                    'partition_columns' value
                          json_array(
                                json_object('name' value 'country', 'type' value
'varchar2(100)'),
                                json_object('name' value 'year', 'type' value 'number'),
                                json_object('name' value 'month', 'type' value 'varchar2(2)')
                          )
         )
    );
END;
/
```

CREATE_EXTERNAL_TABLE Procedure

This procedure creates an external table on files in the Cloud or from files in a directory. This allows you to run queries on external data from Autonomous Database.

Syntax

```
DBMS_CLOUD.CREATE_EXTERNAL_TABLE (
    table_name       IN VARCHAR2,
    credential_name  IN VARCHAR2 DEFAULT NULL,
    file_uri_list    IN CLOB,
    column_list      IN CLOB,
    field_list       IN CLOB DEFAULT,
    format           IN CLOB DEFAULT);
```

Parameters

Parameter	Description
`table_name`	The name of the external table.
`credential_name`	The name of the credential to access the Cloud Object Storage. You can use `'OCI$RESOURCE_PRINCIPAL'` as the `credential_name` when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. This parameter is not used when you specify a directory or a Table Hyperlink URL with `file_uri_list`.

Parameter	Description

file_uri_list	This parameter specifies one of the following:
	• Comma-delimited list of source file URIs.
	• Comma-delimited list of Table Hyperlink URLs.
	• One or more directories and source files.
	Cloud source file URIs
	You can use wildcards as well as regular expressions in the file names in Cloud source file URIs.
	Regular expressions can only be used when the regexuri format parameter is set to TRUE.
	The characters "*" and "?" are considered wildcard characters when the regexuri parameter is set to FALSE. When the regexuri parameter is set to TRUE the characters "*" and "?" are part of the specified regular expression pattern.
	Regular expression patterns are only supported for the file name or subfolder path in your URIs and the pattern matching is identical to that performed by the REGEXP_LIKE function.
	This option is only supported with external tables that are created on a file in the Object Storage.
	For example:
	`file_uri_list => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o(/)*year=[0-9]+(/)*month=[0-9]+(/)*[a-z]+[1-3]??.csv'`
	The format of the URIs depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD URI Formats.
	See REGEXP_LIKE Condition for more information on REGEXP_LIKE condition.
	Table Hyperlink URL
	You can use a Table Hyperlink URL to create an external table in the following ways:
	• Specify a single Autonomous Database Table Hyperlink URL, you can also apply filters and clauses on the data when you read data from the table. For example, you can filter the data using the WHERE clause or sort it using the ORDER BY clause.
	For example:
	`file_uri_list => 'https://dataaccess.adb.us-chicago-1.oraclecloudapps.com/adb/p/K6XExample/data'`
	• Specify a comma-delimited list of Autonomous Database Table Hyperlinks URLs, you must ensure that all included Table Hyperlinks must have the same column names, column order, and column data types in the same schema.
	For example:
	`file_uri_list => 'https://dataaccess.adb.us-chicago-1.oraclecloudapps.com/adb/p/K6XExample/data','https://dataaccess.adb.us-`

Parameter	Description
	chicago-1.oraclecloudapps.com/adb/p/K7XExample/data'
	Note: Table Hyperlink URLs do not support wildcards.
	See About Table Hyperlinks on Autonomous Database and GET_PREAUTHENTICATED_URL Procedure for more information.
	Directory You can specify one directory and one or more file names or use a comma separated list of directories and file names. The format to specify a directory is:'*MY_DIR*:*filename.ext*'. By default the directory name *MY_DIR* is a database object and is case-insensitive. The file name is case sensitive. Regular expressions are not supported when specifying the file names in a directory. You can only use wildcards to specify file names in a directory. The character "*" can be used as the wildcard for multiple characters, and the character "?" can be used as the wildcard for a single character. For example:'*MY_DIR*:*" or '*MY_DIR*:test?' To specify multiple directories, use a comma separated list of directories: For example:'*MY_DIR1*:*, *MY_DIR2*:test?' Use double quotes to specify a case-sensitive directory name. For example:'"*my_dir1*":*, "*my_dir2*":Test?' To include a quote character, use two quotes. For example:'*MY_DIR*:''filename.ext'. This specifies the filename starts with a quote ('').
column_list	Comma-delimited list of column names and data types for the external table.
field_list	Identifies the fields in the source files and their data types. The default value is NULL meaning the fields and their data types are determined by the column_list parameter. This argument's syntax is the same as the field_list clause in regular Oracle Database external tables. For more information about field_list see ORACLE_LOADER Access Driver field_list under field_definitions Clause in *Oracle Database Utilities*.
format	The options describing the format of the source files. For the list of the options and how to specify the values see DBMS_CLOUD Package Format Options. For Avro, ORC, or Parquet format files, see CREATE_EXTERNAL_TABLE Procedure for Avro, ORC, or Parquet Files.

Usage Notes

- The procedure DBMS_CLOUD.CREATE_EXTERNAL_TABLE supports external partitioned files in the supported cloud object storage sources, including:
  - Oracle Cloud Infrastructure Object Storage
  - Azure Blob Storage or Azure Data Lake Storage

  - Amazon S3
  - Amazon S3-Compatible, including: Oracle Cloud Infrastructure Object Storage, Google Cloud Storage, and Wasabi Hot Cloud Storage.
  - GitHub Repository
  - Oracle Cloud Infrastructure Object Storage with Table Hyperlink URLs

  The credential is a table level property; therefore, the external files must be on the same object store.

  See DBMS_CLOUD URI Formats for more information.

- The default record delimiter is `detected newline`. With `detected newline`, `DBMS_CLOUD` tries to automatically find the correct newline character to use as the record delimiter. `DBMS_CLOUD` first searches for the Windows newline character `\r\n`. If it finds the Windows newline character, this is used as the record delimiter for all files in the procedure. If a Windows newline character is not found, `DBMS_CLOUD` searches for the UNIX/Linux newline character `\n`, and if it finds one it uses `\n` as the record delimiter for all files in the procedure. If the source files use a combination of different record delimiters, you may encounter an error such as, "`KUP-04020: found record longer than buffer size supported`". In this case, you need to either modify the source files to use the same record delimiter or only specify the source files that use the same record delimiter.

  See DBMS_CLOUD Package Format Options for information on the `recorddelimiter` format option.

Example

```
BEGIN
   DBMS_CLOUD.CREATE_EXTERNAL_TABLE(
      table_name =>'WEATHER_REPORT_DOUBLE_DATE',
      credential_name =>'OBJ_STORE_CRED',
      file_uri_list =>'&base_URL/
Charlotte_NC_Weather_History_Double_Dates.csv',
      format => json_object('type' value 'csv', 'skipheaders' value '1'),
      field_list => 'REPORT_DATE DATE''mm/dd/yy'',
                     REPORT_DATE_COPY DATE ''yyyy-mm-dd'',
                     ACTUAL_MEAN_TEMP,
                     ACTUAL_MIN_TEMP,
                     ACTUAL_MAX_TEMP,
                     AVERAGE_MIN_TEMP,
                     AVERAGE_MAX_TEMP,
                     AVERAGE_PRECIPITATION',
      column_list => 'REPORT_DATE DATE,
                     REPORT_DATE_COPY DATE,
                     ACTUAL_MEAN_TEMP NUMBER,
                     ACTUAL_MIN_TEMP NUMBER,
                     ACTUAL_MAX_TEMP NUMBER,
                     AVERAGE_MIN_TEMP NUMBER,
                     AVERAGE_MAX_TEMP NUMBER,
                     AVERAGE_PRECIPITATION NUMBER');
   END;
/

SELECT * FROM WEATHER_REPORT_DOUBLE_DATE where
   actual_mean_temp > 69 and actual_mean_temp < 74
```

CREATE_EXTERNAL_TABLE Procedure for Apache Iceberg

This procedure creates external tables for Apache Iceberg tables in the supported configurations.

For a description of supported configurations, see About Querying Apache Iceberg Tables.

Syntax

```
DBMS_CLOUD.CREATE_EXTERNAL_TABLE (
    table_name       IN VARCHAR2,
    credential_name  IN VARCHAR2 DEFAULT NULL,
    file_uri_list    IN CLOB,
    column_list      IN CLOB DEFAULT NULL,
    field_list       IN CLOB DEFAULT NULL,
    format           IN CLOB DEFAULT NULL
);
```

Parameters

Parameter	Description
`table_name`	The name of the external table.
`credential_name`	The name of the credential used to access the data files, the metadata files and the Iceberg Catalog (if used). For AWS and OCI configurations, the credential should be created as described in CREATE_CREDENTIAL Procedure. AWS Amazon Resource Names (ARN) credentials are currently not supported.
`file_uri_list`	Must be NULL if an Iceberg catalog is specified (see `format` parameter below). If an iceberg catalog is not used, then the `file_uri_list` must contain the URI to the iceberg metadata file.
`column_list`	Must be NULL, as the column names and types are automatically derived from Iceberg metadata. The column names match the names found in the underlying data files (Parquet, Avro, ORC). The Oracle data types are derived using the Parquet/Avro/ORC mappings between Iceberg and the Parquet, Avro and ORC data types. Therefore users cannot specify the `column_list`.
`field_list`	Must be NULL, as column names and data types are automatically derived from the Iceberg metadata.
`format`	The `format` parameter has a different structure depending on the type of Iceberg table, AWS or OCI, and what information is used to create the external table, for example information from a data catalog or a direct metadata URI. For examples and further information: see the examples below, Iceberg Support on OCI Data Flow Samples, DBMS_CLOUD URI Formats.

Example AWS Iceberg tables using an AWS Glue Catalog

The `format` parameter when creating tables over an AWS Iceberg table using an AWS Glue Catalog is as follows:

```
format => json_object('access_protocol' value
       json_object('protocol_type' value 'iceberg',
```

```
                    'protocol_config' value
                     json_object('iceberg_catalog_type' value 'aws_glue',
                                 'iceberg_glue_region'  value 'glue region',
                                 'iceberg_table_path'   value
'database_name.table_name')));
```

Where, the `access_protocol` parameter contains a JSON object with two elements as follows:

- `protocol_type`: Must be 'iceberg'
- `protocol_config`: A nested JSON object specifying the iceberg catalog details.
  - `iceberg_catalog_type`: Must be `'aws_glue'`
  - `iceberg_glue_region`: The catalog region, e.g. `'us-west-1'`
  - `iceberg_table_path`: A *`glue database.glue table name`* path.

Example AWS Iceberg table using a metadata file URI

The `format` parameter when creating tables over an AWS Iceberg table using a metadata file URI, is as follows:

```
format => json_object('access_protocol' value
       json_object('protocol_type' value 'iceberg')
```

Example OCI Iceberg table using HadoopCatalog catalog

The `format` parameter when creating tables over an OCI Iceberg table created by OCI Data Flow using HadoopCatalog catalog, is as follows:

```
format => json_object('access_protocol' value
       json_object('protocol_type'   value 'iceberg',
                   'protocol_config' value
                   json_object('iceberg_catalog_type'  value 'hadoop',
                               'iceberg_warehouse'     value '<OCI folder
URI>',
                               'iceberg_table_path'    value
'database_name.table_name')));
```

Where, the `access_protocol` parameter contains a JSON object with two elements as follows:

- `protocol_type`: Must be `'iceberg'`
- `protocol_config`: A nested JSON object specifying the iceberg catalog details.
  - `iceberg_catalog_type`: Must be `'hadoop'`
  - `iceberg_warehouse`: The warehouse directory path used when generating the table, in native URI format.
  - `iceberg_table_path`: The *`database_name.table name`* path used when creating the table.

Example OCI Iceberg table using the URI of the metadata file

The `format` parameter when creating tables over an OCI Iceberg table using the URI of the metadata file, is as follows:

```
format => json_object('access_protocol' value
       json_object('protocol_type' value 'iceberg')
```

Where, the `access_protocol` parameter contains a JSON object with one element as follows:

- `protocol_type`: Must be `'iceberg'`

CREATE_EXTERNAL_TABLE Procedure for Avro, ORC, or Parquet Files

This procedure with the `format` parameter `type` set to the value `avro`, `orc`, or `parquet` creates an external table with either Avro, ORC, or Parquet format files in the Cloud or in a directory.

This allows you to run queries on external data from Autonomous Database.

Syntax

```
DBMS_CLOUD.CREATE_EXTERNAL_TABLE (
    table_name       IN VARCHAR2,
    credential_name  IN VARCHAR2 DEFAULT NULL,
    file_uri_list    IN CLOB,
    column_list      IN CLOB,
    field_list       IN CLOB DEFAULT,
    format           IN CLOB DEFAULT);
```

Parameters

Parameter	Description
`table_name`	The name of the external table.
`credential_name`	The name of the credential to access the Cloud Object Storage. You can use `'OCI$RESOURCE_PRINCIPAL'` as the `credential_name` when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. This parameter is not used when you specify a directory with `file_uri_list`.

Parameter	Description
file_uri_list	This parameter specifies either a comma-delimited list of source file URIs or one or more directories and source files. **Cloud source file URIs** You can use wildcards as well as regular expressions in the file names in Cloud source file URIs. Regular expressions can only be used when the regexuri format parameter is set to TRUE. The characters "*" and "?" are considered wildcard characters when the regexuri parameter is set to FALSE. When the regexuri parameter is set to TRUE the characters "*" and "?" are part of the specified regular expression pattern. Regular expression patterns are only supported for the file name or subfolder path in your URIs and the pattern matching is identical to that performed by the REGEXP_LIKE function. This option is only supported with external tables that are created on a file in the Object Storage. For example: `file_uri_list => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o(/)*year=[0-9]+(/)*month=[0-9]+(/)*[a-z]+[1-3]??.parquet'` The format of the URIs depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD URI Formats. See REGEXP_LIKE Condition for more information on REGEXP_LIKE condition. **Directory** You can specify one directory and one or more file names or use a comma separated list of directories and file names. The format to specify a directory is:'*MY_DIR*:*filename.ext*'. By default the directory name *MY_DIR* is a database object and is case-insensitive. The file name is case sensitive. Regular expressions are not supported when specifying the file names in a directory. You can only use wildcards to specify file names in a directory. The character "*" can be used as the wildcard for multiple characters, and the character "?" can be used as the wildcard for a single character. For example:'*MY_DIR*:*" or '*MY_DIR*:test?' To specify multiple directories, use a comma separated list of directories: For example:'*MY_DIR1*:*, *MY_DIR2*:test?' Use double quotes to specify a case-sensitive directory name. For example:'"*my_dir1*":*, "*my_dir2*":Test?' To include a quote character, use two quotes. For example:'*MY_DIR*:''filename.ext'. This specifies the filename starts with a quote ().

Parameter	Description
column_list	(Optional) This field, when specified, overrides the format->schema parameter which specifies that the schema, columns, and data types, are derived automatically. See the format parameter for details. When the column_list is specified for Avro, ORC, or Parquet source, the column names must match those columns found in the file. Oracle data types must map appropriately to the Avro, ORC, or Parquet data types. For Parquet files, see DBMS_CLOUD Package Parquet to Oracle Data Type Mapping for details. For ORC files, see DBMS_CLOUD Package ORC to Oracle Data Type Mapping for details. For Avro files, see DBMS_CLOUD Package Avro to Oracle Data Type Mapping for details.
field_list	Ignored for Avro, ORC, or Parquet files. The fields in the source match the external table columns by name. Source data types are converted to the external table column data type. For ORC files, see DBMS_CLOUD Package ORC to Oracle Data Type Mapping For Parquet files, see DBMS_CLOUD Package Parquet to Oracle Data Type Mapping for details. For Avro files, see DBMS_CLOUD Package Avro to Oracle Data Type Mapping for details.
format	For Avro, ORC, or Parquet type source files, see DBMS_CLOUD Package Format Options for Avro, ORC, or Parquet for details.

Examples ORC

```
format => '{"type":"orc", "schema": "all"}'
```

```
format => json_object('type' value 'orc', 'schema' value 'first')
```

Examples Avro

```
format => '{"type":"avro", "schema": "all"}'
```

```
format => json_object('type' value 'avro', 'schema' value 'first')
```

Examples Parquet

```
format => '{"type":"parquet", "schema": "all"}'
```

```
format => json_object('type' value 'parquet', 'schema' value 'first')
```

Avro, ORC, or Parquet Column Name Mapping to Oracle Column Names

See DBMS_CLOUD Package Avro, ORC, and Parquet to Oracle Column Name Mapping for information on column name mapping and column name conversion usage in Oracle SQL.

CREATE_EXTERNAL_TEXT_INDEX Procedure

This procedure creates a text index on Object Storage files.

The `CREATE_EXTERNAL_TEXT_INDEX` procedure creates text index on the Object Storage files specified at the `location_uri` location. The index is refreshed at regular intervals, for any new additions or deletions done with files on location URI.

Syntax

```
DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX (
      credential_name  IN  VARCHAR2 DEFAULT NULL,
      location_uri     IN  VARCHAR2,
      index_name       IN  VARCHAR2,
      format           IN  CLOB     DEFAULT NULL
);
```

Parameters

Parameter	Description
`credential_name`	The name of the credential to access the Cloud Object Storage location. For public, pre-authenticated, or pre-signed bucket URIs, a NULL can be specified. See Configure Policies and Roles to Access Resources for more information. If you do not supply a `credential_name` value, the `credential_name` is set to a `NULL` value.
`location_uri`	Specifies the Object Store bucket or folder URI. This parameter is mandatory. The format of the URIs depends on the Cloud Object Storage service. See DBMS_CLOUD URI Formats for more information.
`index_name`	Specifies the name of the index you are building on the files located at the `location_uri` location. This parameter is mandatory.

Parameter	Description
format	Specifies additional configuration options. Options are specified as a JSON string. The supported format options are: refresh_rate: Specifies the frequency in minutes at which the local index is refreshed. New file uploads and deletions result in an index refresh. The default value is 5 minutes. binary_files: Specifies if the contents of the files to be indexed are binary. For example, PDF, MS-Word, The default value is FALSE. json_index: Specifies if the contents of the files to be indexed are JSON or row JSON. The default value is FALSE. json_index_doc_sep: Specifies the separator used for preprocessing row JSON documents. The default value is '\n' (new line). json_index_doc_len: Specifies the length of each JSON document in the object. The default value is 32767 and the maximum value is 200000. file_compression: Specifies the compression format of the JSON files. The default value is gzip. stop_words: Specifies a list of stop words can be supplied when you create indexes. The stop_words value indicates if it is a list of stop words or a table of stop words. When a JSON array is provided the stop words parameter it is treated as a list, otherwise the stop words parameter is treated as a table name whose column "STOP_WORDS" is used to read in the list of stop words. You can specify stop words using the following methods: • JSON Array: For example: format := '{"stop_words": ["king","queen"]}' • Stop word table name: For example: format := '{"stop_words":"STOP_WORDS_TABLE"}' If you do not supply a format parameter, the format is set to a NULL value.

Example

```
BEGIN
DBMS_CLOUD.CREATE_EXTERNAL_TEXT_INDEX (
        credential_name => 'DEFAULT_CREDENTIAL',
        location_uri    => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/ts_data/'
        index_name      => 'EMP',
        format          => JSON_OBJECT ('refresh_rate' value 10)
);
END;
/
```

CREATE_HYBRID_PART_TABLE Procedure

This procedure creates a hybrid partitioned table. This allows you to run queries on hybrid partitioned data from Autonomous Database using database objects and files in the Cloud, or database objects and files in a directory.

Syntax

```
DBMS_CLOUD.CREATE_HYBRID_PART_TABLE (
    table_name           IN VARCHAR2,
    credential_name      IN VARCHAR2,
    partitioning_clause  IN CLOB,
    column_list          IN CLOB,
    field_list           IN CLOB DEFAULT,
    format               IN CLOB DEFAULT);
```

Parameters

Parameter	Description
table_name	The name of the external table.
credential_name	The name of the credential to access the Cloud Object Storage. You can use 'OCI$RESOURCE_PRINCIPAL' as the credential_name when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information.
partitioning_clause	Specifies the complete partitioning clause, including the location information for individual partitions. To use directories, the partitioning clause supports the LOCATION and DEFAULT DIRECTORY values. You can use wildcards as well as regular expressions in the file names in Cloud source file URIs. Regular expressions can only be used when the regexuri format parameter is set to TRUE. The characters "*" and "?" are considered wildcard characters when the regexuri parameter is set to FALSE. When the regexuri parameter is set to TRUE the characters "*" and "?" are part of the specified regular expression pattern. Regular expression patterns are only supported for the file name or subfolder path in your URIs and the pattern matching is identical to that performed by the REGEXP_LIKE function. Regular expression patterns are not supported for directory names. For example: `partitioning_clause => 'partition by range (col1) (partition p1 values less than (1000) external location ( ''https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o(/)*year=[0-9]+(/)*month=[0-9]+(/)*[a-z]+[1-3]??.txt''),….` See REGEXP_LIKE Condition for more information on REGEXP_LIKE condition.
column_list	Comma-delimited list of column names and data types for the external table.

Parameter	Description
field_list	Identifies the fields in the source files and their data types. The default value is NULL meaning the fields and their data types are determined by the *column_list* parameter. This argument's syntax is the same as the field_list clause in regular Oracle external tables. For more information about field_list see *Oracle® Database Utilities*.
format	The options describing the format of the source files. For the list of the options and how to specify the values see DBMS_CLOUD Package Format Options.

Usage Notes

- The procedure DBMS_CLOUD.CREATE_HYBRID_PART_TABLE supports external partitioned files in the supported cloud object storage services, including:
  - Oracle Cloud Infrastructure Object Storage
  - Azure Blob Storage or Azure Data Lake Storage
  - Amazon S3
  - Amazon S3-Compatible, including: Oracle Cloud Infrastructure Object Storage, Google Cloud Storage, and Wasabi Hot Cloud Storage.
  - GitHub Repository

  The credential is a table level property; therefore, the external files must be on the same object store.

  See DBMS_CLOUD URI Formats for more information.
- The procedure DBMS_CLOUD.CREATE_HYBRID_PART_TABLE supports hybrid partitioned files in directories, either in a local file system or in a network file system.
- The external partitioned tables you create with DBMS_CLOUD.CREATE_HYBRID_PART_TABLE include two invisible columns file$path and file$name. These columns help identify which file a record is coming from.
  - file$path: Specifies the file path text up to the beginning of the object name.
  - file$name: Specifies the object name, including all the text that follows the bucket name.

Examples

```
BEGIN
   DBMS_CLOUD.CREATE_HYBRID_PART_TABLE(
      table_name =>'HPT1',
      credential_name =>'OBJ_STORE_CRED',
      format => json_object('delimiter' value ',', 'recorddelimiter' value
'newline', 'characterset' value 'us7ascii'),
      column_list => 'col1 number, col2 number, col3 number',
      partitioning_clause => 'partition by range (col1)
                                (partition p1 values less than (1000)
external location
                                    ( ''&base_URL/file_11.txt'')
                                ,
                                partition p2 values less than (2000)
external location
```

```
                                    ( ''&base_URL/file_21.txt'')
                                  ,
                                  partition p3 values less than (3000)
                                  )'
     );
   END;
/


BEGIN
   DBMS_CLOUD.CREATE_HYBRID_PART_TABLE(
    table_name  => 'HPT1',
    format      => json_object('delimiter'value ',', 'recorddelimiter'value
'newline'),
    column_list => 'NAME VARCHAR2(30), GENDER VARCHAR2(10), BALANCE number',
    partitioning_clause => 'partition by range (B  2  ALANCE)
               (partition p1 values less than (1000) external DEFAULT
DIRECTORY DATA_PUMP_DIR LOCATION (''Scott_male_1000.csv''),
                partition p2 values less than (2000) external DEFAULT
DIRECTORY DATA_PUMP_DIR LOCATION (''Mary_female_3000.csv''),
                partition p3 values less than (3000))' );
END;
/
```

DELETE_ALL_OPERATIONS Procedure

This procedure clears either all data load operations logged in the `user_load_operations` table in your schema or clears all the data load operations of the specified type, as indicated with the `type` parameter.

Syntax

```
DBMS_CLOUD.DELETE_ALL_OPERATIONS (
    type      IN VARCHAR DEFAULT NULL);
```

Parameters

Parameter	Description
`type`	Specifies the type of operation to delete. Type values can be found in the `TYPE` column in the `user_load_operations` table. If no `type` is specified all rows are deleted.

Usage Notes

- This procedure does not delete currently running operations (operations in a "Running" status).
- This procedure deletes any logging tables and log files associated with the operations.

DELETE_FILE Procedure

This procedure removes the specified file from the specified directory on Autonomous Database.

Syntax

```
DBMS_CLOUD.DELETE_FILE (
      directory_name     IN VARCHAR2,
      file_name          IN VARCHAR2,
      force              IN BOOLEAN DEFAULT FALSE);
```

Parameters

Parameter	Description
`directory_name`	The name of the directory on the Autonomous Database instance.
`file_name`	The name of the file to be removed.
`force`	Ignore and do not report errors if the file does not exist. Valid values are: `TRUE` and `FALSE`. The default value is `FALSE`.

> **Note:**
>
> To run `DBMS_CLOUD.DELETE_FILE` with a user other than ADMIN you need to grant write privileges on the directory that contains the file to that user. For example, run the following command as ADMIN to grant write privileges to `adb_user`:
>
> ```
> GRANT WRITE ON DIRECTORY data_pump_dir TO adb_user;
> ```

Example

```
BEGIN
   DBMS_CLOUD.DELETE_FILE(
      directory_name =>  'DATA_PUMP_DIR',
      file_name => 'exp1.dmp' );
   END;
/
```

DELETE_OBJECT Procedure

This procedure deletes the specified object on object store.

Syntax

```
DBMS_CLOUD.DELETE_OBJECT (
      credential_name      IN VARCHAR2,
      object_uri           IN VARCHAR2,
      force                IN BOOLEAN DEFAULT FALSE);
```

Parameters

Parameter	Description
credential_name	The name of the credential to access the Cloud Object Storage. You can use 'OCI$RESOURCE_PRINCIPAL' as the credential_name when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information.
object_uri	Object or file URI for the object to delete. The format of the URI depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD URI Formats.
force	Ignore and do not report errors if object does not exist. Valid values are: TRUE and FALSE. The default value is FALSE.

Example

```
BEGIN
   DBMS_CLOUD.DELETE_OBJECT(
       credential_name => 'DEF_CRED_NAME',
       object_uri => 'https://objectstorage.us-ashburn-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/exp1.dmp' );
   END;
/
```

DELETE_OPERATION Procedure

This procedure clears the data load entries for the specified operation ID logged in the user_load_operations or dba_load_operations tables in your schema.

Syntax

```
DBMS_CLOUD.DELETE_OPERATION (
    id      IN NUMBER);
```

Parameters

Parameter	Description
id	Specifies the operation ID associated with the log file entries you want to delete.

Usage Note

- This procedure deletes any logging tables and log files associated with the operation ID specified in the input.

Example

```
SELECT id FROM user_load_operations WHERE type LIKE '%BAD%';
EXEC DBMS_CLOUD.DELETE_OPERATION(id);
```

DROP_EXTERNAL_TEXT_INDEX Procedure

This procedure drops text index on the Object Storage files.

The `DROP_EXTERNAL_TEXT_INDEX` procedure drops the specified index created with the `CREATE_EXTERNAL_TEXT_INDEX` procedure.

Syntax

```
DBMS_CLOUD.DROP_EXTERNAL_TEXT_INDEX (
      index_name       IN  VARCHAR2,
);
```

Parameters

Parameter	Description
`index_name`	Specifies the name of the index you are dropping. The index name must match the name provided at the time of the index creation. This parameter is mandatory.

Example

```
BEGIN
DBMS_CLOUD.DROP_EXTERNAL_TEXT_INDEX (
        index_name => 'EMP',
);
END;
/
```

EXPORT_DATA Procedure

This procedure exports data from Autonomous Database based on the result of a query. This procedure is overloaded and supports writing files to the cloud or to a directory.

Based on the `format type` parameter, the procedure exports files to the Cloud or to a directory location as text files in CSV, JSON, Parquet, or XML format, or using the ORACLE_DATAPUMP access driver to write data to an Oracle Datapump dump file.

Syntax

```
DBMS_CLOUD.EXPORT_DATA (
      credential_name   IN VARCHAR2 DEFAULT NULL,
      file_uri_list     IN CLOB,
      format            IN CLOB,
      query             IN CLOB);

DBMS_CLOUD.EXPORT_DATA (
      credential_name   IN VARCHAR2 DEFAULT NULL,
      file_uri_list     IN CLOB DEFAULT NULL,
      format            IN CLOB DEFAULT NULL,
```

```
    query             IN CLOB DEFAULT NULL,
    operation_id      OUT NOCOPY NUMBER);
```

Parameters

Parameter	Description
credential_name	The name of the credential to access the Cloud Object Storage. You can use 'OCI$RESOURCE_PRINCIPAL' as the credential_name when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. When the credential parameter is not included, this specifies output to a directory.
file_uri_list	There are different forms, depending on the value of the format parameter and depending on whether you include a credential parameter: • When the format parameter type value is json: The JSON on Object Store or to the specified directory location is saved with a generated file name based on the value of the file_uri_list parameter. See File Naming for Text Output (CSV, JSON, Parquet, or XML) for more information. • When the format parameter type value is datapump, the file_uri_list is a comma-delimited list of the dump files. This specifies the files to be created on the Object Store. Use of wildcard and substitution characters is not supported in the file_uri_list. • When the credential_name parameter is not specified you provide a directory name in file_uri_list. The format of the URIs depend on the Cloud Object Storage service you are using, for details see DBMS_CLOUD URI Formats.
format	A JSON string that provides export format options. Supported option is: • type: The type format option is required and must have one of the values: csv | datapump | json | parquet | xml. See DBMS_CLOUD Package Format Options for EXPORT_DATA.
query	Use this parameter to specify a SELECT statement so that only the required data is exported. The query determines the contents of the files you export as text files CSV, JSON, Parquet, or XML, or as dump files. For example: SELECT warehouse_id, quantity FROM inventories For information with the format type value datapump, see Oracle Data Pump Export Data Filters and Unloading and Loading Data with the ORACLE_DATAPUMP Access Driver for more information. When the format type value is json, each query result is checked and if it is not JSON, as determined with the function: JSON_OBJECT_T.parse(), DBMS_CLOUD.EXPORT_DATA transforms the query to include JSON_OBJECT function to convert the row into JSON. See JSON_OBJECT_T Object Type for more information. For example: SELECT JSON_OBJECT(* RETURNING CLOB) from(SELECT warehouse_id, quantity FROM inventories)
operation_id	Use this parameter to track the progress and final status of the export operation as the corresponding ID in the USER_LOAD_OPERATIONS view.

Usage Notes:

- The `query` parameter value that you supply can be an advanced query, if required, such as a query that includes joins or subqueries.
- Depending on the format parameter specified, `DBMS_CLOUD.EXPORT_DATA` outputs the results of the specified query on the Cloud Object Store or to a directory location in one of these formats:
  - CSV, JSON, Parquet, or XML files.

    See Export Data to Object Store as Text and Export Data to a Directory for more information on using `DBMS_CLOUD.EXPORT_DATA` with CSV, JSON, Parquet, or XML output files.
  - Using the ORACLE_DATAPUMP access driver to write data to a dump file.
- For CSV, JSON, or XML output, by default when a generated file contains 10MB of data a new output file is created. However, if you have less than 10MB of result data you may have multiple output files, depending on the database service and the number of ECPUs (OCPUs if your database uses OCPUs) for the Autonomous Database instance.

  See File Naming for Text Output (CSV, JSON, Parquet, or XML) for more information.

  The default output file chunk size is 10MB for CSV, JSON, or XML. You can change this value with the `format` parameter `maxfilesize` option. See DBMS_CLOUD Package Format Options for EXPORT_DATA for more information.
- For Parquet output, each generated file is less than 128MB and multiple output files may be generated. However, if you have less than 128MB of result data, you may have multiple output files depending on the database service and the number of ECPUs (OCPUs if your database uses OCPUs) for the Autonomous Database instance.

  See File Naming for Text Output (CSV, JSON, Parquet, or XML) for more information.

Usage Notes for ORACLE_DATAPUMP Output (`DBMS_CLOUD.EXPORT_DATA` with `format` parameter `type` option `datapump`):

- `EXPORT_DATA` uses `DATA_PUMP_DIR` as the default logging directory. So the write privilege on `DATA_PUMP_DIR` is required when using `ORACLE_DATAPUMP` output.
- Autonomous Database export using `DBMS_CLOUD.EXPORT_DATA` with `format` parameter `type` option `datapump` only supports Oracle Cloud Infrastructure Object Storage, Oracle Cloud Infrastructure Object Storage Classic object stores or directory output.
- When you specify `DBMS_CLOUD.EXPORT_DATA` with the `format` parameter `type` option `datapump`, the `credential_name` parameter value cannot be an OCI resource principal.
- Oracle Data Pump divides each dump file part into smaller chunks for faster uploads. The Oracle Cloud Infrastructure Object Storage console shows multiple files for each dump file part that you export. The size of the actual dump files will be displayed as zero (0) and its related file chunks as 10mb or less. For example:

  ```
  exp01.dmp
  exp01.dmp_aaaaaa
  exp02.dmp
  exp02.dmp_aaaaaa
  ```

  Downloading the zero byte dump file from the Oracle Cloud Infrastructure console or using the Oracle Cloud Infrastructure CLI will not give you the full dump files. To download the

full dump files from the Object Store, use a tool that supports Swift such as curl, and provide your user login and Swift auth token.

```
curl -O -v -X GET -u 'user1@example.com:auth_token' \
   https://swiftobjectstorage.us-ashburn-1.oraclecloud.com/v1/namespace-
string/bucketname/exp01.dmp
```

If you import a file with the DBMS_CLOUD procedures that support the format parameter type with the value 'datapump', you only need to provide the primary file name. The procedures that support the 'datapump' format type automatically discover and download the chunks.

When you use DBMS_CLOUD.DELETE_OBJECT, the procedure automatically discovers and deletes the chunks when the procedure deletes the primary file.

- The DBMS_CLOUD.EXPORT_DATA procedure creates the dump file(s) from the file_uri_list values that you specify, as follows:
  - As more files are needed, the procedure creates additional files from the file_uri_list.
  - The procedure does not overwrite files. If a dump file in the file_uri_list exists, DBMS_CLOUD.EXPORT_DATA reports an error.
  - DBMS_CLOUD.EXPORT_DATA does not create buckets.
- The number of dump files that DBMS_CLOUD.EXPORT_DATA generates is determined when the procedure runs. The number of dump files that are generated depends on the number of file names you provide in the file_uri_list parameter, as well as on the number of Autonomous Database OCPUs available to the instance, the service level, and the size of the data.

  For example, if you use a 1 OCPU Autonomous Database instance or the low service, then a single dump file is exported with no parallelism, even if you provide multiple file names. If you use a 4 OCPU Autonomous Database instance with the medium or high service, then the jobs can run in parallel and multiple dump files are exported if you provide multiple file names.
- The dump files you create with DBMS_CLOUD.EXPORT_DATA cannot be imported using Oracle Data Pump impdp. Depending on the database, you can use these files as follows:
  - On an Autonomous Database, you can use the dump files with the DBMS_CLOUD procedures that support the format parameter type with the value 'datapump'. You can import the dump files using DBMS_CLOUD.COPY_DATA or you can call DBMS_CLOUD.CREATE_EXTERNAL_TABLE to create an external table.
  - On any other Oracle Database, such as Oracle Database 19c on-premise, you can import the dump files created with the procedure DBMS_CLOUD.EXPORT_DATA using the ORACLE_DATAPUMP access driver. See Unloading and Loading Data with the ORACLE_DATAPUMP Access Driver for more information.
- The query parameter value that you supply can be an advanced query, if required, such as a query that includes joins or subqueries.

Usage Notes for DBMS_CLOUD.EXPORT_DATA with Output to a Directory

- The provided directory must exist and you must be logged in as the ADMIN user or have WRITE access to the directory.
- DBMS_CLOUD.EXPORT_DATA does not create directories.

- The procedure does not overwrite files. For example, if a dump file in the `file_uri_list` exists, `DBMS_CLOUD.EXPORT_DATA` reports an error such as:

```
ORA-31641: unable to create dump file  "/u02/exports/123.dmp"
ORA-27038: created file already exists
```

Examples

The following example shows `DBMS_CLOUD.EXPORT_DATA` with the format `type` parameter with the value `datapump`:

```
BEGIN
   DBMS_CLOUD.EXPORT_DATA(
      credential_name =>'OBJ_STORE_CRED',
      file_uri_list =>'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/exp1.dmp',
      format => json_object('type' value 'datapump', 'compression' value
'basic', 'version' value 'latest'),
      query => 'SELECT warehouse_id, quantity FROM inventories'
     );
   END;
/
```

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

The following example shows `DBMS_CLOUD.EXPORT_DATA` with the format `type` parameter with the value `json`:

```
BEGIN
   DBMS_CLOUD.EXPORT_DATA(
      credential_name => 'OBJ_STORE_CRED',
      file_uri_list   => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/exp1.json',
      query           => 'SELECT * FROM DEPT',
      format          => JSON_OBJECT('type' value 'json', 'compression' value
'gzip'));
     );
   END;
/
```

The following example shows `DBMS_CLOUD.EXPORT_DATA` with the format `type` parameter with the value `xml`:

```
BEGIN
   DBMS_CLOUD.EXPORT_DATA(
      credential_name => 'OBJ_STORE_CRED',
      file_uri_list   => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/exp1.xml',
      query           => 'SELECT * FROM DEPT',
      format          => JSON_OBJECT('type' value 'xml', 'compression' value
'gzip'));
```

```
     );
   END;
/
```

The following example shows `DBMS_CLOUD.EXPORT_DATA` with the format `type` parameter with the value `csv`:

```
BEGIN
   DBMS_CLOUD.EXPORT_DATA(
      credential_name => 'OBJ_STORE_CRED',
      file_uri_list   => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o/exp.csv',
      query           => 'SELECT * FROM DEPT',
      format          => JSON_OBJECT('type' value 'csv', 'delimiter' value
'|', 'compression' value 'gzip', 'header' value true, 'encryption' value
('user_defined_function' value 'ADMIN.decryption_callback')));
     );
   END;
/
```

The following example shows `DBMS_CLOUD.EXPORT_DATA` exporting data to a directory location with the `type` parameter with the value `datapump`:

```
BEGIN
 DBMS_CLOUD.EXPORT_DATA(
    file_uri_list => 'export_dir:sales.dmp',
    format        => json_object('type' value 'datapump'),
    query         => 'SELECT * FROM sales'
 );
END;
/
```

GET_OBJECT Procedure and Function

This procedure is overloaded. The procedure form reads an object from Cloud Object Storage and copies it to Autonomous Database. The function form reads an object from Cloud Object Storage and returns a `BLOB` to Autonomous Database.

Syntax

```
DBMS_CLOUD.GET_OBJECT (
       credential_name      IN VARCHAR2,
       object_uri           IN VARCHAR2,
       directory_name       IN VARCHAR2,
       file_name            IN VARCHAR2 DEFAULT  NULL,
       startoffset          IN NUMBER DEFAULT  0,
       endoffset            IN NUMBER DEFAULT  0,
       compression          IN VARCHAR2 DEFAULT  NULL);


DBMS_CLOUD.GET_OBJECT(
       credential_name      IN VARCHAR2 DEFAULT NULL,
       object_uri           IN VARCHAR2,
       startoffset          IN NUMBER DEFAULT  0,
```

```
      endoffset             IN NUMBER DEFAULT  0,
      compression           IN VARCHAR2 DEFAULT  NULL)
RETURN BLOB;
```

Parameters

Parameter	Description
credential_name	The name of the credential to access the Cloud Object Storage. You can use 'OCI$RESOURCE_PRINCIPAL' as the credential_name when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information.
object_uri	Object or file URI. The format of the URI depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD URI Formats.
directory_name	The name of the directory on the database. 1 You can use a Table Hyperlink URL in the following ways: • Specify a single Autonomous Database Table Hyperlink URL. • Specify comma-delimited list of Autonomous Database Table Hyperlinks URLs, you must ensure that all included Table Hyperlinks must have the same column names, column order, and column data types in the same schema. See About Table Hyperlinks on Autonomous Database and GET_PREAUTHENTICATED_URL Procedure for more information.
file_name	Specifies the name of the file to create. If file name is not specified, the file name is taken from after the last slash in the object_uri parameter. For special cases, for example when the file name contains slashes, use the file_name parameter.
startoffset	The offset, in bytes, from where the procedure starts reading.
endoffset	The offset, in bytes, until where the procedure stops reading.
compression	Specifies the compression used to store the object. When compression is set to 'AUTO' the file is uncompressed (the value 'AUTO' implies the object specified with object_uri is compressed with Gzip).

> **Note:**
>
> To run DBMS_CLOUD.GET_OBJECT with a user other than ADMIN you need to grant WRITE privileges on the directory to that user. For example, run the following command as ADMIN to grant write privileges to adb_user:
>
> ```
> GRANT WRITE ON DIRECTORY data_pump_dir TO adb_user;
> ```

Return Values

The function form reads from Object Store and DBMS_CLOUD.GET_OBJECT returns a BLOB.

Examples

```
BEGIN
   DBMS_CLOUD.GET_OBJECT(
```

```
     credential_name => 'OBJ_STORE_CRED',
     object_uri => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/file.txt',
     directory_name => 'DATA_PUMP_DIR');
END;
/
```

To read character data from a file in Object Store:

```
SELECT to_clob(
     DBMS_CLOUD.GET_OBJECT(
       credential_name => 'OBJ_STORE_CRED',
       object_uri => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/file.txt'))
FROM DUAL;
```

To add an image stored on Object Store in a `BLOB` in the database:

```
DECLARE
   l_blob BLOB := NULL;
BEGIN
   l_blob := DBMS_CLOUD.GET_OBJECT(
     credential_name => 'OBJ_STORE_CRED',
     object_uri => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/MyImage.gif' );
END;
/
```

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

LIST_FILES Function

This function lists the files in the specified directory. The results include the file names and additional metadata about the files such as file size in bytes, creation timestamp, and the last modification timestamp.

Syntax

```
DBMS_CLOUD.LIST_FILES (
    directory_name      IN VARCHAR2)
       RETURN TABLE;
```

Parameters

Parameter	Description
`directory_name`	The name of the directory on the database.

Usage Notes

- To run `DBMS_CLOUD.LIST_FILES` with a user other than ADMIN you need to grant read privileges on the directory to that user. For example, run the following command as ADMIN to grant read privileges to `adb_user`:

  ```
  GRANT READ ON DIRECTORY data_pump_dir TO adb_user;
  ```

- This is a pipelined table function with return type as `DBMS_CLOUD_TYPES.list_object_ret_t`.
- `DBMS_CLOUD.LIST_FILES` does not obtain the checksum value and returns `NULL` for this field.

Example

This is a pipelined function that returns a row for each file. For example, use the following query to use this function:

```
SELECT * FROM DBMS_CLOUD.LIST_FILES('DATA_PUMP_DIR');

OBJECT_NAME       BYTES   CHECKSUM      CREATED              LAST_MODIFIED
------------ ---------- ----------    ---------------------
---------------------
cwallet.sso        2965               2018-12-12T18:10:47Z
2019-11-23T06:36:54Z
```

LIST_OBJECTS Function

This function lists objects in the specified location on object store. The results include the object names and additional metadata about the objects such as size, checksum, creation timestamp, and the last modification timestamp.

Syntax

```
DBMS_CLOUD.LIST_OBJECTS (
       credential_name      IN VARCHAR2,
       location_uri         IN VARCHAR2)
   RETURN TABLE;
```

Parameters

Parameter	Description
`credential_name`	The name of the credential to access the Cloud Object Storage. You can use `'OCI$RESOURCE_PRINCIPAL'` as the `credential_name` when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information.
`location_uri`	Object or file URI. The format of the URI depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD URI Formats.

Usage Notes

- Depending on the capabilities of the object store, `DBMS_CLOUD.LIST_OBJECTS` does not return values for certain attributes and the return value for the field is `NULL` in this case.

  All supported Object Stores return values for the `OBJECT_NAME`, `BYTES`, and `CHECKSUM` fields.

  The following table shows support for the fields `CREATED` and `LAST_MODIFIED` by Object Store:

Object Store	CREATED	LAST_MODIFIED
Oracle Cloud Infrastructure Native	Returns timestamp	Returns timestamp
Oracle Cloud Infrastructure Swift	Returns `NULL`	Returns timestamp
Oracle Cloud Infrastructure Classic	Returns `NULL`	Returns timestamp
Amazon S3	Returns `NULL`	Returns timestamp
Amazon S3-Compatible	Returns `NULL`	Returns timestamp
Azure	Returns timestamp	Returns timestamp
GitHub Repository		

- The checksum value is the MD5 checksum. This is a 32-character hexadecimal number that is computed on the object contents. It is expected to have a different checksum value if `OCI$RESOURCE_PRINCIPAL` credential is used.
- This is a pipelined table function with return type as `DBMS_CLOUD_TYPES.list_object_ret_t`.

Example

This is a pipelined function that returns a row for each object. For example, use the following query to use this function:

```
SELECT * FROM DBMS_CLOUD.LIST_OBJECTS('OBJ_STORE_CRED',
    'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/
bucketname/o/');


OBJECT_NAME   BYTES              CHECKSUM
CREATED         LAST_MODIFIED
------------ ---------- --------------------------------
--------------------- --------------------
cwallet.sso   2965      2339a2731ba24a837b26d344d643dc07
2019-11-23T06:36:54Z
```

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

MOVE_OBJECT Procedure

This procedure moves an object from one Cloud Object Storage bucket or folder to another.

The source and target bucket or folder can be in the same or different Cloud Object store provider.

When the source and target are in distinct Object Stores or have different accounts with the same cloud provider, you can give separate credential names for the source and target locations.

The source credential name is by default also used by the target location when target credential name is not provided.

Syntax

```
DBMS_CLOUD.MOVE_OBJECT (
    source_credential_name  IN  VARCHAR2 DEFAULT NULL,
    source_object_uri       IN  VARCHAR2,
    target_object_uri       IN  VARCHAR2,
    target_credential_name  IN  VARCHAR2 DEFAULT NULL
);
```

Parameters

Parameter	Description
`source_credential_name`	The name of the credential to access the source Cloud Object Storage. You can use `'OCI$RESOURCE_PRINCIPAL'` as the `credential_name` when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. If you do not supply a `source_credential_name` value, the `credential_name` is set to `NULL`.
`source_object_uri`	Specifies URI, that point to the source Object Storage bucket or folder location. This parameter is mandatory. The format of the URIs depends on the Cloud Object Storage service. See DBMS_CLOUD URI Formats for more information.
`target_object_uri`	Specifies the URI for the target Object Storage bucket or folder, where the files need to be moved. This parameter is mandatory. The format of the URIs depends on the Cloud Object Storage service. See DBMS_CLOUD URI Formats for more information.
`target_credential_name`	The name of the credential to access the target Cloud Object Storage location. You can use `'OCI$RESOURCE_PRINCIPAL'` as the `credential_name` when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. If you do not supply a `target_credential_name` value, the `target_object_uri` is set to the `source_credential_name` value.

Example

```
BEGIN
DBMS_CLOUD.MOVE_OBJECT (
    source_credential_name => 'OCI_CRED',
    source_object_uri    => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname1/bgfile.csv',
    target_object_uri    => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname2/myfile.csv'
);
```

```
END;
/
```

PUT_OBJECT Procedure

This procedure is overloaded. In one form the procedure copies a file from Autonomous Database to the Cloud Object Storage. In another form the procedure copies a `BLOB` from Autonomous Database to the Cloud Object Storage.

Syntax

```
DBMS_CLOUD.PUT_OBJECT (
       credential_name      IN VARCHAR2,
       object_uri           IN VARCHAR2,
       directory_name       IN VARCHAR2,
       file_name            IN VARCHAR2
       compression          IN VARCHAR2 DEFAULT  NULL);


DBMS_CLOUD.PUT_OBJECT (
       credential_name      IN VARCHAR2,
       object_uri           IN VARCHAR2,
       contents             IN BLOB
       compression          IN VARCHAR2 DEFAULT  NULL);
```

Parameters

Parameter	Description
`credential_name`	The name of the credential to access the Cloud Object Storage. You can use `'OCI$RESOURCE_PRINCIPAL'` as the `credential_name` when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information.
`object_uri`	Object or file URI. The format of the URI depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD URI Formats.
`directory_name`	The name of the directory on the Autonomous Database. 1
`contents`	Specifies the `BLOB` to copy from Autonomous Database to the Cloud Object Storage.
`file_name`	The name of the file in the specified directory.
`compression`	Specifies the compression used to store the object. Default value: `NULL`

> **Note:**
>
> To run `DBMS_CLOUD.PUT_OBJECT` with a user other than ADMIN you need to grant read privileges on the directory to that user. For example, run the following command as ADMIN to grant read privileges to `adb_user`:
>
> ```
> GRANT READ ON DIRECTORY data_pump_dir TO adb_user;
> ```

Example

To handle `BLOB` data after in-database processing and then store the data directly into a file in the object store:

```
DECLARE
      my_blob_data BLOB;
BEGIN
 /* Some processing producing BLOB data and populating my_blob_data */
DBMS_CLOUD.PUT_OBJECT(
     credential_name => 'OBJ_STORE_CRED',
     object_uri => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/my_new_file',
     contents => my_blob_data));
END;
/
```

Usage Notes

Depending on your Cloud Object Storage, the size of the object you transfer is limited as follows:

Cloud Object Storage Service	Object Transfer Size Limit
Oracle Cloud Infrastructure Object Storage	50 GB
Amazon S3	5 GB
Azure Blob Storage or Azure Data Lake Storage	256 MB
Amazon S3-Compatible	Set by the object store provider. For more information, refer to the provider's documentation.

Oracle Cloud Infrastructure object store does not allow writing files into a public bucket without supplying credentials (Oracle Cloud Infrastructure allows users to download objects from public buckets). Thus, you must supply a credential name with valid credentials to store an object in an Oracle Cloud Infrastructure public bucket using `PUT_OBJECT`.

See DBMS_CLOUD URI Formats for more information.

SYNC_EXTERNAL_PART_TABLE Procedure

This procedure simplifies updating an external partitioned table from files in the Cloud. Run this procedure whenever new partitions are added or when partitions are removed from the Object Store source for the external partitioned table.

Syntax

```
DBMS_CLOUD.SYNC_EXTERNAL_PART_TABLE (
    table_name         IN VARCHAR2,
    schema_name        IN VARCHAR2 DEFAULT,
    update_columns     IN BOOLEAN DEFAULT);
```

Parameters

Parameter	Description
table_name	The name of the target table. The target table needs to be created before you run DBMS_CLOUD.SYNC_EXTERNAL_PART_TABLE.
schema_name	The name of the schema where the target table resides. The default value is NULL meaning the target table is in the same schema as the user running the procedure.
update_columns	The new files may introduce a change to the schema. Updates supported include: new columns, deleted columns. Updates to existing columns, for example a change in the data type throw errors. Default Value: False

VALIDATE_EXTERNAL_PART_TABLE Procedure

This procedure validates the source files for an external partitioned table, generates log information, and stores the rows that do not match the format options specified for the external table in a *badfile* table on Autonomous Database. The overloaded form enables you to use the operation_id parameter.

Syntax

```
DBMS_CLOUD.VALIDATE_EXTERNAL_PART_TABLE (
       table_name                IN VARCHAR2,
       partition_name            IN CLOB DEFAULT,
       subpartition_name         IN CLOB DEFAULT,
       schema_name               IN VARCHAR2 DEFAULT,
       rowcount                  IN NUMBER DEFAULT,
       partition_key_validation  IN BOOLEAN DEFAULT,
       stop_on_error             IN BOOLEAN DEFAULT);


DBMS_CLOUD.VALIDATE_EXTERNAL_PART_TABLE (
       table_name                IN VARCHAR2,
       operation_id              OUT NUMBER,
       partition_name            IN CLOB DEFAULT,
       subpartition_name         IN CLOB DEFAULT,
       schema_name               IN VARCHAR2 DEFAULT,
       rowcount                  IN NUMBER DEFAULT,
       partition_key_validation  IN BOOLEAN DEFAULT,
       stop_on_error             IN BOOLEAN DEFAULT);
```

Parameters

Parameter	Description
table_name	The name of the external table.
operation_id	Use this parameter to track the progress and final status of the load operation as the corresponding ID in the USER_LOAD_OPERATIONS view.
partition_name	If defined, then only a specific partition is validated. If not specified then read all partitions sequentially until rowcount is reached.
subpartition_name	If defined, then only a specific subpartition is validated. If not specified then read from all external partitions or subpartitions sequentially until rowcount is reached.
schema_name	The name of the schema where the external table resides. The default value is NULL meaning the external table is in the same schema as the user running the procedure.
rowcount	Number of rows to be scanned. The default value is NULL meaning all the rows in the source files are scanned.
partition_key_validation	For internal use only. Do not use this parameter.
stop_on_error	Determines if the validate should stop when a row is rejected. The default value is TRUE meaning the validate stops at the first rejected row. Setting the value to FALSE specifies that the validate does not stop at the first rejected row and validates all rows up to the value specified for the rowcount parameter.

VALIDATE_EXTERNAL_TABLE Procedure

This procedure validates the source files for an external table, generates log information, and stores the rows that do not match the format options specified for the external table in a *badfile* table on Autonomous Database. The overloaded form enables you to use the operation_id parameter.

Syntax

```
DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE (
    table_name      IN VARCHAR2,
    schema_name     IN VARCHAR2 DEFAULT,
    rowcount        IN NUMBER DEFAULT,
    stop_on_error   IN BOOLEAN DEFAULT);


DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE(
    table_name      IN VARCHAR2,
    operation_id    OUT NOCOPY NUMBER,
    schema_name     IN VARCHAR2 DEFAULT NULL,
    rowcount        IN NUMBER DEFAULT 0,
    stop_on_error   IN BOOLEAN DEFAULT TRUE);
```

Parameters

Parameter	Description
table_name	The name of the external table.
operation_id	Use this parameter to track the progress and final status of the load operation as the corresponding ID in the USER_LOAD_OPERATIONS view.
schema_name	The name of the schema where the external table resides. The default value is NULL meaning the external table is in the same schema as the user running the procedure.
rowcount	Number of rows to be scanned. The default value is NULL meaning all the rows in the source files are scanned.
stop_on_error	Determines if the validate should stop when a row is rejected. The default value is TRUE meaning the validate stops at the first rejected row. Setting the value to FALSE specifies that the validate does not stop at the first rejected row and validates all rows up to the value specified for the rowcount parameter. If the external table refers to Avro, ORC, or Parquet files then the validate stops at the first rejected row. When the external table specifies the format parameter type set to the value avro, orc, or parquet, the parameter stop_on_error effectively always has the value TRUE. Thus, the table badfile will always be empty for an external table referring to Avro, ORC, or Parquet files.

Usage Notes

- DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE works with both partitioned external tables and hybrid partitioned tables. This potentially reads data from all external partitions until rowcount is reached or stop_on_error applies. You do not have control over which partition, or parts of a partition, is read in which order.

VALIDATE_HYBRID_PART_TABLE Procedure

This procedure validates the source files for a hybrid partitioned table, generates log information, and stores the rows that do not match the format options specified for the hybrid table in a *badfile* table on Autonomous Database. The overloaded form enables you to use the operation_id parameter.

Syntax

```
DBMS_CLOUD.VALIDATE_HYBRID_PART_TABLE (
       table_name                IN VARCHAR2,
       partition_name            IN CLOB DEFAULT,
       subpartition_name         IN CLOB DEFAULT,
       schema_name               IN VARCHAR2 DEFAULT,
       rowcount                  IN NUMBER DEFAULT,
       partition_key_validation  IN BOOLEAN DEFAULT,
       stop_on_error             IN BOOLEAN DEFAULT);


DBMS_CLOUD.VALIDATE_HYBRID_PART_TABLE (
       table_name                IN VARCHAR2,
       operation_id              OUT NUMBER,
       partition_name            IN CLOB DEFAULT,
```

```
subpartition_name           IN CLOB DEFAULT,
schema_name                 IN VARCHAR2 DEFAULT,
rowcount                    IN NUMBER DEFAULT,
partition_key_validation    IN BOOLEAN DEFAULT,
stop_on_error               IN BOOLEAN DEFAULT);
```

Parameters

Parameter	Description
table_name	The name of the external table.
operation_id	Use this parameter to track the progress and final status of the load operation as the corresponding ID in the USER_LOAD_OPERATIONS view.
partition_name	If defined, then only a specific partition is validated. If not specified then read from all external partitions sequentially until rowcount is reached.
subpartition_name	If defined, then only a specific subpartition is validated. If not specified then read from all external partitions or subpartitions sequentially until rowcount is reached.
schema_name	The name of the schema where the external table resides. The default value is NULL meaning the external table is in the same schema as the user running the procedure.
rowcount	Number of rows to be scanned. The default value is NULL meaning all the rows in the source files are scanned.
partition_key_validation	For internal use only. Do not use this parameter.
stop_on_error	Determines if the validate should stop when a row is rejected. The default value is TRUE meaning the validate stops at the first rejected row. Setting the value to FALSE specifies that the validate does not stop at the first rejected row and validates all rows up to the value specified for the rowcount parameter.

DBMS_CLOUD for Bulk File Management

The subprograms for bulk file operations within the DBMS_CLOUD package.

Subprogram	Description
BULK_COPY Procedure	This procedure copies files from one Cloud Object Storage bucket to another.
BULK_DELETE Procedure	The procedure deletes files from Cloud Object Storage bucket or folder.
BULK_DOWNLOAD Procedure	This procedure downloads files from Cloud Object store bucket to a directory in Autonomous Database.
BULK_MOVE Procedure	This procedure moves files from one Cloud Object Storage bucket to another.
BULK_UPLOAD Procedure	This procedure uploads files from a directory in Autonomous Database to the Cloud Object Storage.

- BULK_COPY Procedure
  This procedure bulk copies files from one Cloud Object Storage bucket to another. The overloaded form enables you to use the operation_id parameter.

- BULK_DELETE Procedure
  This procedure bulk deletes files from the Cloud Object Storage. The overloaded form enables you to use the `operation_id` parameter. You can filter the list of files to be deleted using a regular expression pattern compatible with `REGEXP_LIKE` operator.
- BULK_DOWNLOAD Procedure
  This procedure downloads files into an Autonomous Database directory from Cloud Object Storage. The overloaded form enables you to use the `operation_id` parameter. You can filter the list of files to be downloaded using a regular expression pattern compatible with `REGEXP_LIKE` operator.
- BULK_MOVE Procedure
  This procedure bulk moves files from one Cloud Object Storage bucket or folder to another. The overloaded form enables you to use the `operation_id` parameter.
- BULK_UPLOAD Procedure
  This procedure copies files into Cloud Object Storage from an Autonomous Database directory. The overloaded form enables you to use the `operation_id` parameter.

BULK_COPY Procedure

This procedure bulk copies files from one Cloud Object Storage bucket to another. The overloaded form enables you to use the `operation_id` parameter.

You can filter the list of files to be deleted using a regular expression pattern compatible with `REGEXP_LIKE` operator.

The source and target bucket or folder can be in the same or different Cloud Object store provider.

When the source and target are in distinct Object Stores or have different accounts with the same cloud provider, you can give separate credential names for the source and target locations.

The source credential name is by default also used by the target location.

Syntax

```
DBMS_CLOUD.BULK_COPY (
      source_credential_name  IN  VARCHAR2 DEFAULT NULL,
      source_location_uri     IN  VARCHAR2,
      target_location_uri     IN  VARCHAR2,
      target_credential_name  IN  VARCHAR2 DEFAULT NULL,
      regex_filter            IN  VARCHAR2 DEFAULT NULL,
      format                  IN  CLOB     DEFAULT NULL
);

DBMS_CLOUD.BULK_COPY (
      source_credential_name  IN  VARCHAR2 DEFAULT NULL,
      source_location_uri     IN  VARCHAR2,
      target_location_uri     IN  VARCHAR2,
      target_credential_name  IN  VARCHAR2 DEFAULT NULL,
      regex_filter            IN  VARCHAR2 DEFAULT NULL,
      format                  IN  CLOB     DEFAULT NULL,
      operation_id            OUT NUMBER
);
```

Parameters

Parameter	Description
source_credential_name	The name of the credential to access the Cloud Object Storage. You can use 'OCI$RESOURCE_PRINCIPAL' as the credential_name when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. If you do not supply a source_credential_name value, the credential_name is set to NULL.
source_location_uri	Specifies URI, that point to the source Object Storage bucket or folder location. This parameter is mandatory. The format of the URIs depends on the Cloud Object Storage service. See DBMS_CLOUD URI Formats for more information.
target_location_uri	Specifies the URI for the target Object Storage bucket or folder, where the files need to be copied. This parameter is mandatory. The format of the URIs depends on the Cloud Object Storage service. See DBMS_CLOUD URI Formats for more information.
target_credential_name	The name of the credential to access the target Cloud Object Storage location. You can use 'OCI$RESOURCE_PRINCIPAL' as the credential_name when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. If you do not supply a target_credential_name value, the target_location_uri is set to the source_credential_name value.
regex_filter	Specifies the REGEX expression to filter files. The REGEX expression pattern must be compatible with the REGEXP_LIKE operator. If you do not supply a regex_filter value, the regex_filter is set to NULL. See REGEXP_LIKE Condition for more information.

Parameter	Description
format	Specifies the additional configuration options for the file operation. These options are specified as a JSON string. The supported format options are: • logretention: It accepts an integer value that determines the duration in days for which the status table is retained for a bulk operation. The default value is 2 days. • logprefix: It accepts a string value that determines the bulk operation status table name prefix string. The operation type is the default value. For BULK_COPY, the default logprefix value is COPYOBJ. • priority: It accepts a string value that determines the number of file operations performed concurrently. An operation with a higher priority consumes more database resources and should run faster. It accepts the following values: – HIGH: Determines the number of parallel files handled using the database's **ECPU count** (**OCPU count** if your database uses OCPUs) – MEDIUM: Determines the number of simultaneous processes using the concurrency limit for Medium service. The default value is 4. – LOW: Process the files in serial order. The default value is MEDIUM. The maximum number of concurrent file operations is limited to 64. If you do not supply a format value, the format is set to NULL.
operation_id	Use this parameter to track the progress and final status of the load operation as the corresponding ID in the USER_LOAD_OPERATIONS view.

Usage Notes

- An error is returned when the source and target URI point to the same Object Storage bucket or folder.

Example

```
BEGIN
DBMS_CLOUD.BULK_COPY (
     source_credential_name => 'OCI_CRED',
     source_location_uri    => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname1/o',
     target_location_uri    => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname2/o',
     format       => JSON_OBJECT ('logretention' value 7, 'logprefix' value
'BULKOP')
);
END;
/
```

BULK_DELETE Procedure

This procedure bulk deletes files from the Cloud Object Storage. The overloaded form enables you to use the `operation_id` parameter. You can filter the list of files to be deleted using a regular expression pattern compatible with `REGEXP_LIKE` operator.

Syntax

```
 DBMS_CLOUD.BULK_DELETE(
     credential_name  IN  VARCHAR2 DEFAULT NULL,
     location_uri     IN  VARCHAR2,
     regex_filter     IN  VARCHAR2 DEFAULT NULL,
     format           IN  CLOB     DEFAULT NULL
);

DBMS_CLOUD.BULK_DELETE (
     credential_name  IN  VARCHAR2 DEFAULT NULL,
     location_uri     IN  VARCHAR2,
     regex_filter     IN  VARCHAR2 DEFAULT NULL,
     format           IN  CLOB     DEFAULT NULL,
     operation_id     OUT NUMBER
);
```

Parameters

Parameter	Description
`credential_name`	The name of the credential to access the Cloud Object Storage. You can use `'OCI$RESOURCE_PRINCIPAL'` as the `credential_name` when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. If you do not supply a `credential_name` value, the `credential_name` is set to `NULL`.
`location_uri`	Specifies URI, that point to an Object Storage location in the Autonomous Database. This parameter is mandatory. The format of the URIs depends on the Cloud Object Storage service. See DBMS_CLOUD URI Formats for more information.
`regex_filter`	Specifies the REGEX expression to filter files. The REGEX expression pattern must be compatible with the `REGEXP_LIKE` operator. If you do not supply a `regex_filter` value, the `regex_filter` is set to `NULL`. See REGEXP_LIKE Condition for more information.

Parameter	Description
format	Specifies the additional configuration options for the file operation. These options are specified as a JSON string. The supported format options are: • logretention: It accepts an integer value that determines the duration in days for which the status table is retained for a bulk operation . The default value is 2 days. • logprefix: It accepts a string value that determines the bulk operation status table name prefix string. The operation type is the default value. For BULK_DELETE, the default logprefix value is DELETE. • priority: It accepts a string value that determines the number of file operations performed concurrently. An operation with a higher priority consumes more database resources and is completed sooner. It accepts the following values: – HIGH: Determines the number of parallel files handled using the database's **ECPU count** (**OCPU count** if your database uses OCPUs). – MEDIUM: Determines the number of simultaneous processes using the concurrency limit for Medium service. The default value is 4. – LOW: Process the files in serial order. The default value is MEDIUM. The maximum number of concurrent file operations is limited to 64. If you do not supply a format value, the format is set to NULL.
operation_id	Use this parameter to track the progress and final status of the load operation as the corresponding ID in the USER_LOAD_OPERATIONS view.

Example

```
BEGIN
DBMS_CLOUD.BULK_DELETE (
     credential_name => 'OCI_CRED',
     location_uri    => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o',
     format          => JSON_OBJECT ('logretention' value 5, 'logprefix'
value 'BULKDEL')
);
END;
/
```

BULK_DOWNLOAD Procedure

This procedure downloads files into an Autonomous Database directory from Cloud Object Storage. The overloaded form enables you to use the operation_id parameter. You can filter the list of files to be downloaded using a regular expression pattern compatible with REGEXP_LIKE operator.

Syntax

```
DBMS_CLOUD.BULK_DOWNLOAD (
     credential_name  IN  VARCHAR2 DEFAULT NULL,
     location_uri     IN  VARCHAR2,
     directory_name   IN  VARCHAR2,
     regex_filter     IN  VARCHAR2 DEFAULT NULL,
     format           IN  CLOB     DEFAULT NULL
);

DBMS_CLOUD.BULK_DOWNLOAD (
     credential_name  IN  VARCHAR2 DEFAULT NULL,
     location_uri     IN  VARCHAR2,
     directory_name   IN  VARCHAR2,
     regex_filter     IN  VARCHAR2 DEFAULT NULL,
     format           IN  CLOB     DEFAULT NULL,
     operation_id     OUT NUMBER
);
```

Parameters

Parameter	Description
`credential_name`	The name of the credential to access the Cloud Object Storage. You can use `'OCI$RESOURCE_PRINCIPAL'` as the `credential_name` when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. If you do not supply a `credential_name` value, the `credential_name` is set to `NULL`.
`location_uri`	Specifies URI, that point to an Object Storage location in the Autonomous Database. This parameter is mandatory. The format of the URIs depends on the Cloud Object Storage service. See DBMS_CLOUD URI Formats for more information.
`directory_name`	The name of the directory on the Autonomous Database from where you want to download the files. This parameter is mandatory.
`regex_filter`	Specifies the REGEX expression to filter files. The REGEX expression pattern must be compatible with the `REGEXP_LIKE` operator. If you do not supply a `regex_filter` value, the `regex_filter` is set to `NULL`. See REGEXP_LIKE Condition for more information.

Parameter	Description
format	Specifies the additional configuration options for the file operation. These options are specified as a JSON string. The supported format options are: • logretention: It accepts an integer value that determines the duration in days for which the status table is retained for a bulk operation. The default value is 2 days. • logprefix: It accepts a string value that determines the bulk operation status table name prefix string. For BULK_DOWNLOAD, the default logprefix value is DOWNLOAD. The operation type is the default value. • priority: It accepts a string value that determines the number of file operations performed concurrently. An operation with a higher priority consumes more database resources and is completed sooner. It accepts the following values: – HIGH: Determines the number of parallel files handled using the database's **ECPU count** (**OCPU count** if your database uses OCPUs). – MEDIUM: Determines the number of simultaneous processes using the concurrency limit for Medium service. The default value is 4. – LOW: Process the files in serial order. The default value is MEDIUM. The maximum number of concurrent file operations is limited to 64. If you do not supply a format value, the format is set to NULL.
operation_id	Use this parameter to track the progress and final status of the load operation as the corresponding ID in the USER_LOAD_OPERATIONS view.

Example

```
BEGIN
DBMS_CLOUD.BULK_DOWNLOAD (
     credential_name => 'OCI_CRED',
     location_uri    => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o',
     directory_name  => 'BULK_TEST',
     format          => JSON_OBJECT ('logretention' value 7, 'logprefix'
value 'BULKOP')
 );
END;
/
```

BULK_MOVE Procedure

This procedure bulk moves files from one Cloud Object Storage bucket or folder to another. The overloaded form enables you to use the operation_id parameter.

You can filter the list of files to be deleted using a regular expression pattern compatible with REGEXP_LIKE operator.

The source and target bucket or folder can be in the same or different Cloud Object store provider.

When the source and target are in distinct Object Stores or have different accounts with the same cloud provider, you can give separate credential names for the source and target locations.

The source credential name is by default also used by the target location when target credential name is not provided.

The first step in moving files is copying them to the target location, then deleting the source files, once they are successfully copied.

The object is renamed rather than moved if Object Store allows renaming operations between source and target locations.

Syntax

```
DBMS_CLOUD.BULK_MOVE (
      source_credential_name  IN  VARCHAR2 DEFAULT NULL,
      source_location_uri     IN  VARCHAR2,
      target_location_uri     IN  VARCHAR2,
      target_credential_name  IN  VARCHAR2 DEFAULT NULL,
      regex_filter            IN  VARCHAR2 DEFAULT NULL,
      format                  IN  CLOB     DEFAULT NULL
);

DBMS_CLOUD.BULK_MOVE (
      source_credential_name  IN  VARCHAR2 DEFAULT NULL,
      source_location_uri     IN  VARCHAR2,
      target_location_uri     IN  VARCHAR2,
      target_credential_name  IN  VARCHAR2 DEFAULT NULL,
      regex_filter            IN  VARCHAR2 DEFAULT NULL,
      format                  IN  CLOB     DEFAULT NULL,
      operation_id            OUT NUMBER
);
```

Parameters

Parameter	Description
`source_credential_name`	The name of the credential to access the source Cloud Object Storage. You can use `'OCI$RESOURCE_PRINCIPAL'` as the `credential_name` when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. If you do not supply a `source_credential_name` value, the `credential_name` is set to `NULL`.
`source_location_uri`	Specifies URI, that point to the source Object Storage bucket or folder location. This parameter is mandatory. The format of the URIs depends on the Cloud Object Storage service. See DBMS_CLOUD URI Formats for more information.

Parameter	Description
target_location_uri	Specifies the URI for the target Object Storage bucket or folder, where the files need to be moved. This parameter is mandatory. The format of the URIs depends on the Cloud Object Storage service. See DBMS_CLOUD URI Formats for more information.
target_credential_name	The name of the credential to access the target Cloud Object Storage location. You can use 'OCI$RESOURCE_PRINCIPAL' as the credential_name when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. If you do not supply a target_credential_name value, the target_location_uri is set to the source_credential_name value.
regex_filter	Specifies the REGEX expression to filter files. The REGEX expression pattern must be compatible with the REGEXP_LIKE operator. If you do not supply a regex_filter value, the regex_filter is set to NULL. See REGEXP_LIKE Condition for more information.
format	Specifies the additional configuration options for the file operation. These options are specified as a JSON string. The supported format options are: • logretention: It accepts an integer value that determines the duration in days for which the status table is retained for a bulk operation. The default value is 2 days. • logprefix: It accepts a string value that determines the bulk operation status table name prefix string. The operation type is the default value. For BULK_MOVE, the default logprefix value is MOVE. • priority: It accepts a string value that determines the number of file operations performed concurrently. An operation with a higher priority consumes more database resources and is completed sooner. It accepts the following values: – HIGH: Determines the number of parallel files handled using the database's **ECPU count** (**OCPU count** if your database uses OCPUs). – MEDIUM: Determines the number of simultaneous processes using the concurrency limit for Medium service. The default value is 4. – LOW: Process the files in serial order. The default value is MEDIUM. The maximum number of concurrent file operations is limited to 64. If you do not supply a format value, the format is set to NULL.
operation_id	Use this parameter to track the progress and final status of the load operation as the corresponding ID in the USER_LOAD_OPERATIONS view.

Example

```
BEGIN
DBMS_CLOUD.BULK_MOVE (
```

```
     source_credential_name => 'OCI_CRED',
     source_location_uri    => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname1/o',
     target_location_uri    => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname2/o',
     format                 => JSON_OBJECT ('logretention' value 7,
'logprefix' value 'BULKMOVE')
);
END;
/
```

> **Note:**
>
> An error is returned when the source and target URI point to the same Object Storage bucket or folder.

BULK_UPLOAD Procedure

This procedure copies files into Cloud Object Storage from an Autonomous Database directory. The overloaded form enables you to use the `operation_id` parameter.

Syntax

```
DBMS_CLOUD.BULK_UPLOAD (
     credential_name  IN  VARCHAR2 DEFAULT NULL,
     location_uri     IN  VARCHAR2,
     directory_name   IN  VARCHAR2,
     regex_filter     IN  VARCHAR2 DEFAULT NULL,
     format           IN  CLOB     DEFAULT NULL
);

DBMS_CLOUD.BULK_UPLOAD (
     credential_name  IN  VARCHAR2 DEFAULT NULL,
     location_uri     IN  VARCHAR2,
     directory_name   IN  VARCHAR2,
     regex_filter     IN  VARCHAR2 DEFAULT NULL,
     format           IN  CLOB     DEFAULT NULL,
     operation_id     OUT NUMBER
);
```

Parameters

Parameter	Description
`credential_name`	The name of the credential to access the Cloud Object Storage. You can use `'OCI$RESOURCE_PRINCIPAL'` as the `credential_name` when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information. If you do not supply a `credential_name` value, the `credential_name` is set to `NULL`.

Parameter	Description
location_uri	Specifies URI, that points to an Object Storage location to upload files. This parameter is mandatory. The format of the URIs depends on the Cloud Object Storage service. See DBMS_CLOUD URI Formats for more information.
directory_name	The name of the directory on the Autonomous Database from where you upload files. This parameter is mandatory.
regex_filter	Specifies the REGEX expression to filter files. The REGEX expression pattern must be compatible with REGEXP_LIKE operator. If you do not supply a regex_filter value, the regex_filter is set to NULL. See REGEXP_LIKE Condition for more information.
format	Specifies the additional configuration options for the file operation. These options are specified as a JSON string. The supported format options are: • logretention: It accepts an integer value that determines the duration in days for which the status table is retained for a bulk operation. The default value is 2 days. • logprefix: It accepts a string value that determines the bulk operation status table name prefix string. The operation type is the default value. For BULK_UPLOAD, the default logprefix value is UPLOAD. • priority: It accepts a string value that determines the number of file operations performed concurrently. An operation with a higher priority consumes more database resources and is completed sooner. It accepts the following values: – HIGH: Determines the number of parallel files handled using the database's **ECPU count** (**OCPU count** if your database uses OCPUs). – MEDIUM: Determines the number of simultaneous processes using the concurrency limit for Medium service. The default value is 4. – LOW: Process the files in serial order. The default value is MEDIUM. The maximum number of concurrent file operations is limited to 64. If you do not supply a format value, the format is set to NULL.
operation_id	Use this parameter to track the progress and final status of the load operation as the corresponding ID in the USER_LOAD_OPERATIONS view.

Example

```
BEGIN
DBMS_CLOUD.BULK_UPLOAD (
     credential_name => 'OCI_CRED',
     location_uri    => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o',
     directory_name  => 'BULK_TEST',
     format          => JSON_OBJECT ('logretention' value 5, 'logprefix'
```

```
value 'BULKUPLOAD')
 );
END;
/
```

DBMS_CLOUD REST APIs

This section covers the `DBMS_CLOUD` REST APIs provided with Autonomous Database.

REST API	Description
GET_RESPONSE_HEADERS Function	This function returns the HTTP response headers as JSON data in a JSON object in Autonomous Database.
GET_RESPONSE_RAW Function	This function returns the HTTP response in RAW format Autonomous Database. This is useful if the HTTP response is expected to be binary format.
GET_RESPONSE_STATUS_CODE Function	This function returns the HTTP response status code as an integer in Autonomous Database. The status code helps to identify if the request is successful.
GET_RESPONSE_TEXT Function	This function returns the HTTP response in TEXT format (`VARCHAR2` or `CLOB`) in Autonomous Database. Usually, most Cloud REST APIs return JSON response in text format. This function is useful if you expect the HTTP response is in text format.
GET_API_RESULT_CACHE_SIZE Function	This function returns the configured result cache size.
SEND_REQUEST Function and Procedure	This function begins an HTTP request, gets the response, and ends the response in Autonomous Database. This function provides a workflow for sending a Cloud REST API request with arguments and a return response code and payload.
SET_API_RESULT_CACHE_SIZE Procedure	This procedure sets the maximum cache size for current session.

- DBMS_CLOUD REST API Overview
  When you use PL/SQL in your application and you need to call Cloud REST APIs you can use `DBMS_CLOUD.SEND_REQUEST` to send the REST API requests.
- DBMS_CLOUD REST API Constants
  Describes the `DBMS_CLOUD` constants for making `HTTP` requests using `DBMS_CLOUD.SEND_REQUEST`.
- DBMS_CLOUD REST API Results Cache
  You can save `DBMS_CLOUD` REST API results when you set the `cache` parameter to true with `DBMS_CLOUD.SEND_REQUEST`. The `SESSION_CLOUD_API_RESULTS` view describes the columns you can use when REST API results are saved.
- GET_RESPONSE_HEADERS Function
  This function returns the HTTP response headers as JSON data in a JSON object.
- GET_RESPONSE_RAW Function
  This function returns the HTTP response in RAW format. This is useful if the HTTP response is expected to be binary format.
- GET_RESPONSE_STATUS_CODE Function
  This function returns the HTTP response status code as an integer. The status code helps to identify if the request is successful.

- GET_RESPONSE_TEXT Function
  This function returns the HTTP response in TEXT format (VARCHAR2 or CLOB). Usually, most Cloud REST APIs return JSON response in text format. This function is useful if you expect the HTTP response is in text format.
- GET_API_RESULT_CACHE_SIZE Function
  This function returns the configured result cache size. The cache size value only applies for the current session.
- SEND_REQUEST Function and Procedure
  This function and procedure begins an HTTP request, gets the response, and ends the response. This function provides a workflow for sending a cloud REST API request with arguments and the function returns a response code and payload. If you use the procedure, you can view results and response details from the saved results with the SESSION_CLOUD_API_RESULTS view.
- SET_API_RESULT_CACHE_SIZE Procedure
  This procedure sets the maximum cache size for current session. The cache size value only applies for the current session.
- DBMS_CLOUD REST API Examples
  Shows examples using DBMS_CLOUD.SEND_REQUEST to create and delete an Oracle Cloud Infrastructure Object Storage bucket, and an example to list all compartments in the tenancy.

DBMS_CLOUD REST API Overview

When you use PL/SQL in your application and you need to call Cloud REST APIs you can use DBMS_CLOUD.SEND_REQUEST to send the REST API requests.

The DBMS_CLOUD REST API functions allow you to make HTTP requests using DBMS_CLOUD.SEND_REQUEST and obtain and save results. These functions provide a generic API that lets you call any REST API with the following supported cloud services:

- Oracle Cloud Infrastructure
  See API Reference and Endpoints for information on Oracle Cloud Infrastructure REST APIs.
- Amazon Web Services (AWS)
  See Guides and API References for information on Amazon Web Services REST APIs.
- Azure Cloud [1]
  See Azure REST API Reference for information on Azure REST APIs.
- Oracle Cloud Infrastructure Classic
  See All REST Endpoints for information on Oracle Cloud Infrastructure Classic REST APIs.
- GitHub Repository
  See GitHub REST API for more information.

DBMS_CLOUD REST API Constants

Describes the DBMS_CLOUD constants for making HTTP requests using DBMS_CLOUD.SEND_REQUEST.

DBMS_CLOUD supports GET, PUT, POST, HEAD and DELETE HTTP methods. The REST API method to be used for an HTTP request is typically documented in the Cloud REST API documentation.

[1] Support for Azure Cloud REST API calls is limited to the domain "blob.windows.net".

Name	Type	Value
METHOD_DELETE	VARCHAR2(6)	'DELETE'
METHOD_GET	VARCHAR2(3)	'GET'
METHOD_HEAD	VARCHAR2(4)	'HEAD'
METHOD_POST	VARCHAR2(4)	'POST'
METHOD_PUT	VARCHAR2(3)	'PUT'

DBMS_CLOUD REST API Results Cache

You can save DBMS_CLOUD REST API results when you set the cache parameter to true with DBMS_CLOUD.SEND_REQUEST. The SESSION_CLOUD_API_RESULTS view describes the columns you can use when REST API results are saved.

By default DBMS_CLOUD REST API calls do not save results for your session. In this case you use the DBMS_CLOUD.SEND_REQUEST function to return results.

When you use DBMS_CLOUD.SEND_REQUEST and set the cache parameter to TRUE, results are saved and you can view past results in the SESSION_CLOUD_API_RESULTS view. Saving and querying historical results of DBMS_CLOUD REST API requests can help you when you need to work with your previous results in your applications.

For example, to query recent DBMS_CLOUD REST API results, use the view SESSION_CLOUD_API_RESULTS:

```
SELECT timestamp FROM SESSION_CLOUD_API_RESULTS;
```

When you save DBMS_CLOUD REST API results with DBMS_CLOUD.SEND_REQUEST the saved data is only available within the same session (connection). After the session exits, the saved data is no longer available.

Use DBMS_CLOUD.GET_API_RESULT_CACHE_SIZE and DBMS_CLOUD.SET_API_RESULT_CACHE_SIZE to view and set the DBMS_CLOUD REST API cache size, and to disable caching.

- DBMS_CLOUD REST API Results cache_scope Parameter
  When you save DBMS_CLOUD REST API results with DBMS_CLOUD.SEND_REQUEST, access to the results in SESSION_CLOUD_API_RESULTS is provided based on the value of cache_scope.
- DBMS_CLOUD REST API SESSION_CLOUD_API_RESULTS View
  You can save DBMS_CLOUD REST API results when you set the cache parameter to true with DBMS_CLOUD.SEND_REQUEST. The SESSION_CLOUD_API_RESULTS view describes the columns you can use when REST API results are saved.

DBMS_CLOUD REST API Results cache_scope Parameter

When you save DBMS_CLOUD REST API results with DBMS_CLOUD.SEND_REQUEST, access to the results in SESSION_CLOUD_API_RESULTS is provided based on the value of cache_scope.

By default cache_scope is 'PRIVATE' and only the current user of the session can access the results. If you set the cache_scope to 'PUBLIC', then all session users can access the results. The default value for cache_scope specifies that each user can only see DBMS_CLOUD.SEND_REQUEST REST API results generated by the procedures they invoke with invoker's rights. When you invoke DBMS_CLOUD.SEND_REQUEST in a session, there are three possibilities that determines if the current user can see results in the cache, based on the cache_scope value:

- You directly execute `DBMS_CLOUD.SEND_REQUEST` as a top-level statement and the call to `DBMS_CLOUD.SEND_REQUEST` and the REST API results are saved with the same username. In this case you have access to all results with the default value, `'PRIVATE'`, set for `cache_scope`.
- You write a wrapper invoker's rights procedure and as the current user your call with `DBMS_CLOUD.SEND_REQUEST` calls the procedure and the REST API results are saved with the same username. In this case, and you have access to all results with the default value, `'PRIVATE'`, set for `cache_scope`.
- You write a wrapper definer's rights procedure and the procedure is owned by another user. When you call `DBMS_CLOUD.SEND_REQUEST` inside the procedure, the results are saved with the username of the procedure owner.

  For this case, a different definer's rights user is invoking `DBMS_CLOUD.SEND_REQUEST`, and the REST API results are saved with that definers procedure's owner. For this case, by default when `cache_scope` is `PRIVATE'`, the invoker's session cannot see the results.

  If the definer's procedure owner wants to make the results available to any invoking session user, then they must set `cache_scope` to `'PUBLIC'` in the `DBMS_CLOUD.SEND_REQUEST`.

DBMS_CLOUD REST API SESSION_CLOUD_API_RESULTS View

You can save `DBMS_CLOUD` REST API results when you set the `cache` parameter to true with `DBMS_CLOUD.SEND_REQUEST`. The `SESSION_CLOUD_API_RESULTS` view describes the columns you can use when REST API results are saved.

The view `SESSION_CLOUD_API_RESULTS` is the view created if you cache results with `DBMS_CLOUD.SEND_REQUEST`. You can query historical results which belong to your user session. When the session ends, the data in the `SESSION_CLOUD_API_RESULTS` is purged.

Column	Description
`URI`	The `DBMS_CLOUD` REST API request URL
`TIMESTAMP`	The `DBMS_CLOUD` REST API response timestamp
`CLOUD_TYPE`	The `DBMS_CLOUD` REST API cloud type, such as Oracle Cloud Infrastructure, AMAZON_S3, and AZURE_BLOB
`REQUEST_METHOD`	The `DBMS_CLOUD` REST API request method, such as `GET`, `PUT`, `HEAD`
`REQUEST_HEADERS`	The `DBMS_CLOUD` REST API request headers
`REQUEST_BODY_TEXT`	The `DBMS_CLOUD` REST API request body in `CLOB`
`RESPONSE_STATUS_CODE`	The `DBMS_CLOUD` REST API response status code, such as `200(OK)`, `404(Not Found)`
`RESPONSE_HEADERS`	The `DBMS_CLOUD` REST API response headers
`RESPONSE_BODY_TEXT`	The `DBMS_CLOUD` REST API response body in `CLOB`
`SCOPE`	The `cache_scope` set by `DBMS_CLOUD.SEND_REQUEST`. Valid values are `PUBLIC` or `PRIVATE`.

GET_RESPONSE_HEADERS Function

This function returns the HTTP response headers as JSON data in a JSON object.

Syntax

```
DBMS_CLOUD.GET_RESPONSE_HEADERS(
       resp           IN DBMS_CLOUD_TYPES.resp)
   RETURN JSON_OBJECT_T;
```

Parameters

Parameter	Description
resp	HTTP Response type returned from DBMS_CLOUD.SEND_REQUEST.

Exceptions

Exception	Error	Description
invalid_response	ORA-20025	Invalid response type object passed to DBMS_CLOUD.GET_RESPONSE_HEADERS.

GET_RESPONSE_RAW Function

This function returns the HTTP response in RAW format. This is useful if the HTTP response is expected to be binary format.

Syntax

```
DBMS_CLOUD.GET_RESPONSE_RAW(
       resp           IN DBMS_CLOUD_TYPES.resp)
   RETURN BLOB;
```

Parameters

Parameter	Description
resp	HTTP Response type returned from DBMS_CLOUD.SEND_REQUEST.

Exceptions

Exception	Error	Description
invalid_response	ORA-20025	Invalid response type object passed to DBMS_CLOUD.GET_RESPONSE_RAW.

GET_RESPONSE_STATUS_CODE Function

This function returns the HTTP response status code as an integer. The status code helps to identify if the request is successful.

Syntax

```
DBMS_CLOUD.GET_RESPONSE_STATUS_CODE(
       resp           IN DBMS_CLOUD_TYPES.resp)
   RETURN PLS_INTEGER;
```

Parameters

Parameter	Description
resp	HTTP Response type returned from DBMS_CLOUD.SEND_REQUEST.

Exceptions

Exception	Error	Description
invalid_response	ORA-20025	Invalid response type object passed to DBMS_CLOUD.GET_RESPONSE_STATUS_CODE.

GET_RESPONSE_TEXT Function

This function returns the HTTP response in TEXT format (VARCHAR2 or CLOB). Usually, most Cloud REST APIs return JSON response in text format. This function is useful if you expect the HTTP response is in text format.

Syntax

```
DBMS_CLOUD.GET_RESPONSE_TEXT(
       resp           IN DBMS_CLOUD_TYPES.resp)
   RETURN CLOB;
```

Parameters

Parameter	Description
resp	HTTP Response type returned from DBMS_CLOUD.SEND_REQUEST.

Exceptions

Exception	Error	Description
invalid_response	ORA-20025	Invalid response type object passed to DBMS_CLOUD.GET_RESPONSE_TEXT.

GET_API_RESULT_CACHE_SIZE Function

This function returns the configured result cache size. The cache size value only applies for the current session.

Syntax

```
DBMS_CLOUD.GET_API_RESULT_CACHE_SIZE()
   RETURN NUMBER;
```

SEND_REQUEST Function and Procedure

This function and procedure begins an HTTP request, gets the response, and ends the response. This function provides a workflow for sending a cloud REST API request with arguments and the function returns a response code and payload. If you use the procedure, you can view results and response details from the saved results with the `SESSION_CLOUD_API_RESULTS` view.

Syntax

```
DBMS_CLOUD.SEND_REQUEST(
       credential_name    IN VARCHAR2,
       uri                IN VARCHAR2,
       method             IN VARCHAR2,
       headers            IN CLOB DEFAULT NULL,
       async_request_url  IN VARCHAR2 DEFAULT NULL,
       wait_for_states    IN DBMS_CLOUD_TYPES.wait_for_states_t DEFAULT NULL,
       timeout            IN NUMBER DEFAULT 0,
       cache              IN PL/SQL BOOLEAN DEFAULT FALSE,
       cache_scope        IN VARCHAR2 DEFAULT 'PRIVATE',
       body               IN BLOB DEFAULT NULL)
   RETURN DBMS_CLOUD_TYPES.resp;

DBMS_CLOUD.SEND_REQUEST(
       credential_name    IN VARCHAR2,
       uri                IN VARCHAR2,
       method             IN VARCHAR2,
       headers            IN CLOB DEFAULT NULL,
       async_request_url  IN VARCHAR2 DEFAULT NULL,
       wait_for_states    IN DBMS_CLOUD_TYPES.wait_for_states_t DEFAULT NULL,
       timeout            IN NUMBER DEFAULT 0,
       cache              IN PL/SQL BOOLEAN DEFAULT FALSE,
       cache_scope        IN VARCHAR2 DEFAULT 'PRIVATE',
       body               IN BLOB DEFAULT NULL);
```

Parameters

Parameter	Description
`credential_name`	The name of the credential for authenticating with the corresponding cloud native API. You can use `'OCI$RESOURCE_PRINCIPAL'` as the `credential_name` when resource principal is enabled. See ENABLE_RESOURCE_PRINCIPAL for more information.
`uri`	HTTP URI to make the request.
`method`	HTTP Request Method: `GET`, `PUT`, `POST`, `HEAD`, `DELETE`. Use the `DBMS_CLOUD` package constant to specify the method. See DBMS_CLOUD REST API Constants for more information.

Parameter	Description
headers	HTTP Request headers for the corresponding cloud native API in JSON format. The authentication headers are set automatically, only pass custom headers.
async_request_url	An asynchronous request URL. To obtain the URL select your request API from the list of APIs (see https://docs.cloud.oracle.com/en-us/iaas/api/). Then, navigate to find the API for your request in the left pane. For example, Database Services API → Autonomous Database → StopAutonomousDatabase. This page shows the API home (and shows the base endpoint). Then, append the base endpoint with the relative path obtained for your work request WorkRequest link.
wait_for_states	Wait for states is a status of type: DBMS_CLOUD_TYPES.wait_for_states_t. The following are valid values for expected states: 'ACTIVE', 'CANCELED', 'COMPLETED', 'DELETED', 'FAILED', 'SUCCEEDED'. Multiple states are allowed for wait_for_states. The default value for wait_for_states is to wait for any of the expected states: 'ACTIVE', 'CANCELED', 'COMPLETED', 'DELETED', 'FAILED', 'SUCCEEDED'.
timeout	Specifies the timeout, in seconds, for asynchronous requests with the parameters async_request_url and wait_for_states. Default value is 0. This indicates to wait for completion of the request without any timeout.
cache	If TRUE specifies the request should be cached in REST result API cache. The default value is FALSE, which means REST API requests are not cached.
cache_scope	Specifies whether everyone can have access to this request result cache. Valid values: "PRIVATE" and "PUBLIC". The default value is "PRIVATE".
body	HTTP Request Body for PUT and POST requests.

Exceptions

Exception	Error	Description
invalid_req_method	ORA-20023	Request method passed to DBMS_CLOUD.SEND_REQUEST is invalid.
invalid_req_header	ORA-20024	Request headers passed to DBMS_CLOUD.SEND_REQUEST are not in valid JSON format.

Usage Notes

- If you are using Oracle Cloud Infrastructure, you must use a Signing Key based credential value for the credential_name. See CREATE_CREDENTIAL Procedure for more information.
- The optional parameters async_request_url, wait_for_states, and timeout allow you to handle long running requests. Using this asynchronous form of send_request, the function waits for the completion status specified in wait_for_states before returning. With these parameters in the send request, you pass the expected return states in the wait_for_states parameter, and you use the async_request_url parameter to specify an

associated work request, the request does not return immediately. Instead, the request probes the `async_request_url` until the return state is one of the expected states or the `timeout` is exceeded (`timeout` is optional). If no `timeout` is specified, the request waits until a state found in `wait_for_states` occurs.

SET_API_RESULT_CACHE_SIZE Procedure

This procedure sets the maximum cache size for current session. The cache size value only applies for the current session.

Syntax

```
DBMS_CLOUD.SET_API_RESULT_CACHE_SIZE(
       cache_size          IN NUMBER);
```

Parameters

Parameter	Description
`cache_size`	Set the maximum cache size to the specified value `cache_size`. If the new maximum cache size is smaller than the current cache size, older records are dropped until the number of rows is equal to the specified maximum cache size. The maximum value is `10000`. If the cache size is set to `0`, caching is disabled in the session. The default cache size is `10`.

Exceptions

Exception	Error	Description
`invalid API result cache size`	`ORA-20032`	The minimum value is 0 and the maximum value is 10000. This exception is shown when the input value is less than 0 or is larger than 10000.

Example

```
EXEC DBMS_CLOUD.SET_API_RESULT_CACHE_SIZE(101);
```

DBMS_CLOUD REST API Examples

Shows examples using `DBMS_CLOUD.SEND_REQUEST` to create and delete an Oracle Cloud Infrastructure Object Storage bucket, and an example to list all compartments in the tenancy.

> **Note:**
>
> These examples show Oracle Cloud Infrastructure request APIs and require that you use a Signing Key based credential for the `credential_name`. Oracle Cloud Infrastructure Signing Key based credentials include the `private_key` and `fingerprint` arguments.
>
> For example:
>
> ```
> BEGIN
>    DBMS_CLOUD.CREATE_CREDENTIAL (
>        credential_name => 'OCI_KEY_CRED',
>        user_ocid       =>
> 'ocid1.user.oc1..aaaaaaaauq54mi7zdyfhw33ozkwuontjceel7fok5nq3bf2vwetkp
> qsoa',
>        tenancy_ocid    =>
> 'ocid1.tenancy.oc1..aabbbbbbaafcue47pqmrf4vigneebgbcmmoy5r7xvoypicjqqg
> e32ewnrcyx2a',
>        private_key     =>
> 'MIIEogIBAAKCAQEAtUnxbmrekwgVac6FdWeRzoXvIpA9+0r1.....wtnNpESQQQ0QLGPD
> 8NM//JEBg=',
>        fingerprint     =>
> 'f2:db:f9:18:a4:aa:fc:94:f4:f6:6c:39:96:16:aa:27');
> END;
> /
> ```
>
> See CREATE_CREDENTIAL Procedure for information on `DBMS_CLOUD.CREATE_CREDENTIAL`.

Create Bucket Example

Shows an example using `DBMS_CLOUD.SEND_REQUEST` with HTTP `POST` method to create an object store bucket named *`bucketname`*.

See CreateBucket for details on the Oracle Cloud Infrastructure Object Storage Service API for this example.

```
SET SERVEROUTPUT ON
DECLARE
  resp DBMS_CLOUD_TYPES.resp;
BEGIN
  -- Send request
  resp := DBMS_CLOUD.send_request(
            credential_name => 'OCI_KEY_CRED',
            uri => 'https://objectstorage.region.oraclecloud.com/n/namespace-
string/b/',
            method => DBMS_CLOUD.METHOD_POST,
            body => UTL_RAW.cast_to_raw(
                        JSON_OBJECT('name' value 'bucketname',
                                    'compartmentId' value 'compartment_OCID'))
          );

  -- Response Body in TEXT format
  dbms_output.put_line('Body: ' || '------------' || CHR(10) ||
```

```
  DBMS_CLOUD.get_response_text(resp) || CHR(10));

  -- Response Headers in JSON format
  dbms_output.put_line('Headers: ' || CHR(10) || '------------' || CHR(10) ||
  DBMS_CLOUD.get_response_headers(resp).to_clob || CHR(10));

  -- Response Status Code
  dbms_output.put_line('Status Code: ' || CHR(10) || '------------' ||
CHR(10) ||
  DBMS_CLOUD.get_response_status_code(resp));

END;
/
```

Notes:

- In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.
- Where: *region* is an endpoint region. See Object Storage API reference in API Reference and Endpoints for more information. For example, where *region* is: `us-phoenix-1`.

Delete Bucket Example

Shows an example using `DBMS_CLOUD.SEND_REQUEST` with HTTP `DELETE` method to delete an object store bucket named *bucketname*.

See DeleteBucket for details on the Oracle Cloud Infrastructure Object Storage Service API for this example.

```
SET SERVEROUTPUT ON
DECLARE
  resp DBMS_CLOUD_TYPES.resp;
BEGIN
  -- Send request
  resp := DBMS_CLOUD.send_request(
            credential_name => 'OCI_KEY_CRED',
            uri => 'https://objectstorage.region.oraclecloud.com/n/namespace-
string/b/bucketname',
            method => DBMS_CLOUD.METHOD_DELETE
          );

  -- Response Body in TEXT format
  dbms_output.put_line('Body: ' || '------------' || CHR(10) ||
  DBMS_CLOUD.get_response_text(resp) || CHR(10));

  -- Response Headers in JSON format
  dbms_output.put_line('Headers: ' || CHR(10) || '------------' || CHR(10) ||
  DBMS_CLOUD.get_response_headers(resp).to_clob || CHR(10));

  -- Response Status Code
  dbms_output.put_line('Status Code: ' || CHR(10) || '------------' ||
CHR(10) ||
  DBMS_CLOUD.get_response_status_code(resp));
```

```
END;
/
```

Notes:

- In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.
- Where: *region* is an endpoint region. See Object Storage API reference in API Reference and Endpoints for more information. For example, where *region* is: `us-phoenix-1`.

List Compartments Example

Shows an example using `DBMS_CLOUD.SEND_REQUEST` with HTTP `GET` method to list all compartments in the tenancy (root compartment). This example shows how to pass request headers in the `DBMS_CLOUD.SEND_REQUEST`.

See ListCompartments for details on the Oracle Cloud Infrastructure Identity and Access Management Service API for this example.

```
--
-- List compartments
--
DECLARE
  resp DBMS_CLOUD_TYPES.resp;
  root_compartment_ocid VARCHAR2(512) := '&1';
BEGIN
  -- Send request
  dbms_output.put_line('Send Request');
  resp := DBMS_CLOUD.send_request(
            credential_name => 'OCI_KEY_CRED',
            uri => 'https://identity.region.oraclecloud.com/20160918/
compartments?compartmentId=' || root_compartment_ocid,
            method => DBMS_CLOUD.METHOD_GET,
            headers => JSON_OBJECT('opc-request-id' value 'list-compartments')
          );
  dbms_output.put_line('Body: ' || '------------' || CHR(10) ||
DBMS_CLOUD.get_response_text(resp) || CHR(10));
  dbms_output.put_line('Headers: ' || CHR(10) || '------------' || CHR(10) ||
DBMS_CLOUD.get_response_headers(resp).to_clob || CHR(10));
  dbms_output.put_line('Status Code: ' || CHR(10) || '------------' ||
CHR(10) || DBMS_CLOUD.get_response_status_code(resp));
  dbms_output.put_line(CHR(10));
END;
/
```

Where: *region* is an endpoint region. See Identity and Access Management (IAM) API reference in API Reference and Endpoints for more information. For example, where *region* is: `uk-london-1`.

Asynchronous Request Example

Shows an example using `DBMS_CLOUD.SEND_REQUEST` with HTTP `POST` method to perform the Autonomous Database stop operation and wait for status. This example shows how to use

DBMS_CLOUD.SEND_REQUEST with the async_request_url, wait_for_states, and timeout parameters.

```
--
-- Sent Work Request Autonomous Database Stop Request with Wait for Status
DECLARE
    l_resp DBMS_CLOUD_TYPES.resp;
    l_resp_json JSON_OBJECT_T;
    l_key_shape JSON_OBJECT_T;
    l_body JSON_OBJECT_T;
    status_array DBMS_CLOUD_TYPES.wait_for_states_t;
BEGIN
  status_array := DBMS_CLOUD_TYPES.wait_for_states_t('SUCCEEDED');
  l_body := JSON_OBJECT_T('{}');
  l_body.put('autonomousDatabaseId', 'ocid');
-- Send request
  dbms_output.put_line(l_body.to_clob);
  dbms_output.put_line('Send Request');
  l_resp := DBMS_CLOUD.send_request(
                        credential_name    => 'NATIVE_CRED_OCI',
                        uri                => 'https://
database.region.oraclecloud.com/20160918/autonomousDatabases/ocid/actions/
stop',
                        method             => DBMS_CLOUD.METHOD_POST,
                        body               =>
UTL_RAW.cast_to_raw(l_body.to_clob),
                        async_request_url  => 'https://
iaas.region.oraclecloud.com/20160918/workRequests',
                        wait_for_states    => status_array,
                        timeout            => 600
                    );
   dbms_output.put_line('resp body: '||DBMS_CLOUD.get_response_text(l_resp));
   dbms_output.put_line('resp headers: '||
DBMS_CLOUD.get_response_headers(l_resp).to_clob);
END;
/
```

Where: *region* is an endpoint region. See Identity and Access Management (IAM) API reference in API Reference and Endpoints for more information. For example, where *region* is: uk-london-1.

The *ocid* is the Oracle Cloud Infrastructure resource identifier. See Resource Identifiers for more information.

DBMS_CLOUD URI Formats

Describes the format of the source file URIs in operations with DBMS_CLOUD. The format depends on the object storage service you are using.

DBMS_CLOUD guarantees secure communication and any URI that you specify must use HTTPS, with https:// as the prefix for the URI.

- Oracle Cloud Infrastructure Object Storage Native URI Format
- Oracle Cloud Infrastructure Object Storage Swift URI Format

- Oracle Cloud Infrastructure Object Storage URI Format Using Pre-Authenticated Request URL
- URI Format Using Public URL
- Oracle Cloud Infrastructure Object Storage Classic URI Format
- Amazon S3 URI Format
- Azure Blob Storage or Azure Data Lake Storage URI Format
- Amazon S3 Compatible URI Format
- GitHub Raw URL Format
  `DBMS_CLOUD` supports GitHub Raw URLs to access data from a GitHub Repository.
- Additional Customer-Managed URI Formats

Oracle Cloud Infrastructure Object Storage Native URI Format

If your source files reside on Oracle Cloud Infrastructure Object Storage in the commercial realm (OC1), it is recommended that you use the following URI format which uses Object Storage Dedicated Endpoints. See Object Storage Dedicated Endpoints, for further information.

```
https://namespace-string.objectstorage.region.oci.customer-oci.com/n/
namespace-string/b/bucketname/o/filename
```

> **Note:**
>
> OCI Object Store dedicated endpoint URLs are only supported in commercial realms (OC1).

If your source files reside on Oracle Cloud Infrastructure Object Storage and are not in the commercial realm (OC1), you must use the following format:

```
https://objectstorage.region.oraclecloud.com/n/namespace-string/b/bucket/o/
filename
```

For example, in the commercial realm (OC1) the Native URI for the file `channels.txt` in the *bucketname* bucket in the `Phoenix` data center is:

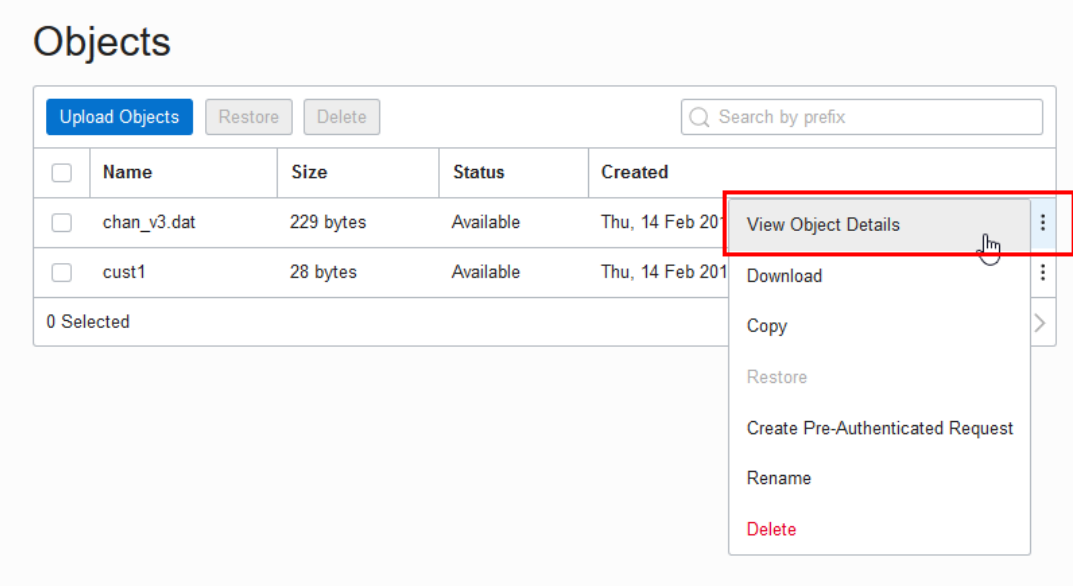
```
https://namespace.objectstorage.region.oci.customer-oci.com/n/namespace/b/
bucketname/o/channels.txt
```

In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

You can find the URI from the Oracle Cloud Infrastructure Object Storage "Object Details" in the right hand side ellipsis menu in the Object Store:

1. Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.

2. From the Oracle Cloud Infrastructure left navigation menu click **Core Infrastructure**. Under **Object Storage**, click **Object Storage**.
3. Under List Scope, select a **Compartment**.
4. From the **Name** column, select a bucket.
5. In the Objects area, click **View Object Details**.



6. On the **Object Details** page, the **URL Path (URI)** field shows the URI to access the object.

> **Note:**
>
> The source files need to be stored in an Object Storage tier bucket. Autonomous Database does not support buckets in the Archive Storage tier. See Overview of Object Storage for more information.

Oracle Cloud Infrastructure Object Storage Swift URI Format

If your source files reside on Oracle Cloud Infrastructure Object Storage in the commercial realm (OC1), it is recommended that you use the following URI format which uses Object Storage Dedicated Endpoints. See Object Storage Dedicated Endpoints, for further information.

```
https://namespace-string.swiftobjectstorage.region.oci.customer-oci.com/v1/
namespace-string/bucket/filename
```

> **Note:**
>
> OCI Object Store dedicated endpoint URLs are only supported in the commercial realms (OC1).

If your source files reside on Oracle Cloud Infrastructure Object Storage and are not in the commercial realm (OC1), you must use the following format:

```
https://swiftobjectstorage.region.oraclecloud.com/v1/namespace-string/bucket/
filename
```

For example, in the commercial realm (OC1) the Swift URI for the file `channels.txt` in the *bucketname* bucket in the `Phoenix` data center is:

```
https://namespace-string.swiftobjectstorage.us-phoenix-1.oci.customer-
oci.com/v1/namespace-string/bucketname/channels.txt
```

In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

> **Note:**
>
> The source files need to be stored in an Object Storage tier bucket. Autonomous Database does not support buckets in the Archive Storage tier. See Overview of Object Storage for more information.

Oracle Cloud Infrastructure Object Storage URI Format Using Pre-Authenticated Request URL

If your source files reside on the Oracle Cloud Infrastructure Object Storage you can use Oracle Cloud Infrastructure pre-authenticated URIs. When you create a pre-authenticated request, a unique URL is generated. You can then provide the unique URL to users in your organization, partners, or third parties to access the Object Storage resource target identified in the pre-authenticated request.

> **Note:**
>
> Carefully assess the business requirement for and the security ramifications of pre-authenticated access. When you create the pre-authenticated request URL, note the **Expiration** and the **Access type** to make sure they are appropriate for your use. A pre-authenticated request URL gives anyone who has the URL access to the targets identified in the request for as long as the request is active. In addition to considering the operational needs of pre-authenticated access, it is equally important to manage its distribution.

If your source files reside on Oracle Cloud Infrastructure Object Storage in the commercial realm (OC1), it is recommended that you use the following URI format which uses Object Storage Dedicated Endpoints. See Object Storage Dedicated Endpoints, for further information.

```
https://namespace-string.objectstorage.region.oci.customer-oci.com/p/encrypted_string/n/
namespace-string/b/bucket/o/filename
```

> **Note:**
>
> OCI Object Store dedicated endpoint URLs are only supported in the commercial realms (OC1).

If your source files reside on Oracle Cloud Infrastructure Object Storage and are not in the commercial realm (OC1), you must use the following format:

```
https://objectstorage.region.oraclecloud.com.com/p/encrypted_string/n/namespace-string/b/
bucket/o/filename
```

For example, in the commercial realm (OC1) a sample pre-authenticated URI for the file `channels.txt` in the *bucketname* bucket in the `Phoenix` data center is:

```
https://namespace-string.objectstorage.us-phoenix-1.oci.customer-oci.com/p/2xN-
uDtWJNsiD910UCYGue/n/namespace-string/b/bucketname/o/channels.txt
```

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

You can use a pre-authenticated URL in any `DBMS_CLOUD` procedure that takes a URL to access files in Oracle Cloud Infrastructure object store, without the need to create a credential. You need to either specify the `credential_name` parameter as `NULL` or not supply a `credential_name` parameter.

For example:

```
BEGIN
  DBMS_CLOUD.COPY_DATA(
     table_name =>'CHANNELS',
     file_uri_list =>'https://objectstorage.us-phoenix-1.oraclecloud.com/p/
unique-pre-authenticated-string/n/namespace-string/b/bucketname/o/
channels.txt',
     format => json_object('delimiter' value ',') );
END;
/
```

> **Note:**
>
> A list of mixed URLs is valid. If the URL list contains both pre-authenticated URLs and URLs that require authentication, `DBMS_CLOUD` uses the specified `credential_name` to access the URLs that require authentication and for the pre-authenticated URLs the specified `credential_name` is ignored.

See Using Pre-Authenticated Requests for more information.

URI Format Using Public URL

If your source files reside on an Object Store that provides public URLs, you can use public URLs with `DBMS_CLOUD` procedures. Public means the Object Storage service supports anonymous, unauthenticated access to the Object Store files. See your Cloud Object Storage service for details on how to make an object public in a supported Object Store.

> **Note:**
>
> Carefully assess the business requirement for and the security ramifications of using public URLs. When you use public URLs, due to the file content not being authenticated, make sure this is appropriate for your use.

You can use a public URL in any `DBMS_CLOUD` procedure that takes a URL to access files in your object store, without the need to create a credential. You need to either specify the `credential_name` parameter as `NULL` or not supply a `credential_name` parameter.

For example the following uses `DBMS_CLOUD.COPY_DATA` without a `credential_name`:

```
BEGIN
  DBMS_CLOUD.COPY_DATA(
     table_name =>'CHANNELS',
     file_uri_list =>'https://objectstorage.us-ashburn-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/chan_v3.dat',
     format => json_object('delimiter' value ',') );
END;
/
```

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

> **Note:**
>
> A list of mixed URLs is valid. If the URL list contains both public URLs and URLs that require authentication, `DBMS_CLOUD` uses the specified `credential_name` to access the URLs that require authentication and for the public URLs the specified `credential_name` is ignored.

See Public Buckets for information on using Oracle Cloud Infrastructure public buckets.

Oracle Cloud Infrastructure Object Storage Classic URI Format

If your source files reside in Oracle Cloud Infrastructure Object Storage Classic, see the REST page for a description of the URI format for accessing your files: About REST URLs for Oracle Cloud Infrastructure Object Storage Classic Resources.

Amazon S3 URI Format

If your source files reside in Amazon S3, see the following for a description of the URI format for accessing your files: Accessing a bucket.

For example the following refers to the file `channels.txt` in the `adb` bucket in the `us-west-2` region.

```
https://s3-us-west-2.amazonaws.com/adb/channels.txt
```

You can use a presigned URL in any `DBMS_CLOUD` procedure that takes a URL to access files in Amazon S3 object store, without the need to create a credential. To use a presigned URL in

any DBMS_CLOUD procedure, either specify the credential_name parameter as NULL, or do not supply a credential_name parameter.

See Share an Object with Others for more information.

> **Note:**
>
> DBMS_CLOUD supports the standard Amazon S3 endpoint syntax to access your buckets. DBMS_CLOUD does not support Amazon S3 legacy endpoints. See Legacy Endpoints for more information.

Azure Blob Storage or Azure Data Lake Storage URI Format

If your source files reside in Azure Blob Storage or Azure Data Lake Storage, see the following for a description of the URI format for accessing your files: Resource URI Syntax.

For example the following refers to the file channels.txt in the adb container in the storage account adb_user:

```
https://adb_user.blob.core.windows.net/adb/channels.txt
```

> **Note:**
>
> You can use Shared Access Signatures (SAS) URL in any DBMS_CLOUD procedure that takes a URL to access files in Azure Blob Storage or Azure Data Lake Storage, without the need to create a credential. To use a Shared Access Signature (SAS) URL, either specify the credential_name parameter as NULL, or do not supply a credential_name parameter.
> See Grant Limited Access to Azure Storage Resources Using Shared Access Signatures (SAS) for more information.

Amazon S3 Compatible URI Format

DBMS_CLOUD supports object storage service implementations that support Amazon S3 compatible URLs, including the following services:

- Oracle Cloud Infrastructure Object Storage with Amazon S3 compatible URL
- Google Cloud Storage with Amazon S3 compatible URL
- Wasabi Hot Cloud Storage with Amazon S3 compatible URL

> **Note:**
>
> To use DBMS_CLOUD with an Amazon S3 compatible object store you need to provide valid credentials. See CREATE_CREDENTIAL Procedure for more information.

If your source files reside on a service that supports Amazon S3 compatible URIs, use the following URI format to access your files:

- **Oracle Cloud Infrastructure Object Storage S3 Compatible URL**

  If your source files reside on Oracle Cloud Infrastructure Object Storage in the commercial realm (OC1), it is recommended that you use the object URL and bucket URL formats listed below for commercial realm (OC1). See Object Storage Dedicated Endpoints, for further information.

  > **Note:**
  >
  > OCI Object Store dedicated endpoint URLs are only supported in the commercial realms (OC1).

  Object URL Formats

  – Supported only in the commercial realm (OC1):

  ```
  https://mynamespace.compat.objectstorage.region.oci.customer-oci.com/
  bucket_name/object_name
  ```

  – Supported in all zones:

  ```
  https://mynamespace.compat.objectstorage.region.oraclecloud.com/
  bucket_name/object_name
  ```

  Bucket URL Formats:

  – Supported only in the commercial realm (OC1):

  ```
  https://mynamespace.compat.objectstorage.region.oci.customer-oci.com/
  bucket_name
  ```

  – Supported in all zones:

  ```
  https://mynamespace.compat.objectstorage.region.oraclecloud.com/
  bucket_name
  ```

  See Amazon S3 Compatibility and Object Storage Service API for more information.

- **Google Cloud Storage S3 Compatible URL**

  Object URL Format:

  ```
  https://bucketname.storage.googleapis.com/object_name
  ```

  Bucket URL Format:

  ```
  https://bucketname.storage.googleapis.com/
  ```

  See Migrating from Amazon S3 to Cloud Storage and Request Endpoints for more information.

- **Wasabi S3 Compatible URL**

Object URL Format:

```
https://bucketname.s3.region.wasabisys.com/object_name
```

Bucket URL Format:

```
https://bucketname.s3.region.wasabisys.com/
```

See Wasabi S3 API Reference and Service URLs for Wasabi's Storage Regions for more information.

GitHub Raw URL Format

DBMS_CLOUD supports GitHub Raw URLs to access data from a GitHub Repository.

> **Note:**
>
> For DBMS_CLOUD access with GitHub Raw URLs, repository access is limited to read-only functionality. The DBMS_CLOUD APIs such as DBMS_CLOUD.PUT_OBJECT that write data are not supported with DBMS_CLOUD APIs on a GitHub Repository.
> As an alternative, use DBMS_CLOUD_REPO.PUT_FILE to upload data to a GitHub Repository.

Use GitHub Raw URLs with DBMS_CLOUD APIs to access source files that reside on a GitHub Repository. When you browse to a file on GitHub and click the **Raw** link, this shows the GitHub Raw URL. The raw.githubusercontent.com domain provides unprocessed versions of files stored in GitHub repositories.

For example, using DBMS_CLOUD.GET_OBJECT:

```
BEGIN
 DBMS_CLOUD.GET_OBJECT(
   credential_name => 'MY_CRED',
   object_uri      => 'https://raw.githubusercontent.com/myaccount/myrepo/
master/data-management-library/autonomous-database/adb-loading.csv',
   directory_name  => 'DATA_PUMP_DIR'
  );
END;
/
```

For example, using DBMS_CLOUD.CREATE_EXTERNAL_TABLE:

```
BEGIN
  DBMS_CLOUD.CREATE_EXTERNAL_TABLE(
    credential_name => 'MY_CRED',
    table_name      => 'EMPLOYEES_EXT',
    file_uri_list   => 'https://raw.githubusercontent.com/myaccount/myrepo/
master/data-management-library/autonomous-database/*.csv',
    column_list     => 'name varchar2(30), gender varchar2(30), salary
number',
    format          => JSON_OBJECT('type' value 'csv')
```

```
  );
END;
/
SELECT * FROM employees_ext;
```

DBMS_CLOUD procedures that take a URL to access a GitHub Repository do not require credentials with public visibility GitHub repositories. To use a public visibility URL you can specify the credential_name parameter as NULL or not supply a credential_name parameter. See Setting repository visibility for more information.

Additional Customer-Managed URI Formats

In addition to the pre-configured, recognized URIs with their fully-qualified domain names (FQDNs), DBMS_CLOUD cannot determine the proper authentication scheme for customer-managed endpoints URIs. In those cases, DBMS_CLOUD relies on the proper URI scheme to identify the authentication scheme for the customer-managed endpoint.

URI Scheme	Authentication Type	Access Method Description	URI Example
basic://	Basic authentication	Username and password stored in database credential object is used to authenticate the HTTP request	basic://api.github.com/users/myaccount
bearer://	Bearer token authentication	Bearer token stored in the password field in database credential object is used to specify the Authorization header for the HTTP request	bearer://api.sendgrid.com/v3/resource
oci://	OCI native	OCI signing key obtained from database credential object stored and used to sign requests using the OCI authentication protocol	oci://objectstorage.us-ashburn-1.oraclecloud.com
public://	No authentication	Public URLs	public://cms.data.gov/
s3://	Amazon Web Services S3-compatible	Access key and secret key obtained from the username/password field of database credential object, and S3-compatible authentication performed for the HTTP request.	s3://bucket.myprivatesite.com/file1.csv

Examples:

Customer-managed endpoint using S3-compatible authentication.

This example shows how for new URIs, customers can add the public or private host name pattern using DBMS_NETWORK_ACL_ADMIN package. The code block, executed by user ADMIN, enables HTTPS access for user SCOTT to endpoints in domain *.myprivatesite.com. It then shows how user SCOTT accesses the newly enabled endpoint. Note that credential MY_CRED for

user SCOTT must store the access key and secret key for S3-compatible authentication performed for the HTTP request indicated by the URI prefix.

```
BEGIN
   DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
         host => '*.myprivatesite.com',
         ace => xs$ace_type(privilege_list => xs$name_list('http'),
                            principal_name => 'SCOTT',
                            principal_type => xs_acl.ptype_db),
         private_target => TRUE   );
END;
/

BEGIN
   DBMS_CLOUD.get_object(
       credential_name => 'MY_CRED',
       object_uri      => 's3://bucket.myprivatesite.com/file1.csv',
       directory_name  => 'MY_DIR'  );
END;
/
```

Customer-managed endpoint with public access

This example shows how to register the SCOTT user to access public REST APIs. The ADMIN user creates a network ACL for the host to provide access to SCOTT user.

```
BEGIN
   DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
         host => 'data.cms.gov',
         ace => xs$ace_type(privilege_list => xs$name_list('http'),
                            principal_name => 'SCOTT',
                            principal_type => xs_acl.ptype_db)
   );
END;
/

SELECT DBMS_CLOUD.get_response_text(
          DBMS_CLOUD.send_request(
              uri     => 'public://data.cms.gov/provider-data/api/1/datastore/
imports/a',
              method  => DBMS_CLOUD.METHOD_GET,
              headers => JSON_OBJECT('Accept' VALUE 'application/json')
          )
       )
   FROM DUAL;
/
```

DBMS_CLOUD Package Format Options

The format argument in DBMS_CLOUD specifies the format of source files.

The two ways to specify the format argument are:

```
format => '{"format_option" : "format_value" }'
```

And:

```
format => json_object('format_option' value 'format_value'))
```

Examples:

```
format => json_object('type' VALUE 'CSV')
```

To specify multiple format options, separate the values with a ",".

For example:

```
format => json_object('ignoremissingcolumns' value 'true', 'removequotes' value 'true',
                          'dateformat' value 'YYYY-MM-DD-HH24-MI-SS', 'blankasnull'
value 'true', 'logretention' value 7)
```

> **Note:**
>
> For Avro, ORC, or Parquet format options, see DBMS_CLOUD Package Format Options for Avro, ORC, or Parquet.

As noted in the **Format Option** column, a limited set of format options are valid with `DBMS_CLOUD.COPY_COLLECTION` or with `DBMS_CLOUD.COPY_DATA` when the `format type` is JSON.

Format Option	Description	Syntax
`access_protocol`	Specifies the type of Apache Iceberg table, such as AWS or OCI Object Storage, and what information is used to create the external table, for example information from a data catalog or from a direct metadata URI.	See CREATE_EXTERNAL_TABLE Procedure for Apache Iceberg, for details on the `access_protocol` syntax.
`blankasnull`	When set to `true`, loads fields consisting of spaces as null.	`blankasnull : true` Default value: `False`
`characterset` Valid with `format` JSON and `COPY_DATA`	Specifies the characterset of source files	`characterset`: *string* Default value: Database characterset
`columnpath` Only use with `format` JSON and `COPY_DATA`	Array of JSON path expressions that correspond to the fields that need to be extracted from the JSON records. Each of the JSON path expressions in the array should follow the rules described in SQL/JSON Path Expressions. Only use with format JSON and `DBMS_CLOUD.COPY_DATA`.	JSON Array of json path expressions expressed in string format. For example: 'columnpath' value `'["$.WEATHER_STATION_ID", "$.WEATHER_STATION_NAME"]'`
`compression` Option valid with JSON data	Specifies the compression type of the source file. ZIP archiving format is not supported. Specifying the value `auto` checks for the compression types: `gzip, zlib, zstd, bzip2`.	`compression`: `auto|gzip|zlib|zstd|bzip2` Default value: Null value meaning no compression.
`conversionerrors`	If a row is rejected because of data type conversion errors, the related columns are stored as null or the row is rejected.	`conversionerrors : reject_record | store_null` Default value: `reject_record`

Format Option	Description	Syntax

`dateformat`	Specifies the date format in the source file. The format option `AUTO` searches for the following formats:
	
	`J`
	`MM-DD-YYYYBC`
	`MM-DD-YYYY`
	`YYYYMMDD HHMISS`
	`YYMMDD HHMISS`
	`YYYY.DDD`
	`YYYY-MM-DD`
		`dateformat` : *string*
		Default value: Database date format

`delimiter`	Specifies the field delimiter. To use a special character as the delimiter, specify the HEX value of the ASCII code of the character. For example, the following specifies the TAB character as the delimiter: `format => json_object('delimiter' value 'X''9''')`	`delimiter` : *character* Default value `|` (pipe character)

Format Option	Description	Syntax
detectfieldorder	Specifies that the fields in the external data files are in a different order than the columns in the table. Detect the order of fields using the first row of each external data file and map it to the columns of the table. The field names in external data files are compared in case insensitive manner with the names of the table columns. This format option is applicable for the following procedures: • DBMS_CLOUD.COPY_DATA • DBMS_CLOUD.CREATE_EXTERNAL_TABLE • DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE • DBMS_CLOUD.CREATE_HYBRID_PART_TABLE Restrictions for detectfieldorder: • Field names in the data file must appear in the first record line and must not contain any white spaces between the field names. • The field delimiter in the field names record must be the same as the field delimiter for the data in the file. • Quoted field names are not supported. The field names in data files are compared, in case insensitive manner, with the names of the external table columns. • Embedded field delimiters are not allowed in the field names. • The number of columns in the table must match the number of fields in the data files. • This format option is not applicable for Bigdata or Oracle Data Pump formats, as those formats have precise column metadata information in the binary file format. The text formats, CSV, JSON, Parquet, or XML can benefit from this automatic field order detection when the first line contains the field names. See FIELD NAMES and the description for ALL FILES for more information.	detectfieldorder: true Default value: false
enablelogs	The format option enablelogs is used with the following DBMS_CLOUD procedures: • COPY_DATA • COPY_COLLECTION • EXPORT_DATA enablelogs specifies a boolean value, when set to TRUE, logs are generated. When set to FALSE, logs are not generated. For example: format => JSON_OBJECT('enablelogs' value FALSE)	enablelogs: false Default value: true

Format Option	Description	Syntax
`encryption`	The format option `encryption` specifies the encryption and decryption options to export and import data to and from the Object Store. Use `encryption` to specify the following parameters to encrypt and decrypt: • `user_defined_function`: Specifies a fully qualified user defined function to decrypt or encrypt the specified BLOB (binary large object). It returns a decrypted or encrypted BLOB. `user_defined_function` is mutually exclusive with other parameters for `encryption`. For example, `ADMIN.DECRYPTION_CALLBACK`. • `type`: Specifies the `DBMS_CRYPTO` encryption algorithm to decrypt or encrypt. `type` accepts values in the *Block Cipher Algorithms* + *Block Cipher Chaining Modifiers* + *Block Cipher Padding Modifiers* format. Supported Block Cipher Algorithms are: – `DBMS_CRYPTO.ENCRYPT_AES256` Supported Block Cipher Chaining Modifiers are: – `DBMS_CRYPTO.CHAIN_CBC` – `DBMS_CRYPTO.CHAIN_CFB` – `DBMS_CRYPTO.CHAIN_ECB` – `DBMS_CRYPTO.CHAIN_OFB` Supported Block Cipher Padding Modifiers are: – `DBMS_CRYPTO.PAD_PKCS5` – `DBMS_CRYPTO.PAD_NONE` – `DBMS_CRYPTO.PAD_ZERO` – `DBMS_CRYPTO.PAD_ORCL` • `credential_name`: Specifies the credential used to store the encryption key. The Block Cipher Chaining Modifiers and Block Cipher Padding Modifiers values defaults to `DBMS_CRYPTO.CHAIN_CBC` and `DBMS_CRYPTO.PAD_PKCS5`, if you do not specify values for these parameters. The format option `encryption` is used with the following `DBMS_CLOUD` procedures: • Used to pass parameters to decrypt for these procedures: – `DBMS_CLOUD.COPY_DATA` – `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` – `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` – `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE` For `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE` this option is only applicable to the Object Storage files. – `DBMS_CLOUD.COPY_COLLECTION`	`encryption:`*value* Where *value* is a JSON string that provides additional parameters for encryption: `type`: *value* Specifies the encryption type. `credential_name`: *value* Specifies the credential used to store the encryption key. `user_defined_function`: *value* Specifies a fully qualified user-defined function to decrypt or encrypt the specified BLOB (binary large object).

Format Option	Description	Syntax
	• Used to pass parameters to encrypt for these procedure: – DBMS_CLOUD.EXPORT_DATA For example: `format => JSON_OBJECT('encryption' value json_object ('type' value DBMS_CRYPTO.ENCRYPT_AES256 + DBMS_CRYPTO.CHAIN_CBC + DBMS_CRYPTO.PAD_PKCS5, 'credential_name' value 'ENCRYPTION_CRED'))`	
`endquote`	Data can be enclosed between two delimiters, specified with `quote` and `endquote`. The `quote` and `endquote` characters are removed during loading when specified. For example: `format => JSON_OBJECT('quote' value '(', 'endquote' value ')')`	`endquote:character` Default value: Null, meaning no `endquote`.
`escape`	The character "\" is used as the escape character when specified.	`escape : true` Default value: `false`
`ignoreblanklines` Option valid with JSON data	Blank lines are ignored when set to true.	`ignoreblanklines : true` Default value: `False`
`ignoremissingcolumns`	If there are more columns in the `field_list` than there are in the source files, the extra columns are stored as null.	`ignoremissingcolumns : true` Default value `False`

Format Option	Description	Syntax

`implicit_partition_columns`	Enable implicit partitioning and specify the partition column names by using the `implicit_partition_columns` format option with `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`.
	Implicit partitioning is enabled in the following ways:
	• Use `implicit_partition_columns` to provide a list of partition columns and specify the `implicit_partition_type`. For example:
	`format => '{"implicit_partition_type":"hive", "implicit_partition_columns": ["country","year","month"]}'`
	• Use `implicit_partition_columns` to provide a list of partition columns without providing the partition type. The partition type is automatically detected as hive or non-hive. For example:
	`format => '{"implicit_partition_columns": ["country","year","month"]}'`
	• Use `implicit_partition_type` to provide the type of partition columns without providing a list of partition columns. The automatic discovery of the partition keys in HIVE-style partitioned data is triggered to determine column names. For example:
	`format => '{"partition_type":"hive"}'`
	See `implicit_partition_type` for the option description and syntax.
		`implicit_partition_columns` : *array of strings*
		Default value: If `implicit_partition_type` is specified, the column names are derived through automatic discovery of the partition keys in HIVE-style partitioned data. Otherwise, the default is null and implicit partitioning is not enabled.

`implicit_partition_type`	Enable implicit partitioning and specify the data types of partition columns by using the `implicit_partition_type` format option with `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`. See `implicit_partition_columns` for further information on enabling implicit partitioning and examples.	`implicit_partition_type` : hive Default value: If `implicit_partition_columns` is specified, the type is automatically detected as hive or non-hive. Otherwise, the default is null and implicit partitioning is not enabled.
`jsonpath` Only use with `COPY_COLLECTION`	JSON path to identify the document to load. This option is valid only for JSON collection data with `DBMS_CLOUD.COPY_COLLECTION`.	`jsonpath`: *string* Default value: Null
`keyassignment` Only use with `COPY_COLLECTION`	Specifies whether a new collection is created as a mongo-compatible collection or as a SODA collection. When the value is set to `embedded_oid`, a new collection is created as a mongo-compatible collection. By default this parameter is not set, meaning a new collection is created as a SODA collection.	`keyassignment`: `embedded_oid` Default: `keyassignment` is not set

Format Option	Description	Syntax
`keypath` Only use with `COPY_COLLECTION`	Specifies an attribute in the data to be loaded as the `'_id'` value. If `keypath` is specified then you must also specify the `keyassignment` value as `embedded_oid`. Set the value to a path, for example, `'$.mykey'`, to pick the value of the path as `'_id'` value. This parameter is optional and is only valid for loading into mongo-compatible collections. If not specified, Oracle generates a 12-byte unique system ID and populates that as the `'_id'` attribute, if an `'_id'` attribute is not already present in the data being loaded.	`keypath`: *`string`* Default: `keypath` is not set. When `keypath` is set, the default *string* value is `NULL`.
`language`	Specifies a language name (for example, FRENCH), from which locale-sensitive information can be derived.	`language`: *string* Default value: Null See Locale Data in *Oracle Database Globalization Support Guide* for a listing of Oracle-supported languages.

`logdir`	Specifies a string value that determines the directory object name where the `logfile_table` or `badfile_table` files are saved.
	By default, the `logdir` is not case-sensitive, but the case is reserved when the specified value is enclosed in double-quotes.
	For example:
	`format => JSON_OBJECT ('logdir' value 'test_log')`
	The `logdir` format option specified in the above example saves the `logfile_table` or `badfile_table` files in the `TEST_LOG` directory object.
	`format => JSON_OBJECT ('logdir' value '"test_log"')`
	The `logdir` format option specified in the above example saves the `logfile_table` or `badfile_table` files in the `test_log` directory object.
		`logdir`: *string*
		Default value: `DATA_PUMP_DIR`

`logprefix`	Specifies a string value that determines the prefix for the `logfile_table` and `badfile_table` files. The log table name format is: *logprefix*`$`*operation_id* By default, the logprefix is in upper case, but the case is reserved when the specified value is enclosed in double-quotes. For example: `format => JSON_OBJECT ('logprefix' value 'TEST')` Log files then use the `TEST` prefix, such as: `TEST$2_LOG` and `TEST$2_BAD`.	`logprefix`: *string* Default value: `COPY`

Format Option	Description	Syntax
`logretention`	Specifies a positive integer duration, in days, for which the `logfile_table` and `badfile_table` files are retained. Valid values: `0` to `99999` For example: `format => JSON_OBJECT ('logretention' value 7)`	`logretention`: *number* Default value: `2`
`maxdocsize` This option is valid only with JSON data	Maximum size of JSON documents.	`maxdocsize`: *number* Default value: 1 Megabyte Maximum allowed value: 2 Gigabytes
`numericcharacters`	Specifies the characters to use as the group separator and decimal character. *decimal_character*: The decimal separates the integer portion of a number from the decimal portion. *group_separator*: The group separator separates integer groups (that is, thousands, millions, billions, and so on).	`numericcharacters:` '*decimal_character group_separator*' Default value: `"."` See NLS_NUMERIC_CHARACTERS in *Oracle Database Globalization Support Guide* for more information.
`numberformat`	Specifies the number format model. Number format models cause the number to be rounded to the specified number of significant digits. A number format model is composed of one or more number format elements. This is used in combination with `numericcharacters`.	`numberformat:` *number_format_model* Default value: is derived from the setting of the `NLS_TERRITORY` parameter See Number Format Models in *SQL Language Reference* for more information.

Format Option	Description	Syntax
partition_columns	The format option partition_columns is used with DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE to specify the column names and data types of partition columns when the partition columns are derived from the file path, depending on the type of data file, structured or unstructured: • When the DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE includes the column_list parameter and the data file is unstructured, such as with CSV text files, partition_columns does not include the data type. For example, use a format such as the following for this type of partition_columns specification: '"partition_columns": ["state","zipcode"]' The data type is not required because it is specified in the DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE column_list parameter. • When the DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE does not include the column_list parameter and the data files are structured, such as Avro, ORC, or Parquet files, the partition_columns option includes the data type. For example, the following shows a partition_columns specification: '"partition_columns":[{"name":"country", "type":"varchar2(10)"}, {"name":"year", "type":"number"}, {"name":"month", "type":"varchar2(10)"}]' If the data files are unstructured and the type sub-clause is specified with partition_columns, the type sub-clause is ignored. For object names that are not based on hive format, the order of the partition_columns specified columns must match the order as they appear in the object name in the file_uri_list.	
quote	Specifies the quote character for the fields, the quote characters are removed during loading when specified.	quote: *character* Default value: Null meaning no quote

Format Option	Description	Syntax

`recorddelimiter`
Option valid with JSON data
	Specifies the record delimiter.
	By default, `DBMS_CLOUD` tries to automatically find the correct newline character as the delimiter. It first searches the file for the Windows newline character "`\r\n`". If it finds the Windows newline character, this is used as the record delimiter for all files in the procedure. If a Windows newline character is not found, it searches for the UNIX/Linux newline character "`\n`" and if it finds one it uses "`\n`" as the record delimiter for all files in the procedure.
	Specify this argument explicitly if you want to override the default behavior, for example:
	`format => json_object('recorddelimiter' VALUE '''\r\n''')`
	To indicate that there is no record delimiter you can specify a `recorddelimiter` that does not occur in the input file. For example, to indicate that there is no delimiter, specify the control character 0x01 (SOH) as a value for the `recorddelimiter` and set the `recorddelimiter` value to "`0x''01''`" (this character does not occur in JSON text). For example:
	`format => '{"recorddelimiter" : "0x''01''"}'`
	The `recorddelimiter` is set once per procedure call. If you are using the default value, `detected newline`, then all files use the same record delimiter, if one is detected.
		`recorddelimiter`: *character*
		Default value: detected newline

Format Option	Description	Syntax
`regexuri`	The format option `regexuri` is used with the following DBMS_CLOUD procedures: • `COPY_COLLECTION` • `COPY_DATA` • `CREATE_EXTERNAL_TABLE` • `CREATE_EXTERNAL_PART_TABLE` • `CREATE_HYBRID_PART_TABLE` When the value of `regexuri` is set to `TRUE`, you can use wildcards as well as regular expressions in the file names in Cloud source file URIs. The characters "*" and "?" are considered wildcard characters when the `regexuri` parameter is set to `FALSE`. When the `regexuri` parameter is set to `TRUE` the characters "*" and "?" are part of the specified regular expression pattern. Regular expression patterns are only supported for the file name or subfolder path in your URIs and the pattern matching is identical to that performed by the `REGEXP_LIKE` function. Regular expression patterns are not supported for directory names. For external tables, this option is only supported with the tables that are created on a file in the Object Storage. For example: `format => JSON_OBJECT('regexuri' value TRUE)` See REGEXP_LIKE Condition for more information on `REGEXP_LIKE` condition.	`regexuri`: `True` Default value : `False`
`rejectlimit`	The operation will error out after specified number of rows are rejected.	`rejectlimit`: *number* Default value: `0`
`removequotes`	Removes any quotes that are around any field in the source file.	`removequotes`: `true` Default value: `False`
`skipheaders`	Specifies how many rows should be skipped from the start of the file.	`skipheaders`: *number* Default value: 0 if not specified, 1 if specified without a value
`territory`	Specifies a territory name to further determine input data characteristics.	`territory`: *string* Default value: Null See Locale Data in *Oracle Database Globalization Support Guide* for a listing of Oracle-supported territories.

`timestampformat`	Specifies the timestamp format in the source file. The format option `AUTO` searches for the following formats:
	`YYYY-MM-DD HH:MI:SS.FF`
	`YYYY-MM-DD HH:MI:SS.FF3`
	`YYYY-MM-DD HH24:MI:SS.FF3`
	`MM/DD/YYYY HH:MI:SS.FF3`
		`timestampformat` : *string*
		Default value: Database timestamp format
		The string can contain wildcard characters such as "$".

Format Option	Description	Syntax

`timestampltzformat`	Specifies the timestamp with local timezone format in the source file. The format option `AUTO` searches for the following formats:
	`DD Mon YYYY HH:MI:SS.FF TZR`
	`MM/DD/YYYY HH:MI:SS.FF TZR`
	`YYYY-MM-DD HH:MI:SS+/-TZR`
	`YYYY-MM-DD HH:MI:SS.FF3`
	`DD.MM.YYYY HH:MI:SS TZR`
		`timestampltzformat` : *string*
		Default value: Database timestamp with local timezone format

`timestamptzformat`	Specifies the timestamp with timezone format in the source file. The format option `AUTO` searches for the following formats:
	`DD Mon YYYY HH:MI:SS.FF TZR`
	`MM/DD/YYYY HH:MI:SS.FF TZR`
	`YYYY-MM-DD HH:MI:SS+/-TZR`
	`YYYY-MM-DD HH:MI:SS.FF3`
	`DD.MM.YYYY HH:MI:SS TZR`
		`timestamptzformat`: *string*
		Default value: Database timestamp with timezone format

`trimspaces`	Specifies how the leading and trailing spaces of the fields are trimmed. See the description of trim_spec.	`trimspaces`: `rtrim| ltrim| notrim| lrtrim| ldrtrim` Default value: `notrim`
`truncatecol`	If the data in the file is too long for a field, then this option will truncate the value of the field rather than reject the row.	`truncatecol`:`true` Default value: `False`

Format Option	Description	Syntax
`type`	Specifies the source file type. See the description of `CSV` in field_definitions Clause If the `type` is `datapump`, then the only other valid format option is `rejectlimit`. If the `type` is `datapump`, then the only Object Stores supported are Oracle Cloud Infrastructure Object Storage and Oracle Cloud Infrastructure Object Storage Classic. See DBMS_CLOUD Package Format Options for Avro, ORC, or Parquet for `type` values `avro`, `orc`, or `parquet`. For JSON data with `DBMS_CLOUD.COPY_COLLECTION` `type` has two valid values: `json` (default) and `ejson`. For `DBMS_CLOUD.COPY_COLLECTION` these values both specify that the input is JSON data. The value `ejson` causes extended objects in the textual JSON input data to be translated to scalar JSON values in the native binary JSON collection. The value `json` does not perform this transformation and all objects in the input data are converted to binary JSON format. For JSON data with `DBMS_CLOUD.COPY_DATA` `type` has one valid value: `json`. This value specifies that the input is JSON data. • See COPY_COLLECTION Procedure for information on loading JSON objects. • See Textual JSON Objects That Represent Extended Scalar Values for information about extended JSON objects, which are interpreted when parameter `type` has value `ejson`.	`type: csv|csv with embedded|csv without embedded |avro||datapump| orc|parquet` **Note:** Not all `DBMS_CLOUD` procedures support all of these types. `csv` is the same as `csv without embedded`. Default value: Null For JSON data there are two valid `type` values for use with `DBMS_CLOUD.COPY_COLLECTION`: `json|ejson` In this case the default value is `json`. For JSON data with `DBMS_CLOUD.COPY_DATA`, only `json` is valid.
`unpackarrays` Only use with `COPY_COLLECTION`	When set to `true`, if a loaded document is an array, then the contents of the array are loaded as documents rather than the array itself. This only applies to the top-level array. When set to `true`, the entire array is inserted as a single document. This option is valid only for JSON collection data with `DBMS_CLOUD.COPY_COLLECTION`.	`unpackarrays: true` Default value: `False`

DBMS_CLOUD Package Format Options for EXPORT_DATA

Describes the valid format parameter options for `DBMS_CLOUD.EXPORT_DATA` with text file formats, CSV, JSON, Parquet, or XML, and for Oracle Data Pump.

These are the valid `format` parameters for use with `DBMS_CLOUD.EXPORT_DATA`. You specify text file output when you use the `format type` option and the value is one of: `csv`, `json`, `parquet`, or `xml`. This also shows the `format` options when the `format type` is `datapump`.

The two ways to specify the format argument are:

```
format => '{"format_option" : "format_value" }'
```

And:

```
format => json_object('format_option' value 'format_value'))
```

Examples:

```
format => json_object('type' VALUE 'json')
```

To specify multiple format options, separate the values with a ",".

For example:

```
format => json_object('compression' value 'gzip', 'type' value 'json')
```

This table covers the format options for `DBMS_CLOUD.EXPORT_DATA` when the `format` parameter `type` option is one of: CSV, JSON, Parquet, or XML. For other procedures and other output types, see DBMS_CLOUD Package Format Options for the list of format options.

Format Option	Description	Syntax

`compression`	Specifies the compression type of the source file.
	Note: ZIP archiving format is not supported.
	When the `format type` is `csv`, `json`, or `xml`, the default compression is Null, meaning no compression.
	When the `format type` is `parquet`, the default compression is `snappy`.
	When the `format type` is `datapump` you can specify supported Oracle Data Pump access parameters:
	• `compression`: The valid values are: `BASIC`, `LOW`, `MEDIUM`, and `HIGH`.
	• `version`: The valid values are: `COMPATIBLE`, `LATEST`, and a specified *`version_number`*.
		When the `type` is:`csv` | `json` | `xml`
		`compression:gzip | zlib | zstd | bzip2`
		Default value: Null value meaning no compression.
		When the `type` is `parquet`
		`compression: gzip` | `snappy`
		Default value: `snappy`
		When the type is `datapump`
		`compression: BASIC|LOW|MEDIUM|HIGH`

`delimiter`	Specifies a custom field delimiter. `format => json_object('delimiter' value '|')` The delimiter value cannot be an ASCII code or an escape character. **Note:** This option only applies with `csv type`.	`delimiter` : *character* Default value , (comma)
`endquote`	Specifies that fields can be enclosed between two delimiters, with `quote` and `endquote`. If `endquote` is not specified, then the `quote` character will be used by default as the `endquote` character. For example: `format => JSON_OBJECT('quote' value '(', 'endquote' value ')')` **Note:** This option only applies with `csv type`.	`endquote`:*character* Default value: Null, meaning no `endquote`.

Format Option	Description	Syntax
`escape`	Specifies the occurrence of quote character in the field value using `"\"` character.	`escape : true` Default value: `false`

> **Note:**
>
> This option only applies with `csv type`.

Format Option	Description	Syntax
`encryption`	The format option `encryption` specifies the encryption and decryption options to export and import data to and from the Object Store. Use `encryption` to specify the following parameters to encrypt and decrypt: • `user_defined_function`: Specifies a fully qualified user-defined function to decrypt or encrypt the specified BLOB (binary large object). It returns a decrypted or encrypted BLOB. This parameter is mutually exclusive with other parameters for `encryption`. For example, `ADMIN.DECRYPTION_CALLBACK`. • `type`: Specifies the built-in encryption algorithm to decrypt or encrypt. `user_defined_function` and `type` are mutually exclusive. `type` accepts values in the *Block Cipher Algorithms* + *Block Cipher Chaining Modifiers* + *Block Cipher Padding Modifiers* format. Supported Block Cipher Algorithms are: – `DBMS_CRYPTO.ENCRYPT_AES256` Supported Block Cipher Chaining Modifiers are: – `DBMS_CRYPTO.CHAIN_CBC` – `DBMS_CRYPTO.CHAIN_CFB` – `DBMS_CRYPTO.CHAIN_ECB` – `DBMS_CRYPTO.CHAIN_OFB` Supported Block Cipher Padding Modifiers are: – `DBMS_CRYPTO.PAD_PKCS5` – `DBMS_CRYPTO.PAD_NONE` – `DBMS_CRYPTO.PAD_ZERO` – `DBMS_CRYPTO.PAD_ORCL` • `credential_name`: Specifies the credential used to store the encryption key. The Block Cipher Chaining Modifiers and Block Cipher Padding Modifiers values defaults to `DBMS_CRYPTO.CHAIN_CBC` and `DBMS_CRYPTO.PAD_PKCS5`, if you do not specify values for these parameters. The format option `encryption` is used with the following `DBMS_CLOUD` procedures: • Used to pass parameters to decrypt for these procedures: – `DBMS_CLOUD.COPY_DATA` – `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` – `DBMS_CLOUD.CREATE_HYBRID_TABLE` – `DBMS_CLOUD.COPY_COLLECTION` • Used to pass parameters to encrypt for these procedure: – `DBMS_CLOUD.EXPORT_DATA` For example: `format => JSON_OBJECT('encryption' value json_object ('type' value DBMS_CRYPTO.ENCRYPT_AES256 +`	`encryption:`*value* Where *value* is a JSON string that provides additional parameters for encryption: `type`: *value* Specifies the encryption type. `credential_name`: *value* Specifies the credential used to store the encryption key. `user_defined_function`: *value* Specifies a fully qualified user-defined function to decrypt or encrypt the specified BLOB (binary large object).

Format Option	Description	Syntax
	`DBMS_CRYPTO.CHAIN_CBC + DBMS_CRYPTO.PAD_PKCS5, 'credential_name' value 'ENCRYPTION_CRED'))`	
`header`	Writes column names as the first line in output files of `csv` type. The `header` option can accept a `boolean` or a `string` value. The valid values are: • `false`: Skips the header row. • `true`: Includes the header row. The column names are based on the `SELECT` statement in the `query` parameter. You must specify column aliases in the `SELECT` statement when using virtual columns or expressions. • *`String to define custom header names`*: Enables you to define header rows with custom names. The number of columns and delimiters in the string value must match the number of columns and delimiters in the `SELECT` statement. The default delimiter is comma (,). For example: `format => JSON_OBJECT('type' value 'csv', 'delimiter' value '|', 'compression' value 'gzip', 'header' value true)` **Note:** This option only applies with `csv type`.	`header: true| false| String to define custom header names` Default value: `false`
`fileextension`	Custom file extension to override the default choice for the format type. This applies to text formats with `DBMS_CLOUD.EXPORT_DATA`: CSV, JSON, Parquet, or XML. If the specified string does not start with period (dot), then a dot is automatically inserted before the file extension in the final file name. If no file extension is desired, use the value: `fileextension ='none'`	Valid values: Any file extension. Default value: Depends on the `format type` option: • CSV format: `.csv` • JSON format: `.json` • PARQUET format: `.parquet` • XML format: `.xml`
`maxfilesize`	Number in bytes for maximum size of output generated. This applies to text based formats for exporting data with `DBMS_CLOUD.EXPORT_DATA` when the `format type` option is set to,`csv`, `json`, or `xml`. Note: This option is not valid when the `format type` option is `parquet`.	Minimum value: `10485760` (10 MB) Maximum value: 1 GB Default value: `10485760` (10 MB)

Format Option	Description	Syntax
`quote`	In CSV format, fields can be enclosed between two delimiters. Specify the delimiters with `quote` and `endquote`. If `endquote` is not specified, then the `quote` character will be used by default as the `endquote` character. **Note:** This option only applies with `csv type`.	`quote`: *character* Default value: Null meaning do not enclose fields with quotes.
`trimspaces`	Specifies how the leading and trailing spaces of the fields are trimmed for CSV format. Trim spaces is applied before quoting the field, if the `quote` parameter is specified. See the description of trim_spec. **Note:** This option only applies with `csv type`.	`trimspaces: rtrim| ltrim| notrim| lrtrim| ldrtrim` Default value: `notrim`
`type`	Specifies the output file type. `csv`: specifies Character Separated Values (CSV) format that allows you to export query results as a set of column values separated by any custom character. `json`: specifies to export the query results as JSON files. `parquet`: specifies to export the query results as Parquet files. `xml`: specifies to export query results as rows of valid XML documents. Each row is encapsulated in a root XML tag of `<RECORD> </RECORD>`. The query result is automatically transformed into XML format using XMLFOREST SQL function. Use Column Aliases to customize the XML tag names for columns. When the `format type` is `datapump` you can specify supported Oracle Data Pump access parameters: • `compression`: The valid values are: `BASIC`, `LOW`, `MEDIUM`, and `HIGH`. • `version`: The valid values are: `COMPATIBLE`, `LATEST`, and a specified *`version_number`*. See access_parameters Clause for more information.	`type: csv|datapump|json| parquet|xml`

DBMS_CLOUD Avro, ORC, and Parquet Support

This section covers the `DBMS_CLOUD` Avro, ORC, and Parquet support provided with Autonomous Database.

- DBMS_CLOUD Package Format Options for Avro, ORC, or Parquet
  The format argument in DBMS_CLOUD specifies the format of source files.
- DBMS_CLOUD Package Avro to Oracle Data Type Mapping
  Describes the mapping of Avro data types to Oracle data types.
- DBMS_CLOUD Package ORC to Oracle Data Type Mapping
  Describes the mapping of ORC data types to Oracle data types.
- DBMS_CLOUD Package Parquet to Oracle Data Type Mapping
  Describes the mapping of Parquet data types to Oracle data types.
- DBMS_CLOUD Package Oracle Data Type to Parquet Mapping
  Describes the mapping of Oracle data types to Parquet data types.
- DBMS_CLOUD Package Avro, ORC, and Parquet Complex Types
  Describes the mapping of Avro, ORC, and Parquet complex data types to Oracle data types.
- DBMS_CLOUD Package Avro, ORC, and Parquet to Oracle Column Name Mapping
  Describes rules for how Avro, ORC, and Parquet column names are converted to Oracle column names.

DBMS_CLOUD Package Format Options for Avro, ORC, or Parquet

The format argument in DBMS_CLOUD specifies the format of source files.

The two ways to specify the format argument are:

```
format => '{"format_option" : "format_value" }'
```

And:

```
format => json_object('format_option' value 'format_value'))
```

Examples:

```
format => json_object('type' VALUE 'CSV')
```

To specify multiple format options, separate the values with a ",".

For example:

```
format => json_object('ignoremissingcolumns' value 'true', 'removequotes' value 'true',
'dateformat' value 'YYYY-MM-DD-HH24-MI-SS', 'blankasnull' value 'true')
```

Format Option	Description	Syntax
maxvarchar	Controls the default maximum varchar2() length for character columns when creating external tables on Avro, ORC, or Parquet format files. The options are: • standard: Specifies that the length limit for varchar2() columns is 4000 bytes. This is equivalent to auto on systems with MAX_STRING_SIZE = standard. On systems with MAX_STRING_SIZE = extended, when you create an external table on Avro, ORC, or Parquet and you specify maxvarchar as standard, this shortens character columns to varchar2(4000). • extended: Specifies that the length limit for varchar2() columns is 32767 bytes. This is equivalent to auto on systems with MAX_STRING_SIZE = extended. This value raises an error on systems with MAX_STRING_SIZE = standard. • auto: The system automatically picks the varchar2() maximum byte size, based on the MAX_STRING_SIZE value. The value auto is the default value for maxvarchar.	maxvarchar: standard | extended | auto Default value:auto
regexuri	When the value of regexuri is set to TRUE, you can use wildcards as well as regular expressions in the file names in Cloud source file URIs. The characters "*" and "?" are considered wildcard characters when the regexuri parameter is set to FALSE. When the regexuri parameter is set to TRUE the characters "*" and "?" are part of the specified regular expression pattern. Regular expression patterns are only supported for the file name or subfolder path in your URIs and the pattern matching is identical to that performed by the REGEXP_LIKE function. Regular expression patterns are not supported for directory names. For external tables, this option is only supported with the tables that are created on a file in the Object Storage. For example: format => JSON_OBJECT('regexuri' value TRUE) See REGEXP_LIKE Condition for more information on REGEXP_LIKE condition.	regexuri: True Default value:False
type	Specifies the file type.	type : avro | orc | parquet

Format Option	Description	Syntax
schema	When schema is set to first or all, the external table columns and data types are automatically derived from the Avro, ORC, or Parquet file metadata. The column names will match those found in Avro, ORC, or Parquet. The data types are converted from Avro, ORC, or Parquet data types to Oracle data types. All columns are added to the table. The value first specifies to use the metadata from the first Avro, ORC, or Parquet file in the file_uri_list to auto generate the columns and their data types. Use first if all of the files have the same schema. The value all specifies to use the metadata from all Avro, ORC, or Parquet files in the file_uri_list to auto generate the columns and their data types. Use all (slower) if the files may have different schemas. **Default**: If column_list is specified, then the schema value, if specified is ignored. If column_list is not specified then the schema default value is first. **Note**: For Avro, ORC, or Parquet format files the schema format option is not available and the column_list parameter must be specified for partitioned external tables using the DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE procedure.	schema : first | all

DBMS_CLOUD Package Avro to Oracle Data Type Mapping

Describes the mapping of Avro data types to Oracle data types.

> **Note:**
>
> Complex types, such as maps, arrays, and structs are supported starting with Oracle Database 19c. See DBMS_CLOUD Package Avro, ORC, and Parquet Complex Types for information on using Avro complex types.

Avro Type	Oracle Type
INT	NUMBER(10)
LONG	NUMBER(19)
BOOL	NUMBER(1)
UTF8 BYTE_ARRAY	RAW(2000)
FLT	BINARY_FLOAT
DBL	BINARY_DOUBLE
DECIMAL(p)	NUMBER(p)
DECIMAL(p,s)	NUMBER(p,s)
DATE	DATE
STRING	VARCHAR2
TIME_MILLIS	VARCHAR2(20 BYTE)
TIME_MICROS	VARCHAR2(20 BYTE)
TIMESTAMP_MILLIS	TIMESTAMP(3)

Avro Type	Oracle Type
TIMESTAMP_MICROS	TIMESTAMP(6)
ENUM	VARCHAR2(*n*) Where: "*n*" is the actual maximum length of the AVRO ENUM's possible values
DURATION	RAW(2000)
FIXED	RAW(2000)
NULL	VARCHAR2(1) BYTE

See DBMS_CLOUD Package Avro, ORC, and Parquet Complex Types for information on using Avro complex types.

DBMS_CLOUD Package ORC to Oracle Data Type Mapping

Describes the mapping of ORC data types to Oracle data types.

See DBMS_CLOUD Package Avro, ORC, and Parquet Complex Types for information on using ORC complex types.

ORC Type	Oracle Type	More Information
array	VARCHAR2(*n*) JSON format	DBMS_CLOUD Package Avro, ORC, and Parquet Complex Types
bigint (64 bit)	NUMBER(19)	
binary	BLOB	
boolean (1 bit)	NUMBER(1)	
char	CHAR(*n*)	
date	DATE	
double	BINARY_DOUBLE	
float	BINARY_FLOAT	
int (32 bit)	NUMBER(10)	
list	VARCHAR2(*n*) JSON format	DBMS_CLOUD Package Avro, ORC, and Parquet Complex Types
map	VARCHAR2(*n*) JSON format	DBMS_CLOUD Package Avro, ORC, and Parquet Complex Types
smallint (16 bit)	NUMBER(5)	
string	VARCHAR2(4000) or VARCHAR2(32767)	Depending on the value of the format parameter option `maxvarchar` and on the `MAX_STRING_SIZE` value. See DBMS_CLOUD Package Format Options for Avro, ORC, or Parquet for more information.
struct	VARCHAR2(*n*) JSON format	DBMS_CLOUD Package Avro, ORC, and Parquet Complex Types
timestamp	TIMESTAMP	
tinyint (8 bit)	NUMBER(3)	
union	VARCHAR2(*n*) JSON format	DBMS_CLOUD Package Avro, ORC, and Parquet Complex Types

ORC Type	Oracle Type	More Information
varchar	VARCHAR2(*n*)	

DBMS_CLOUD Package Parquet to Oracle Data Type Mapping

Describes the mapping of Parquet data types to Oracle data types.

> **Note:**
>
> Complex types, such as maps, arrays, and structs are supported starting with Oracle Database 19c. See DBMS_CLOUD Package Avro, ORC, and Parquet Complex Types for information on using Parquet complex types.

Parquet Type	Oracle Type
UINT_64	NUMBER(20)
INT_64	NUMBER(19)
UINT_32	NUMBER(10)
INT_32	NUMBER(10)
UINT_16	NUMBER(5)
INT_16	NUMBER(5)
UINT_8	NUMBER(3)
INT_8	NUMBER(3)
BOOL	NUMBER(1)
UTF8 BYTE_ARRAY	VARCHAR2(4000 BYTE)
FLT	BINARY_FLOAT
DBL	BINARY_DOUBLE
DECIMAL(p)	NUMBER(p)
DECIMAL(p,s)	NUMBER(p,s)
DATE	DATE
STRING	VARCHAR2(4000) or VARCHAR2(32767) Depending on the value of the format parameter option `maxvarchar` and on the Autonomous Database `MAX_STRING_SIZE` value. See DBMS_CLOUD Package Format Options for Avro, ORC, or Parquet for more information.
TIME_MILLIS	VARCHAR2(20 BYTE)
TIME_MILLIS_UTC	VARCHAR2(20 BYTE)
TIME_MICROS	VARCHAR2(20 BYTE)
TIME_MICROS_UTC	VARCHAR2(20 BYTE)
TIMESTAMP_MILLIS	TIMESTAMP(3)
TIMESTAMP_MILLIS_UTC	TIMESTAMP(3)
TIMESTAMP_MICROS	TIMESTAMP(6)
TIMESTAMP_MICROS_UTC	TIMESTAMP(6)
TIMESTAMP_NANOS	TIMESTAMP(9)

See DBMS_CLOUD Package Avro, ORC, and Parquet Complex Types for information on using Parquet complex types.

DBMS_CLOUD Package Oracle Data Type to Parquet Mapping

Describes the mapping of Oracle data types to Parquet data types.

Oracle Type	Parquet Type
BINARY_DOUBLE	DBL
BINARY_FLOAT	FLT
DATE	DATE
NUMBER(p,s)	DECIMAL(p,s)
NUMBER(p)	DECIMAL(p)
TIMESTAMP(3)	TIMESTAMP_MILLIS
TIMESTAMP(3)	TIMESTAMP_MILLIS_UTC
TIMESTAMP(6)	TIMESTAMP_MICROS
TIMESTAMP(6)	TIMESTAMP_MICROS_UTC
TIMESTAMP(9)	TIMESTAMP_NANOS
VARCHAR2(4000)	STRING

NLS Session Parameters

The NLS session parameters `NLS_DATE_FORMAT`, `NLS_TIMESTAMP_FORMAT`, `NLS_TIMESTAMP_TZ_FORMAT` and `NLS_NUMERIC_CHARACTERS` define how the date, timestamp, timestamp with time zone format, and radix separator for timestamp with decimal marker should be shown when a table with those column types are queried.

In addition, when you export data using `DBMS_CLOUD.EXPORT_DATA` and specify Parquet output, Autonomous Database reads the values of these parameters from the `NLS_SESSION_PARAMETERS` table. Autonomous Database uses these values to convert the Oracle data types `DATE` or `TIMESTAMP` to Parquet types.

The `NLS_SESSION_PARAMETERS` parameters support an `RR` format mask (two character year specification).

The `RR` format mask for the year is not supported for these parameters when you export data to Parquet with `DBMS_CLOUD.EXPORT_DATA`. An application error is raised if you attempt to export to parquet and the `NLS_SESSION_PARAMETERS` are set to use the `RR` format mask (the default value for the `RR` format depends on the value of the `NLS_TERRITORY` parameter).

When one of the parameters `NLS_DATE_FORMAT`, `NLS_TIMESTAMP_FORMAT` or `NLS_TIMESTAMP_TZ_FORMAT` uses the `RR` format mask, you must change the format value to supported value to export data to Parquet with `DBMS_CLOUD.EXPORT_DATA`. For example:

```
ALTER SESSION SET NLS_DATE_FORMAT = "MM/DD/YYYY";
ALTER SESSION SET NLS_TIMESTAMP_FORMAT = 'YYYY-MM-DD HH:MI:SS.FF';
ALTER SESSION SET  NLS_TIMESTAMP_TZ_FORMAT='YYYY-MM-DD HH:MI:SS.FF TZH:TZM';
```

After you change the value, you can verify the change by querying the `NLS_SESSION_PARAMETERS` view:

```
SELECT value FROM NLS_SESSION_PARAMETERS
       WHERE parameter IN
('NLS_DATE_FORMAT','NLS_TIMESTAMP_FORMAT','NLS_TIMESTAMP_TZ_FORMAT');
```

If `NLS_DATE_FORMAT` is set, it applies to the columns with `DATE` datatype. If `NLS_TIMESTAMP_FORMAT` is set, it applies to the columns with `TIMESTAMP` datattype. If `NLS_TIMESTAMP_TZ_FORMAT` is set, it applies to the columns with `TIMESTAMP WITH TIME ZONE` datatype.

See Date and Time Parameters and NLS Data Dictionary Views for more information.

DBMS_CLOUD Package Avro, ORC, and Parquet Complex Types

Describes the mapping of Avro, ORC, and Parquet complex data types to Oracle data types.

Autonomous Database supports complex data types, including the following complex types:

- struct
- list
- map
- union
- array

When you specify a source file type of Avro, ORC, or Parquet and the source file includes complex columns, Autonomous Database queries return JSON for the complex columns. This simplifies processing of query results; you can use Oracle's powerful JSON parsing features consistently across the file types and data types. The following table shows the format for the complex types in Autonomous Database:

> **Note:**
>
> The complex fields map to `VARCHAR2` columns and `VARCHAR2` size limits apply.

Type	Parquet	ORC	Avro	Oracle
List: sequence of values	List	List	Array	`VARCHAR2` (JSON format)
Map: list of objects with single key	Map	Map	Map	`VARCHAR2` (JSON format)
Union: values of different type	Not Available	Union	Union	`VARCHAR2` (JSON format)
Object: zero or more key-value pairs	Struct	Struct	Record	`VARCHAR2` (JSON format)

If your ORC, Parquet, or Avro source files contain complex types, then you can query the JSON output for these common complex types. For example, the following shows an ORC file, `movie-info.orc`, with a complex type (the same complex type handling applies for Parquet and Avro source files).

Consider the `movie-info.orc` file with the following schema:

```
id    int
original_title string
overview       string
poster_path    string
release_date   string
```

```
vote_count     int
runtime        int
popularity     double
genres         array<struct<id:int,name:string>
```

Notice that each movie is categorized by multiple `genres` using an array of `genres`. The `genres` array is an `array` of `structs` and each item has an `id` (`int`) and a `name` (`string`). The `genres` array is considered a complex type. You can create a table over this ORC file using `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` as follows:

```
BEGIN
DBMS_CLOUD.CREATE_EXTERNAL_TABLE(
        table_name =>'movie_info',
        credential_name =>'OBJ_STORE_CRED',
        file_uri_list =>'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
mytenancy/b/movies/o/movie-info.orc',
        format => '{"type":"orc", "schema": "first"}');
END;
/
```

When you create the external table the database automatically generates the columns based on the schema in the ORC file (if you are using Avro or Parquet, the same applies). For this example, the `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` creates a table in your database as follows:

```
CREATE TABLE "ADMIN"."MOVIE_INFO"
    ( "ID"
      NUMBER(10,0),
      "ORIGINAL_TITLE"  VARCHAR2(4000 BYTE) COLLATE "USING_NLS_COMP",
      "OVERVIEW"        VARCHAR2(4000 BYTE) COLLATE "USING_NLS_COMP",
      "POSTER_PATH"     VARCHAR2(4000 BYTE) COLLATE "USING_NLS_COMP",
      "RELEASE_DATE"    VARCHAR2(4000 BYTE) COLLATE "USING_NLS_COMP",
      "VOTE_COUNT"      NUMBER(10,0),
      "RUNTIME"         NUMBER(10,0),
      "POPULARITY"      BINARY_DOUBLE,
      "GENRES"          VARCHAR2(4000 BYTE) COLLATE "USING_NLS_COMP"
    )  DEFAULT COLLATION "USING_NLS_COMP"
    ORGANIZATION EXTERNAL
     ( TYPE      ORACLE_BIGDATA
       DEFAULT DIRECTORY "DATA_PUMP_DIR"
       ACCESS PARAMETERS
       ( com.oracle.bigdata.credential.name=OBJ_STORE_CRED
         com.oracle.bigdata.fileformat=ORC
   )
       LOCATION
        (
      'https://objectstorage.us-phoenix-1.oraclecloud.com/n/mytenancy/b/
movies/o/movie-info.orc'
           )
     )
   REJECT LIMIT UNLIMITED
   PARALLEL;
 )
```

Now you can query the movie data:

```
SELECT original_title, release_date, genres
     FROM movie_info
     WHERE release_date > '2000'
     ORDER BY original_title;
```

This produces the following output:

```
original_title              release_date   genres
(500) Days of Summer        2009           [{"id":3,"name":"Drama"},
{"id":6,"name":"Comedy"},{"id":17,"name":"Horror"},{"id":19,"name":"Western"},
{"id":18,"name":"War"},{"id":15,"name":"Romance"}]
10,000 BC                   2008           [{"id":6,"name":"Comedy"}]
11:14                       2003           [{"id":9,"name":"Thriller"},
{"id":14,"name":"Family"}]
127 Hours                   2010           [{"id":6,"name":"Comedy"},
{"id":3,"name":"Drama"}]
13 Going on 30              2004           [{"id":6,"name":"Comedy"},
{"id":3,"name":"Drama"},{"id":18,"name":"War"},{"id":15,"name":"Romance"}]
1408                        2007           [{"id":45,"name":"Sci-Fi"},
{"id":6,"name":"Comedy"},{"id":17,"name":"Horror"},{"id":6,"name":"Comedy"},
{"id":18,"name":"War"}]
```

Notice that the complex type `genres` is returned as a JSON array.

To make the JSON data more useful, you can transform the column using Oracle's JSON functions. For example, you can use the JSON "." notation as well as the more powerful transform functions such as `JSON_TABLE`.

See Simple Dot-Notation Access to JSON Data for information on "." notation.

See SQL/JSON Function JSON_TABLE for information on `JSON_TABLE`.

The following example shows a query on the table that takes each value of the array and turns the value into a row in the result set:

```
SELECT original_title, release_date, m.genre_name, genres
    FROM movie_info mi,
       JSON_TABLE(mi.genres, '$.name[*]'
        COLUMNS (genre_name VARCHAR2(25) PATH
      '$')
                 ) AS m
 WHERE rownum < 10;
```

The `JSON_TABLE` creates a row for each value of the array, think outer join, and the struct is parsed to extract the name of the genre. This produces the following output:

```
original_title                  release_date         genre_name        genres
(500) Days of Summer            2009                 Drama
[{"id":3,"name":"Drama"},{"id":6,"name":"Comedy"},

{"id":17,"name":"Horror"},{"id":19,"name":"Western"},
```

```
{"id":18,"name":"War"},{"id":15,"name":"Romance"}]
(500) Days of Summer              2009                  Comedy
[{"id":3,"name":"Drama"},{"id":6,"name":"Comedy"},

{"id":17,"name":"Horror"},{"id":19,"name":"Western"},

{"id":18,"name":"War"},{"id":15,"name":"Romance"}]
(500) Days of Summer              2009                  Horror
[{"id":3,"name":"Drama"},{"id":6,"name":"Comedy"},

{"id":17,"name":"Horror"},{"id":19,"name":"Western"},

{"id":18,"name":"War"},{"id":15,"name":"Romance"}]
(500) Days of Summer              2009                  Western
[{"id":3,"name":"Drama"},{"id":6,"name":"Comedy"},

{"id":17,"name":"Horror"},{"id":19,"name":"Western"},

{"id":18,"name":"War"},{"id":15,"name":"Romance"}]
(500) Days of Summer              2009                  War
[{"id":3,"name":"Drama"},{"id":6,"name":"Comedy"},

{"id":17,"name":"Horror"},{"id":19,"name":"Western"},

{"id":18,"name":"War"},{"id":15,"name":"Romance"}]
(500) Days of Summer              2009                  Romance
[{"id":3,"name":"Drama"},{"id":6,"name":"Comedy"},

{"id":17,"name":"Horror"},{"id":19,"name":"Western"},

{"id":18,"name":"War"},{"id":15,"name":"Romance"}]
10,000 BC                         2008                  Comedy
[{"id":6,"name":"Comedy"}]
11:14                             2003                  Family
[{"id":9,"name":"Thriller"},{"id":14,"name":"Family"}]
11:14                             2003                  Thriller
[{"id":9,"name":"Thriller"},{"id":14,"name":"Family"}]
127 Hours                         2010                  Comedy
[{"id":6,"name":"Comedy"},{"id":3,"name":"Drama"}]
127 Hours                         2010                  Drama
[{"id":6,"name":"Comedy"},{"id":3,"name":"Drama"}]
13 Going on 30                    2004                  Romance
[{"id":6,"name":"Comedy"},{"id":3,"name":"Drama"},

{"id":18,"name":"War"},{"id":15,"name":"Romance"}]
13 Going on 30                    2004                  Comedy
[{"id":6,"name":"Comedy"},{"id":3,"name":"Drama"},

{"id":18,"name":"War"},{"id":15,"name":"Romance"}]
13 Going on 30                    2004                  War
[{"id":6,"name":"Comedy"},{"id":3,"name":"Drama"},

{"id":18,"name":"War"},{"id":15,"name":"Romance"}]
13 Going on 30                    2004                  Drama
[{"id":6,"name":"Comedy"},{"id":3,"name":"Drama"},
```

```
{"id":18,"name":"War"},{"id":15,"name":"Romance"}]
```

DBMS_CLOUD Package Avro, ORC, and Parquet to Oracle Column Name Mapping

Describes rules for how Avro, ORC, and Parquet column names are converted to Oracle column names.

The following are supported for Avro, ORC, and Parquet column names, but may require use of double quotes for Oracle SQL references in external tables. Thus, for ease of use and to avoid having to use double quotes when referencing column names, if possible do not use the following in Avro, ORC, and Parquet column names:

- Embedded blanks
- Leading numbers
- Leading underscores
- Oracle SQL reserved words

The following table shows various types of Avro, ORC, and Parquet column names, and rules for using the column names in Oracle column names in external tables.

Avro, ORC, or Parquet Name	CREATE TABLE Name	Oracle CATALOG	Valid SQL	Notes

part, Part, or PART	part, Part, PART	PART	`select part`
			`select Part`
			`select paRt`
			`select PART`
				Oracle implicitly uppercases unquoted column names

Ord No	"Ord No"	Ord No	`select "Ord No"`	Double quotes are required when there are embedded blanks, which also preserves the character case
__index_key__	"__index_key__"	__index_key__	`select "__index_key__"`	Double quotes are required when there is a leading underscore, which also preserves the character case
6Way	"6Way"	6Way	`select "6Way"`	Double quotes are required when there is a leading numeric digit, which also preserves the character case

Avro, ORC, or Parquet Name	CREATE TABLE Name	Oracle CATALOG	Valid SQL	Notes
create, Create, or CREATE, and so on. (any case variation) partition, Partition, PARTITION, and so on (for an Oracle Reserved word)	"CREATE" "PARTITION"	CREATE PARTITION	`select "CREATE"` `select "PARTITION"`	Double quotes are required around Oracle SQL Reserved words. These are forced to uppercase, but must always be double-quoted when used anywhere in SQL

rowid, Rowid, ROWid, and so on (for ROWID see notes)	rowid		`select "rowid"`
			`select "Rowid"`
			`select "ROWid"`
			`select "rowid"`
				For ROWID, any mixed or lower-case variation of ROWID preserves the case and must always be double-quoted and use the original case variations. Due to the inherent conflict with Oracle ROWID for the table, if you specify upper-case ROWID, it is automatically stored as lower-case "rowid" and must always be double-quoted when referenced.

Notes:

- In general a column name in an external table can be referenced without double quotes.
- Unless there is an embedded blank, a leading underscore ("_") or leading numeric digit ("0" through "9") in the column name, the original case of the column name is preserved, and it must always be referenced with double quotes and using the original case (upper, lower or mixed-case) of the Avro, ORC, or Parquet column name.
- After using `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` to create an external table with the format specified as `avro`, `orc`, or `parquet`, use the `DESCRIBE` command in SQL*Plus to view the table's column names.
- When Oracle SQL Reserved Words are used in Avro, ORC, or Parquet column names, they must always be double-quoted when referenced anywhere in SQL. See Oracle SQL Reserved Words for more information.

DBMS_CLOUD Exceptions

The following table describes exceptions for `DBMS_CLOUD`.

Exception	Code	Description
`reject_limit`	20003	The reject limit of an external table was reached.
`credential_not_exist`	20004	A credential object does not exist.
`table_not_exist`	20005	A table does not exist.
`unsupported_obj_store`	20006	An unsupported object store URI was provided.
`iden_too_long`	20008	An identifier is too long.
`invalid_format`	20009	A format argument is not valid.
`missing_credential`	20010	Mandatory credential object information was not specified.
`invalid_object_uri`	20011	An invalid object URI was provided.
`invalid_partitioning_clause`	20012	An partitioning clause is missing or was not provided.
`unsupported_feature`	20013	An unsupported feature was used that is not existent in the current database version.
`part_not_exist`	20014	A partition or subpartition does not exist, or a table is not a partitioned external table or hybrid partitioned table.
`invalid_table_name`	20016	An invalid table name was used.
`invalid_schema_name`	20017	An invalid schema name was used.
`invalid_dir_name`	20018	An invalid directory name was used.
`invalid_file_name`	20019	An invalid file name was used.
`invalid_cred_attribute`	20020	Invalid credential attributes were specified.
`table_exist`	20021	A table already exists.
`credential_exist`	20022	A credential object already exists.
`invalid_req_method`	20023	A request method is either too long or invalid.
`invalid_req_header`	20024	An invalid request header was specified.
`file_not_exist`	20025	A file does not exist.
`invalid_response`	20026	An HTTP response was not valid.
`invalid_operation`	20027	An invalid task class or ID was specified.
`invalid_user_name`	20028	An invalid username was specified.

DBMS_CLOUD_ADMIN Package

The `DBMS_CLOUD_ADMIN` package provides administrative routines for configuring a database.

- Summary of DBMS_CLOUD_ADMIN Subprograms
  This section covers the `DBMS_CLOUD_ADMIN` subprograms provided with Autonomous Database.
- DBMS_CLOUD_ADMIN Exceptions
  The following table describes exceptions for `DBMS_CLOUD_ADMIN`.

Summary of DBMS_CLOUD_ADMIN Subprograms

This section covers the `DBMS_CLOUD_ADMIN` subprograms provided with Autonomous Database.

Subprogram	Description
ATTACH_FILE_SYSTEM Procedure	This procedure attaches a file system in a directory on your database.
CANCEL_WORKLOAD_CAPTURE Procedure	This procedure cancels the current workload capture.
CREATE_DATABASE_LINK Procedure	This procedure creates a database link to a target database. There are options to create a database link to another Autonomous Database instance, to an Oracle Database that is not an Autonomous Database, or to a non-Oracle Database using Oracle-managed heterogeneous connectivity.
DETACH_FILE_SYSTEM Procedure	This procedure detaches a file system from a directory on your database.
DISABLE_APP_CONT Procedure	This procedure disables database application continuity for the session associated with the specified service name in Autonomous Database.
DISABLE_EXTERNAL_AUTHENTICATION Procedure	This procedure disables external authentication for the Autonomous Database instance.
DISABLE_OPERATOR_ACCESS Procedure	Immediately revokes Cloud Operator access on the Autonomous Database Database instance.
DISABLE_PRINCIPAL_AUTH Procedure	This procedure revokes principal based authentication for the specified provider and applies to the ADMIN user or to the specified user.
DISABLE_RESOURCE_PRINCIPAL Procedure	This procedure disables resource principal credentials for the database or for the specified schema. With a user name specified, other than ADMIN, the procedure revokes the specified schema access to the resource principal credential.
DROP_DATABASE_LINK Procedure	This procedure drops a database link.
ENABLE_APP_CONT Procedure	This procedure enables database application continuity for the session associated with the specified service name in Autonomous Database.
ENABLE_AWS_ARN Procedure	This procedure enables a user to create AWS ARN credentials in Autonomous Database.
ENABLE_EXTERNAL_AUTHENTICATION Procedure	This procedure enables a user to logon to Autonomous Database using the specified external authentication scheme.
ENABLE_FEATURE Procedure	This procedure enables the specified feature on the Autonomous Database instance.
ENABLE_OPERATOR_ACCESS Procedure	Grants the Cloud Operator access to an Autonomous Database instance for a specified number of hours.
ENABLE_PRINCIPAL_AUTH Procedure	This procedure enables principal authentication for the specified provider and applies to the ADMIN user or the specified user.

Subprogram	Description
ENABLE_RESOURCE_PRINCIPAL Procedure	This procedure enables resource principal credential and creates the credential `OCI$RESOURCE_PRINCIPAL`. With a user name specified, other than ADMIN, the procedure grants the specified schema access to the resource principal credential.
FINISH_WORKLOAD_CAPTURE Procedure	This procedure stops the workload capture and uploads capture files to object storage.
PREPARE_REPLAY Procedure	This procedure prepares replay for the refreshable clone.
PURGE_FLASHBACK_ARCHIVE Procedure	This procedure purges historical data from the Flashback Data Archive.
REPLAY_WORKLOAD Procedure	This procedure is overloaded. It initiates the workload replay.
SET_FLASHBACK_ARCHIVE_RETENTION Procedure	This procedure enables ADMIN users to modify the retention period for Flashback Time Travel `flashback_archive`.
START_WORKLOAD_CAPTURE Procedure	This procedure initiates a workload capture.

- ATTACH_FILE_SYSTEM Procedure
  This procedure attaches a file system in the database.
- CANCEL_WORKLOAD_CAPTURE Procedure
  This procedure cancels any ongoing workload capture on the database.
- CREATE_DATABASE_LINK Procedure
- DETACH_FILE_SYSTEM Procedure
  This procedure detaches a file system from the database.
- DISABLE_APP_CONT Procedure
  This procedure disables database application continuity for the session associated with the specified service name in Autonomous Database.
- DISABLE_EXTERNAL_AUTHENTICATION Procedure
  Disables user authentication with external authentication schemes for the database.
- DISABLE_FEATURE Procedure
  This procedure disables the specified feature on the Autonomous Database instance.
- DISABLE_OPERATOR_ACCESS Procedure
  This procedure immediately revokes Cloud Operator access on the Autonomous Database instance.
- DISABLE_PRINCIPAL_AUTH Procedure
  This procedure revokes principal based authentication for a specified provider on Autonomous Database and applies to the ADMIN user or to the specified user.
- DISABLE_RESOURCE_PRINCIPAL Procedure
  Disable resource principal credential for the database or for the specified schema.
- DROP_DATABASE_LINK Procedure
  This procedure drops a database link.
- ENABLE_APP_CONT Procedure
  This procedure enables database application continuity for the session associated with the specified service name in Autonomous Database.

- ENABLE_AWS_ARN Procedure
  This procedure enables an Autonomous Database instance to use Amazon Resource Names (ARNs) to access AWS resources.
- ENABLE_EXTERNAL_AUTHENTICATION Procedure
  Enable users to login to the database with external authentication schemes.
- ENABLE_FEATURE Procedure
  This procedure enables the specified feature on the Autonomous Database instance.
- ENABLE_OPERATOR_ACCESS Procedure
  Oracle Cloud Operations does not access your Autonomous Database instance and access is disallowed by default. When access is required to troubleshoot or mitigate an issue, you can allow a cloud operator access to the database schemas for a limited time.
- ENABLE_PRINCIPAL_AUTH Procedure
  This procedure enables principal authentication on Autonomous Database for the specified provider and applies to the ADMIN user or the specified user.
- ENABLE_RESOURCE_PRINCIPAL Procedure
  Enable resource principal credential for the database or for the specified schema. This procedure creates the credential `OCI$RESOURCE_PRINCIPAL`.
- FINISH_WORKLOAD_CAPTURE Procedure
  This procedure finishes the current workload capture, stops any subsequent workload capture requests to the database, and uploads the capture files to Object Storage.
- PREPARE_REPLAY Procedure
  The `PREPARE_REPLAY` procedure prepares the refreshable clone for a replay.
- PURGE_FLASHBACK_ARCHIVE Procedure
  This procedure enables ADMIN users to purge historical data from Flashback Data Archive. You can either purge all historical data from Flashback Data Archive `flashback_archive` or selective data based on timestamps or System Change Number.
- REPLAY_WORKLOAD Procedure
  This procedure initiates a workload replay on your Autonomous Database instance. The overloaded form enables you to replay the capture files from an Autonomous Database instance, on-premises database, or other cloud service databases.
- SET_FLASHBACK_ARCHIVE_RETENTION Procedure
  This procedure allows ADMIN users to modify the retention period for Flashback Data Archive `flashback_archive`.
- START_WORKLOAD_CAPTURE Procedure
  This procedure initiates a workload capture on your Autonomous Database instance.

ATTACH_FILE_SYSTEM Procedure

This procedure attaches a file system in the database.

The `DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM` procedure attaches a file system in your database and stores information about the file system in the `DBA_CLOUD_FILE_SYSTEMS` view.

Syntax

```
DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM (
    file_system_name         IN VARCHAR2,
    file_system_location     IN VARCHAR2,
    directory_name           IN VARCHAR2,
    description              IN VARCHAR2 DEFAULT NULL,
```

```
    params                    IN CLOB DEFAULT NULL
);
```

Parameters

Parameter	Description
`file_system_name`	Specifies the name of the file system. This parameter is mandatory.
`file_system_location`	Specifies the location of the file system. The value you supply with `file_system_location` consists of a Fully Qualified Domain Name (FQDN) and a file path in the form: *FQDN*`:`*file_path*. For example: • FQDN: `myhost.sub000445.myvcn.oraclevcn.com` For Oracle Cloud Infrastructure File Storage set the FQDN in **Show Advanced Options** when you create a file system. See Creating File Systems for more information. • File Path: `/results` This parameter is mandatory.
`directory_name`	Specifies the directory name for the attached file system. The directory must exist. This parameter is mandatory.
`description`	(Optional) Provides a description of the task.
`params`	A JSON string that provides an additional parameter for the file system. • `nfs_version`: Specifies the NFS version to use when NFS is attached (NFSv3 or NFSv4). Valid values: 3, 4. Default value: 3

Examples:

Attach to an NFSv3 file system:

```
BEGIN
   DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM (
    file_system_name      => 'FSS',
    file_system_location  => 'myhost.sub000445.myvcn.oraclevcn.com:/results',
    directory_name        => 'FSS_DIR',
    description           => 'Source NFS for sales data'
  );
END;
/
```

Attach to an NFSv4 file system:

```
BEGIN
   DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM (
    file_system_name      => 'FSS',
    file_system_location  => 'myhost.sub000445.myvcn.oraclevcn.com:/results',
    directory_name        => 'FSS_DIR',
    description           => 'Source NFS for sales data',
    params                => JSON_OBJECT('nfs_version' value 4)
);
```

```
END;
/
```

Usage Notes

- To run this procedure you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.
- You must have `WRITE` privilege on the directory object in the database to attach a file system using `DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM`.
- The `DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM` procedure can only attach a private File Storage Service in databases with Private Endpoints enabled.

  See OCI File Storage Service and Configure Network Access with Private Endpoints for more information.
- The `DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM` procedure looks up the Network File System hostname on the customer's virtual cloud network (VCN). The error `"ORA-20000: Mounting NFS fails"` is returned if the hostname specified in the location cannot be located.
- Oracle Cloud Infrastructure File Storage uses NFSv3 to share
- If you attach to non-Oracle Cloud Infrastructure File Storage systems, the procedure supports NFSv3 and NFSv4
- If you have an attached NFS server that uses NFSv3 and the NFS version is updated to NFSv4 in the NFS server, you must run `DBMS_CLOUD_ADMIN.DETACH_FILE_SYSTEM` and then `DBMS_CLOUD_ADMIN.ATTACH_FILE_SYSTEM` (using the `params` parameter with `nfs_version` set to 4). This attaches NFS with the matching protocol so that Autonomous Database can access the NFSv4 Server. Without detaching and then reattaching, the NFS server will be inaccessible and you may see an error such as: `"Protocol not supported"`.

CANCEL_WORKLOAD_CAPTURE Procedure

This procedure cancels any ongoing workload capture on the database.

This procedure cancels the current workload capture and enables refresh on the refreshable clone.

Syntax

```
DBMS_CLOUD_ADMIN.CANCEL_WORKLOAD_CAPTURE;
```

Example

```
BEGIN
   DBMS_CLOUD_ADMIN.CANCEL_WORKLOAD_CAPTURE;
END;
/
```

Usage Note

- To run this procedure you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.

CREATE_DATABASE_LINK Procedure

This procedure creates a database link to a target database in the schema calling the API.

The overloaded forms support the following:

- When you use the `gateway_params` parameter, this enables you to create a database link with Oracle-managed heterogeneous connectivity where the link is to a supported non-Oracle database.
- When you use the `rac_hostnames` parameter, this enables you to create a database link from an Autonomous Database on a private endpoint to a target Oracle RAC database. In this case, you use the `rac_hostnames` parameter to specify the host names of one or more individual nodes of the target Oracle RAC database.

Syntax

```
DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
       db_link_name          IN VARCHAR2,
       hostname              IN VARCHAR2,
       port                  IN NUMBER,
       service_name          IN VARCHAR2,
       ssl_server_cert_dn    IN VARCHAR2 DEFAULT,
       credential_name       IN VARCHAR2 DEFAULT,
       directory_name        IN VARCHAR2 DEFAULT,
       gateway_link          IN BOOLEAN DEFAULT,
       public_link           IN BOOLEAN DEFAULT,
       private_target        IN BOOLEAN DEFAULT
       gateway_params        IN CLOB DEFAULT);


DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
       db_link_name          IN VARCHAR2,
       rac_hostnames         IN CLOB,
       port                  IN NUMBER,
       service_name          IN VARCHAR2,
       ssl_server_cert_dn    IN VARCHAR2 DEFAULT,
       credential_name       IN VARCHAR2 DEFAULT,
       directory_name        IN VARCHAR2 DEFAULT,
       gateway_link          IN BOOLEAN DEFAULT,
       public_link           IN BOOLEAN DEFAULT,
       private_target        IN BOOLEAN DEFAULT);
```

Parameters

Parameter	Description
db_link_name	The name of the database link to create.

Parameter	Description
hostname	The hostname for the target database. Specifying localhost for hostname as is not allowed. When you specify a connection with Oracle-managed heterogeneous connectivity by supplying the gateway_params parameter, note the following: • When the db_type value is google_bigquery the hostname is not used and you can provide value such as example.com. • When the db_type value is snowflake the hostname is the Snowflake account identifier. To find your Snowflake account identifier, see Account Identifier Formats by Cloud Platform and Region. Use this parameter or rac_hostnames, do not use both. The DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK input should mention the scan name of the target Dedicated Autonomous Database as the 'hostname' parameter.
rac_hostnames	Specifies hostnames for the target Oracle RAC database. The value is a JSON array that specifies one or more individual host names for the nodes of the target Oracle RAC database. Multiple host names can be passed in JSON, separated by a ",". For example: `'["sales1-svr1.domain", "sales1-svr2.domain", "sales1-svr3.domain"]'` When the target is an Oracle RAC database, use the rac_hostnames parameter to specify one or more hostnames with DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK. This allows you to take advantage of the high availability capabilities of Oracle RAC. Using an IP address, a SCAN IP, or a SCAN hostname in the rac_hostnames value is not supported. When you specify a list of host names in the rac_hostnames parameter, CREATE_DATABASE_LINK uses all of the specified host names as addresses in the connect string. If one of the specified hosts is not available on the target Oracle RAC database, Autonomous Database automatically attempts to connect using another host name from the list. Use this parameter or hostname, do not use both. Specifying localhost for a rac_hostname value is not allowed.

Parameter	Description
port	Specifies the port for the connections to the target database. When you specify a connection with Oracle-managed heterogeneous connectivity using the gateway_params parameter, set the port based on the db_type value: • awsredshift: use port 5439 • azure: use port 1433 • db2: use port 2500 for Db2 versions >= 11.5.6 • db2: use port 5000 for Db2 versions <= 11.5.5 • google_analytics: use port 443 • google_bigquery: use port 443 • google_drive: use port 443 • hive: use port 433 • mongodb: use port 27017 • mysql: use port 3306 • mysql_community: use port 3306 • postgres: use port 5432 • salesforce: use port 19937 • servicenow: use port 443 • sharepoint: use port 443 • snowflake: use port 443 • youtube: use port 443 See Oracle-Managed Heterogeneous Connectivity Database Types and Ports for more information.
service_name	The service_name for the database to link to. For a target Autonomous Database, find the service name by one of the following methods: • Look in the tnsnames.ora file in the wallet.zip that you download from an Autonomous Database for your connection. • Click **Database connection** on the Oracle Cloud Infrastructure Console. In the **Connection Strings** area, each connection string includes a service_name entry with the connection string for the corresponding service. When both Mutual TLS (mTLS) and TLS connections are allowed, under **TLS authentication** select **TLS** to view the TNS names and connection strings for connections with TLS authentication. See View TNS Names and Connection Strings for an Autonomous Database Instance for more information. • Query V$SERVICES view. For example: `SELECT name FROM V$SERVICES;` When you specify a connection with Oracle-managed heterogeneous connectivity using the gateway_params parameter, the service_name is the database name of the non-Oracle database.

Parameter	Description
ssl_server_cert_dn	The DN value found in the server certificate. Oracle-managed heterogeneous connectivity is preconfigured with a wallet that contains most of the common trusted root and intermediate SSL certificates. The ssl_server_cert_dn must be NULL when you supply the gateway_params parameter or do not include the ssl_server_cert_dn parameter (the default value is NULL). **Public Endpoint Link to an Autonomous Database Target without a Wallet:** To connect to an Autonomous Database target on a public endpoint without a wallet (TLS): • The directory_name parameter must be NULL. • The ssl_server_cert_dn parameter must be NULL or do not include this parameter (the default value is NULL). **Private Endpoint Link without a Wallet:** To connect to an Oracle Database on a private endpoint without a wallet: • The target database must be on a private endpoint. • The directory_name parameter must be NULL. • The ssl_server_cert_dn parameter must be NULL or do not include this parameter (the default is NULL). • The private_target parameter must be TRUE.
credential_name	The name of a stored credential created with DBMS_CLOUD.CREATE_CREDENTIAL. This is the credentials to access the target database.
directory_name	The directory for the cwallet.sso file. The default value for this parameter is 'data_pump_dir'. Oracle-managed heterogeneous connectivity is preconfigured with a wallet that contains most of the common trusted root and intermediate SSL certificates. The directory_name parameter is not required when you supply the gateway_params parameter. **Public Endpoint Link to an Autonomous Database Target without a Wallet:** To connect to an Autonomous Database on a public endpoint without a wallet (TLS): • The directory_name parameter must be NULL. • The ssl_server_cert_dn parameter must be NULL or do not include this parameter (the default value is NULL). In addition, to connect to an Autonomous Database with TCP, the ssl_server_cert_dn parameter must be NULL or do not include this parameter (the default value is NULL). **Private Endpoint Link without a Wallet:** To connect to a target Oracle Database on a private endpoint without a wallet: • The target database must be on a private endpoint. • The directory_name parameter must be NULL. • The ssl_server_cert_dn parameter must be NULL or do not include this parameter (the default value is NULL). • The private_target parameter must be TRUE.

Parameter	Description
gateway_link	Indicates if the database link is created to another Oracle Database or to an Oracle Database Gateway. If gateway_link is set to FALSE, this specifies a database link to another Autonomous Database or to another Oracle Database. If gateway_link is set to TRUE, this specifies a database link to a non-Oracle system. This creates a connect descriptor in the database link definition that specifies (HS=OK). When gateway_link is set to TRUE and gateway_params is NULL, this specifies a database link to a customer-managed Oracle gateway. The default value for this parameter is FALSE.
public_link	Indicates if the database link is created as a public database link. To run DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK with this parameter set to TRUE, the user invoking the procedure must have EXECUTE privilege on the credential associated with the public database link and must have the CREATE PUBLIC DATABASE LINK system privilege. The EXECUTE privilege on the credential can be granted either by the ADMIN user or by the credential owner. The default value for this parameter is FALSE.
private_target	When a database link accesses a hostname that needs to be resolved in a VCN DNS server, specify the private_target parameter with value TRUE. When private_target is TRUE, the hostname parameter must be a single hostname (on a private endpoint, using an IP address, a SCAN IP, or a SCAN hostname is not supported). The default value for this parameter is FALSE.

Parameter	Description
gateway_params	db_type This parameter specifies the target database type for Oracle-managed heterogeneous connectivity to connect to non-Oracle databases. The db_type value is one of: • awsredshift • azure * See Usage Notes for additional supported gateway_params when db_type is azure. • db2 • google_analytics • google_bigquery * See Usage Notes for additional supported gateway_params when db_type is google_bigquery. • google_drive * See Usage Notes for additional supported gateway_params when db_type is google_drive. • hive * See Usage Notes for additional supported gateway_params when db_type is hive. • mongodb • mysql • postgres • salesforce * See Usage Notes for additional supported gateway_params when db_type is salesforce. • servicenow * See Usage Notes for additional supported gateway_params when db_type is servicenow. • sharepoint * See Usage Notes for additional supported gateway_params when db_type is sharepoint. • snowflake * See Usage Notes for additional supported gateway_params when db_type is snowflake. • youtube * See Usage Notes for additional supported gateway_params when db_type is youtube. • NULL When gateway_params is NULL and gateway_link is set to TRUE, this specifies a database link to a customer-managed Oracle gateway. Specify the parameter with the json_object form. For example: gateway_params => json_object('db_type' value 'awsredshift') See Oracle-Managed Heterogeneous Connectivity Database Types and Ports for required port values for each database type. When gateway_params is NULL and private_target is TRUE, if directory_name is NULL, a TCP-based database link is created. When gateway_params is NULL and private_target is TRUE, if directory_name is NULL, a TCPS-based database link is created. enable_ssl: For a database on a private endpoint, the gateway_params parameter supports the optional enable_ssl parameter. Set this parameter to

Parameter	Description
	`true` to support SSL/TLS connections for a remote target database on a private endpoint. By default `enable_ssl` is `false`. For example: `gateway_params => JSON_OBJECT( 'db_type' value 'snowflake', 'enable_ssl' value true)`

Usage Notes

- When you specify the `gateway_params` parameter, for some `db_type` values, additional `gateway_params` parameters are supported:

db_type	Additional gateway_params Values
`azure`	When the `db_type` value is `azure` there are two optional parameters, `auth_method` and `domain` to support NTLM/AD authentication. When NTLMv2 is configured, set `auth_method=10` and set `domain` to the Windows domain value. To use an Azure Active Directory password, set `auth_method=13`. Do not use `auth_method=13` to access an Oracle on-premises Database. See Microsoft SQL Server Authentication Method for more information.
`google_analytics`	When the `db_type` is `google_analytics`, the credential you specify must be a Google OAuth credential (`gcp_oauth2`). See CREATE_CREDENTIAL Procedure for more information.
`google_bigquery`	When the `db_type` is `google_bigquery`, the credential you specify must be a Google OAuth credential (`gcp_oauth2`) See CREATE_CREDENTIAL Procedure for more information. When `db_type` is `google_bigquery`, the parameter `project` is valid. This parameter specifies the project name for `google_bigquery` and is required. The table name you specify when you use `SELECT` with Google BigQuery must be in quotes. For example: `SELECT * FROM "sales"@GOOGLE_BIGQUERY_LINK`

db_type	Additional gateway_params Values

google_drive	For db_type google_drive, the credential you supply with DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK must be an OAuth type credential using the params parameter with gcp_oauth2 values specified (client_id, client_secret, and refresh_token).
	To connect to Google Drive and get data you must supply the gateway parameters directory_name and file_name. These parameters specify a model file (auto REST file) that maps the JSON response to the relational model. The model file specifies the endpoints, table mapping, and HTTP response code for processing the JSON response. See Model file syntax and Example Model file for more information.
	When you use gateway_params parameter with db_type google_drive you must supply the gateway_params parameter db_type with the value 'google_drive' and the directory_name and file_name parameters, along with OAuth type credentials:
	The directory_name parameter specifies the directory with the REST config file. You could create this directory as follows:
	`create or replace directory google_drive_dir as 'GOOGLE_DRIVE_DIR';`
	Obtain and download the auto REST file to the specified directory. For example:
	`exec DBMS_CLOUD.get_object('google_drive_dir_cred', 'https://objectstorage.<...>/ google_drive.rest','GOOGLE_DRIVE_DIR');`
	Set the file_name value to the name of the auto REST file you downloaded, "google_drive.rest". Then you can use the REST config file with OAuth 2.0.

hive	When db_type is hive, the parameter http_path is valid. This parameter specifies the HttpPath value, if required, to connect to the Hive instance.
salesforce	When the db_type is salesforce, the parameter: security_token is valid. A security token is a case-sensitive alphanumeric code. Supplying a security_token value is required to access Salesforce. For example: `gateway_params => JSON_OBJECT( 'db_type' value 'salesforce', 'security_token' value 'security_token_value' )` See Reset Your Security Token for more information.

db_type	Additional gateway_params Values

servicenow	To connect to ServiceNow and get data you must supply the gateway parameters directory_name and file_name. These parameters specify a model file (REST config file) that maps the JSON response to the relational model. The model file specifies the endpoints, table mapping, and HTTP response code for processing the JSON response. See Model file syntax and Example Model file for more information.
	When you use gateway_params parameter with db_type servicenow, there are two supported options:
	– Basic Authentication: you must supply the gateway_params parameter db_type with the value 'servicenow', and supply the directory_name and file_name parameters along with username/password type credentials.
	– OAuth 2.0 Authentication: you must supply the gateway_params parameter db_type with the value 'servicenow', and the directory_name, file_name, and token_uri parameters, along with OAuth type credentials.
	The directory_name parameter specifies the directory with the ServiceNow REST config file. You could create this directory as follows:
	`create or replace directory servicenow_dir as 'SERVICENOW_DIR';`
	Obtain and download the ServiceNow REST config file to the specified directory. For example:
	`exec DBMS_CLOUD.get_object('servicenow_dir_cred', 'https://objectstorage.<...>/servicenow.rest','SERVICENOW_DIR');`
	Set the file_name value to the name of the REST config file you downloaded, "servicenow.rest".
	Then you can use the ServiceNow REST config file with either basic authentication or OAuth2.0.
	Then you can use the ServiceNow REST config file with either basic authentication or OAuth2.0. See HETEROGENEOUS_CONNECTIVITY_INFO View for samples.

db_type	Additional gateway_params Values
sharepoint	When you use gateway_params parameter with db_type sharepoint, also specify values for auth_uri, scope, service_url, and token_uri. For db_type sharepoint, the credential you supply with DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK must be an OAuth type credential using the params parameter with gcp_oauth2 values specified (client_id, client_secret, and refresh_token). See CREATE_CREDENTIAL Procedure for more information. Generate the refresh_token for *tenant_id*.sharepoint.com/.default offline_access. See the following for more information: Determine the scope and OAuth 2.0 authentication. Set values for gateway_params for db_type sharepoint: – auth_uri: Set the auth_uri from Azure. See Obtain application client information and endpoints for more information. – scope: Set the scope. See Connection option descriptions for more information. – service_url: Set the service_url. See Service URL for more information. – token_uri: Get the token_uri from Azure. See Obtain application client information and endpoints for more information.

snowflake	When the db_type is SNOWFLAKE, the optional parameters: role, schema, and warehouse are valid. These values specify a different schema, role, or warehouse value, other than the default. For example:
	`gateway_params => JSON_OBJECT(`
	`'db_type' value 'snowflake',`
	`'role' value 'ADMIN',`
	`'schema' value 'PUBLIC',`
	`'warehouse' value 'TEST' )`

db_type	Additional gateway_params Values

youtube	To connect to Youtube and get data you must supply the gateway parameters directory_name and file_name. These parameters specify a model file (REST config file) that maps the JSON response to the relational model. The model file specifies the endpoints, table mapping, and HTTP response code for processing the JSON response. See Model file syntax and Example Model file for more information.
	When you use gateway_params parameter with db_type youtube, the credential you specify must be a Google OAuth credential (gcp_oauth2) See CREATE_CREDENTIAL Procedure for more information.
	OAuth 2.0 Authentication: you must supply the gateway_params parameter db_type with the value youtube, and the directory_name and file_name parameters, along with OAuth type credentials.
	The directory_name parameter specifies the directory with the ServiceNow REST config file. You could create this directory as follows:
	`create or replace directory youtube_dir as 'YOUTUBE_DIR';`
	Obtain and download the ServiceNow REST config file to the specified directory. For example:
	`exec DBMS_CLOUD.get_object('youtube_dir_cred', 'https://objectstorage.<...>/ youtube.rest','YOUTUBE_DIR');`
	Set the file_name value to the name of the REST config file you downloaded, "youtube.rest".
	Then you can use the ServiceNow REST config file with OAuth2.0. See HETEROGENEOUS_CONNECTIVITY_INFO View for samples.

- To run DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK with a user other than ADMIN, you need to grant EXECUTE and CREATE DATABASE LINK privileges to that user. For example, run the following command as ADMIN to grant privileges to adb_user:

```
GRANT EXECUTE ON DBMS_CLOUD_ADMIN TO adb_user;
GRANT CREATE DATABASE LINK TO adb_user;
```

  In addition, when you create a Database Link in a schema other than the ADMIN schema, for example in a schema named adb_user, the adb_user schema must own the credential you use with DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK.

- Only one wallet file is valid per directory specified with the directory_name parameter. You can only upload one cwallet.sso at a time to the directory you choose for wallet files. This means with a cwallet.sso in a directory, you can only create database links to the databases for which the wallet in that directory is valid. To use multiple cwallet.sso files with database links you need to create additional directories and put each cwallet.sso in a different directory.

  See Create Directory in Autonomous Database for information on creating directories.

- To create a database link to an Autonomous Database, set `GLOBAL_NAMES` to `FALSE` on the source database (non-Autonomous Database).

```
SQL> ALTER SYSTEM SET GLOBAL_NAMES = FALSE;

System altered.

SQL> SHOW PARAMETER GLOBAL_NAMES
NAME                      TYPE        VALUE
------------------------  ----------- -----------
global_names              boolean     FALSE
```

- When the `private_target` parameter is `TRUE`, the `hostname` parameter specifies a private host inside the VCN.

Examples

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
     credential_name => 'DB_LINK_CRED',
     username => 'adb_user',
     password => 'password');
  DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
     db_link_name => 'SALESLINK',
     hostname => 'adb.eu-frankfurt-1.oraclecloud.com',
     port => '1522',
     service_name => 'example_medium.adb.example.oraclecloud.com',
     ssl_server_cert_dn => 'CN=adb.example.oraclecloud.com,OU=Oracle BMCS
FRANKFURT,O=Oracle Corporation,L=Redwood City,ST=California,C=US',
     credential_name => 'DB_LINK_CRED');
END;
/


BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'AWS_REDSHIFT_LINK_CRED',
    username => 'NICK',
    password => 'password'
  );
DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
          db_link_name => 'AWSREDSHIFT_LINK',
          hostname => 'example.com',
          port => '5439',
          service_name => 'example_service_name',
          ssl_server_cert_dn => NULL,
          credential_name => 'AWS_REDSHIFT_LINK_CRED',
          gateway_params => JSON_OBJECT('db_type'  value 'awsredshift'));
END;
/


BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'PRIVATE_ENDPOINT_CRED',
    username => 'db_user',
```

```
    password => 'password'
  );
DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
          db_link_name => 'PRIVATE_ENDPOINT_DB_LINK',
          hostname => 'exampleHostname',
          port => '1521',
          service_name => 'exampleServiceName',
          credential_name => 'PRIVATE_ENDPOINT_CRED',
          ssl_server_cert_dn => NULL,
          directory_name => NULL,
          private_target => TRUE);
END;
/


BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'GOOGLE_BIGQUERY_CRED',
    params => JSON_OBJECT( 'gcp_oauth2' value JSON_OBJECT(
          'client_id' value 'client_id',
          'client_secret' value 'client_secret',
          'refresh_token' value 'refresh_token' )));

DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
          db_link_name => 'GOOGLE_BIGQUERY_LINK',
          hostname => 'example.com',
          port => '443',
          service_name => 'example_service_name',
          credential_name => 'GOOGLE_BIGQUERY_CRED',
          gateway_params => JSON_OBJECT(
                    'db_type' value 'google_bigquery',
                    'project' value 'project_name1' ));
END;
/
```

The table name you specify when you use `SELECT` with Google BigQuery or Google Analytics must be in quotes. For example:

```
SELECT * FROM "sales"@GOOGLE_BIGQUERY_LINK
```

Use the `rac_hostnames` parameter with a target Oracle RAC database on a private endpoint.

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
     credential_name => 'DB_LINK_CRED1',
     username => 'adb_user',
     password => 'password');
      DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK(
     db_link_name => 'SALESLINK',
     rac_hostnames => '["sales1-svr1.example.adb.us-
ashburn-1.oraclecloud.com",
                        "sales1-svr2.example.adb.us-
ashburn-1.oraclecloud.com",
                        "sales1-svr3.example.adb.us-
ashburn-1.oraclecloud.com"]',
```

```
     port => '1522',
     service_name => 'example_high.adb.oraclecloud.com',
     ssl_server_cert_dn => 'CN=adb.example.oraclecloud.com,OU=Oracle BMCS
FRANKFURT,O=Oracle Corporation,L=Redwood City,ST=California,C=US',
     credential_name => 'DB_LINK_CRED1',
     directory_name => 'EXAMPLE_WALLET_DIR',
     private_target => TRUE);
END;
/
```

DETACH_FILE_SYSTEM Procedure

This procedure detaches a file system from the database.

The `DBMS_CLOUD_ADMIN.DETACH_FILE_SYSTEM` procedure detaches a file system from your database. In addition to that, the `DBMS_CLOUD_ADMIN.DETACH_FILE_SYSTEM` procedure also removes the information about the file system from the `DBA_CLOUD_FILE_SYSTEMS` view.

Syntax

```
DBMS_CLOUD_ADMIN.DETACH_FILE_SYSTEM(
    file_system_name         IN VARCHAR2
);
```

Parameters

Parameter	Description
`file_system_name`	Specifies the name of the file system. This parameter is mandatory.

Example:

```
BEGIN
   DBMS_CLOUD_ADMIN.DETACH_FILE_SYSTEM (
    file_system_name      => 'FSS'
  );
END;
/
```

Usage Notes

- To run this procedure, you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.
- You must have the `WRITE` privilege on the directory object in the database, to detach a file system from a directory using the `DBMS_CLOUD_ADMIN.DETACH_FILE_SYSTEM` procedure.
- The `DBMS_CLOUD_ADMIN.DETACH_FILE_SYSTEM` procedure can only detach a private File Storage Service in databases with Private Endpoints enabled.

  See OCI File Storage Service and Configure Network Access with Private Endpoints for more information.

- The `DBMS_CLOUD_ADMIN.DETACH_FILE_SYSTEM` procedure looks up the Network File System hostname on the customer's virtual cloud network (VCN). The error `"ORA-20000: Mounting NFS fails"` is returned if the hostname specified in the location cannot be located.

DISABLE_APP_CONT Procedure

This procedure disables database application continuity for the session associated with the specified service name in Autonomous Database.

Syntax

```
DBMS_CLOUD_ADMIN.DISABLE_APP_CONT(
    service_name      IN VARCHAR2);
```

Parameters

Parameter	Description
`service_name`	The `service_name` for the Autonomous Database service. To find service names: • Look in the `tnsnames.ora` file in the `wallet.zip` that you download from an Autonomous Database for your connection. • Click **Database connection** on the Oracle Cloud Infrastructure Console. In the **Connection strings** area, each connection string includes a `service_name` entry that contains the connection string for the corresponding service. When both Mutual TLS (mTLS) and TLS connections are allowed, under **TLS authentication** select **TLS** to view the TNS names and connection strings for connections with TLS authentication. See View TNS Names and Connection Strings for an Autonomous Database Instance for more information. • Query `V$SERVICES` view. For example: `SELECT name FROM V$SERVICES;`

Usage Notes

See Overview of Application Continuity for more information on Application Continuity.

Example

```
BEGIN
    DBMS_CLOUD_ADMIN.DISABLE_APP_CONT(
        service_name => 'nv123abc1_adb1_high.adb.oraclecloud.com' );
END;
/
```

Verify the value as follows:

```
SELECT name, failover_type FROM DBA_SERVICES;

NAME                                                  FAILOVER_TYPE
----------------------------------------------------- --------------
nv123abc1_adb1_high.adb.oraclecloud.com
```

DISABLE_EXTERNAL_AUTHENTICATION Procedure

Disables user authentication with external authentication schemes for the database.

Syntax

```
DBMS_CLOUD_ADMIN.DISABLE_EXTERNAL_AUTHENTICATION;
```

Exceptions

Exception	Error	Description
invalid_ext_auth	ORA-20004	See the accompanying message for a detailed explanation.

Example

```
BEGIN
   DBMS_CLOUD_ADMIN.DISABLE_EXTERNAL_AUTHENTICATION;
END;
/

PL/SQL procedure successfully completed.
```

DISABLE_FEATURE Procedure

This procedure disables the specified feature on the Autonomous Database instance.

Syntax

```
DBMS_CLOUD_ADMIN.DISABLE_FEATURE(
    feature_name    IN  VARCHAR2);
```

Parameters

Parameter	Description
feature_name	Specifies the feature type to be disabled. Supported values are: • 'AUTO_DST_UPGRADE': Disable AUTO DST feature. • 'AUTO_DST_UPGRADE_EXCL_DATA': Disable AUTO DST EXCL DATA feature. • 'ORAMTS': Disable OraMTS feature. • 'OWM': Disable Oracle Workspace Manager. • 'WORKLOAD_AUTO_REPLAY': Disable workload auto replay feature. This parameter is mandatory.

Examples

```
BEGIN
DBMS_CLOUD_ADMIN.DISABLE_FEATURE(
    feature_name => 'ORAMTS');
```

```
END;
/


BEGIN
DBMS_CLOUD_ADMIN.DISABLE_FEATURE(
    feature_name => 'AUTO_DST_UPGRADE');
END;
/


BEGIN
DBMS_CLOUD_ADMIN.DISABLE_FEATURE(
    feature_name => 'AUTO_DST_UPGRADE_EXCL_DATA');
END;
/


BEGIN
DBMS_CLOUD_ADMIN.DISABLE_FEATURE(
    feature_name => 'OWM');
END;
/


BEGIN
DBMS_CLOUD_ADMIN.DISABLE_FEATURE(
    feature_name => 'WORKLOAD_AUTO_REPLAY');
END;
/
```

Usage Notes

- To disable the `OraMTS`, `AUTO_DST_UPGRADE`, `AUTO_DST_UPGRADE_EXCL_DATA`, `OWM`, or `WORKLOAD_AUTO_REPLAY` features for your Autonomous Database instance, you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.
- When both `AUTO_DST_UPGRADE` and `AUTO_DST_UPGRADE_EXCL_DATA` are disabled, if new time zone versions are available the Autonomous Database instance does not upgrade to use the latest available time zone files.
- Query `dba_cloud_config` to verify that `AUTO_DST_UPGRADE` is disabled.

  ```
  SELECT param_name, param_value FROM dba_cloud_config WHERE
         LOWER(param_name) = 'auto_dst_upgrade';

  0 rows selected.
  ```

- Query `dba_cloud_config` to verify that `AUTO_DST_UPGRADE_EXCL_DATA` is disabled.

  ```
  SELECT param_name, param_value FROM dba_cloud_config WHERE
         LOWER(param_name) = 'auto_dst_upgrade_excl_data';

  0 rows selected.
  ```

DISABLE_OPERATOR_ACCESS Procedure

This procedure immediately revokes Cloud Operator access on the Autonomous Database instance.

Syntax

```
DBMS_CLOUD_ADMIN.DISABLE_OPERATOR_ACCESS;
```

Example

```
BEGIN
    DBMS_CLOUD_ADMIN.DISABLE_OPERATOR_ACCESS;
END;
/
```

DISABLE_PRINCIPAL_AUTH Procedure

This procedure revokes principal based authentication for a specified provider on Autonomous Database and applies to the ADMIN user or to the specified user.

Syntax

```
DBMS_CLOUD_ADMIN.DISABLE_PRINCIPAL_AUTH(
       provider    IN VARCHAR2,
       username    IN VARCHAR2 DEFAULT 'ADMIN' );
```

Parameters

Parameter	Description
`provider`	Specifies the type of provider. Valid values: • `AWS` • `AZURE` • `GCP` • `OCI`
`username`	Specifier the user to disable principal based authentication for. A null value is valid for the `username`. If `username` is not specified, the procedure applies for the "`ADMIN`" user.

Usage Notes

- When the provider value is `AZURE` and the `username` is `ADMIN`, the procedure disables Azure service principal based authentication on Autonomous Database and deletes the Azure application on the Autonomous Database instance.
- When the provider value is `AZURE` and the `username` is a user other than the `ADMIN` user, the procedure revokes the privileges from the specified user. The `ADMIN` user and other users that are enabled to use the Azure service principal can continue to use `ADMIN.AZURE$PA` and the application that is created for the Autonomous Database instance remains on the instance.

Examples

```
BEGIN
     DBMS_CLOUD_ADMIN.DISABLE_PRINCIPAL_AUTH(
       provider => 'AZURE',
       username => 'SCOTT');
END;
/


BEGIN
     DBMS_CLOUD_ADMIN.DISABLE_PRINCIPAL_AUTH(
       provider => 'GCP');
END;
/
```

DISABLE_RESOURCE_PRINCIPAL Procedure

Disable resource principal credential for the database or for the specified schema.

Syntax

```
DBMS_CLOUD_ADMIN.DISABLE_RESOURCE_PRINCIPAL(
   username         IN VARCHAR2);
```

Parameter

Parameter	Description
`username`	Specifies an optional user name. The name of the database schema to remove resource principal access. If you do not supply a `username`, the `username` is set to ADMIN and the command removes the `OCI$RESOURCE_PRINCIPAL` credential.

Exceptions

Exception	Error	Description
`resource principal is already disabled`	`ORA-20031`	If you attempt to disable the resource principal when it is already disabled.

Usage Notes

- Resource principal is not available with refreshable clones.
- You must set up a dynamic group and policies for the dynamic group before you call `DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL`.

  See the following for more information on creating policies, creating a dynamic group, and creating rules:

  – Use Resource Principal to Access Oracle Cloud Infrastructure Resources

  – Managing Dynamic Groups

  - Managing Policies

- Verify that a resource principal credential is enabled by querying one of the views: DBA_CREDENTIALS or ALL_TAB_PRIVS.

  For example, as the ADMIN user query the view `DBA_CREDENTIALS`:

```
SELECT owner, credential_name FROM dba_credentials
        WHERE credential_name = 'OCI$RESOURCE_PRINCIPAL' AND owner =
'ADMIN';

OWNER  CREDENTIAL_NAME
-----  ----------------------
ADMIN  OCI$RESOURCE_PRINCIPAL
```

  For example, as a non-ADMIN user query the view `ALL_TAB_PRIVS`:

```
SELECT grantee, table_name, grantor, FROM ALL_TAB_PRIVS
          WHERE  grantee = 'ADB_USER';

GRANTEE   TABLE_NAME GRANTOR
--------- -----------------------------------
ADB_USER  OCI$RESOURCE_PRINCIPAL ADMIN
```

Example

```
EXEC DBMS_CLOUD_ADMIN.DISABLE_RESOURCE_PRINCIPAL();

PL/SQL procedure successfully completed.

SQL> select owner, credential_name from dba_credentials where credential_name
= 'OCI$RESOURCE_PRINCIPAL';

No rows selected.
```

DROP_DATABASE_LINK Procedure

This procedure drops a database link.

Syntax

```
DBMS_CLOUD_ADMIN.DROP_DATABASE_LINK(
        db_link_name      IN VARCHAR2,
        public_link       IN BOOLEAN DEFAULT);
```

Parameters

Parameter	Description
`db_link_name`	The name of the database link to drop.
`public_link`	To run `DBMS_CLOUD_ADMIN.DROP_DATABASE_LINK` with `public_link` set to `TRUE`, you must have the `DROP PUBLIC DATABASE LINK` system privilege. The default value for this parameter is `FALSE`.

Examples

```
BEGIN
    DBMS_CLOUD_ADMIN.DROP_DATABASE_LINK(
        db_link_name => 'SALESLINK' );
END;
/


BEGIN
    DBMS_CLOUD_ADMIN.DROP_DATABASE_LINK(
        db_link_name => 'SALESLINK'
        public_link => TRUE );
END;
/
```

Usage Notes

After you are done using a database link and you run `DBMS_CLOUD_ADMIN.DROP_DATABASE_LINK`, to ensure security of your Oracle database remove any stored wallet files. For example:

- Remove the wallet file in Object Store.
- Use `DBMS_CLOUD.DELETE_FILE` to remove the wallet file from the `data_pump_dir` directory or from the user defined directory where the wallet file was uploaded.

ENABLE_APP_CONT Procedure

This procedure enables database application continuity for the session associated with the specified service name in Autonomous Database.

Syntax

```
DBMS_CLOUD_ADMIN.ENABLE_APP_CONT(
    service_name      IN VARCHAR2);
```

Parameters

Parameter	Description
service_name	The service_name for the Autonomous Database service. To find service names: • Look in the tnsnames.ora file in the wallet.zip that you download from an Autonomous Database for your connection. • Click **Database connection** on the Oracle Cloud Infrastructure Console. In the **Connection strings** area, each connection string includes a service_name entry that contains the connection string for the corresponding service. When both Mutual TLS (mTLS) and TLS connections are allowed, under **TLS authentication** select **TLS** to view the TNS names and connection strings for connections with TLS authentication. See View TNS Names and Connection Strings for an Autonomous Database Instance for more information. • Query V$SERVICES view. For example: `SELECT name FROM V$SERVICES;`

Usage Notes

See Overview of Application Continuity for more information on Application Continuity.

Example

```
BEGIN
    DBMS_CLOUD_ADMIN.ENABLE_APP_CONT(
        service_name => 'nvthp2ht_adb1_high.adb.oraclecloud.com'
    );
END;
/
```

Verify the value as follows:

```
SELECT name, failover_type FROM DBA_SERVICES;

NAME                                                  FAILOVER_TYPE
----------------------------------------------------- -------------
nvthp2ht_adb1_high.adb.oraclecloud.com                TRANSACTION
```

ENABLE_AWS_ARN Procedure

This procedure enables an Autonomous Database instance to use Amazon Resource Names (ARNs) to access AWS resources.

Syntax

```
DBMS_CLOUD_ADMIN.ENABLE_AWS_ARN(
       username     IN VARCHAR2 DEFAULT NULL,
       grant_option IN BOOLEAN DEFAULT FALSE);
```

Parameters

Parameter	Description
username	Name of the user to enable to use Amazon Resource Names (ARNs). A null value is valid for the username. If username is not specified, the procedure applies for the "ADMIN" user.
grant_option	When username is supplied, if grant_option is TRUE the specified username can enable Amazon Resource Names (ARNs) usage for other users.

Example

```
BEGIN
    DBMS_CLOUD_ADMIN.ENABLE_AWS_ARN(
        username => 'adb_user');
END;
/
```

Usage Note

- You must be the ADMIN user to run the procedure DBMS_CLOUD_ADMIN.ENABLE_AWS_ARN.

See Use Amazon Resource Names (ARNs) to Access AWS Resources for more information.

ENABLE_EXTERNAL_AUTHENTICATION Procedure

Enable users to login to the database with external authentication schemes.

Syntax

```
DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
   type         IN VARCHAR2,
   force        IN BOOLEAN DEFAULT FALSE,
   params       IN CLOB DEFAULT NULL
);
```

Parameter

Parameter	Description
type	Specifies the external authentication type. Valid values: or . • 'OCI_IAM' • 'AZURE_AD' • 'CMU' • 'KERBEROS'
force	(Optional) Override a currently enabled external authentication scheme. Valid values are TRUE or FALSE. The default value is FALSE.

Parameter	Description
params	A JSON string that provides additional parameters for the external authentication. CMU parameters: • location_uri: specifies the cloud storage URI for the bucket where files required for CMU are stored. If you specify location_uri there is a fixed name directory object CMU_WALLET_DIR created in the database at the path 'cmu_wallet' to save the CMU configuration files. In this case, you do not need to supply the directory_name parameter. • credential_name: specifies the credentials that are used to download the CMU configuration files from the Object Store to the directory object. Default value is NULL which allows you to provide a Public, Preauthenticated, or pre-signed URL for Object Store bucket or subfolder. • directory_name: specifies the directory name where configuration files required for CMU are stored. If directory_name is supplied, you are expected to copy the CMU configuration files dsi.ora and cwallet.sso to this directory object. KERBEROS parameters: • location_uri: specifies the cloud storage URI for the bucket where the files required for Kerberos are stored. If you specify location_uri there is a fixed name directory object KERBEROS_DIR created in the database to save the Kerberos configuration files. In this case, you do not need to supply the directory_name parameter. • credential_name: specifies the credential that are used to download Kerberos configuration files from the Object Store location to the directory object. Default value is NULL which allows you to provide a Public, Preauthenticated, or pre-signed URL for Object Store bucket or subfolder. • directory_name: specifies the directory name where files required for Kerberos are stored. If directory_name is supplied, you are expected to copy the Kerberos configuration files to this directory object. • kerberos_service_name: specifies a name to use as the Kerberos service name. This parameter is optional. Default value: When not specified, the kerberos_service_name value is set to the Autonomous Database instance's GUID. AZURE_AD parameters: • tenant_id: Tenant ID of the Azure Account. Tenant Id specifies the Autonomous Database instance's Azure AD application registration. • application_id: Azure Application ID created in Azure AD to assign roles/schema mappings for external authentication in the Autonomous Database instance. • application_id_uri: Unique URI assigned to the Azure Application. This it the identifier for the Autonomous Database instance. The name must be domain qualified (this supports cross tenancy resource access). The maximum length for this parameter is 256 characters.

Exceptions

Exception	Error	Description
invalid_ext_auth	ORA-20004	See the accompanying message for a detailed explanation.

Usage Notes

- With `type OCI_IAM`, if the resource principal is not enabled on the Autonomous Database instance, this routine enables resource principal with `DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL`.
- This procedure sets the system parameters `IDENTITY_PROVIDER_TYPE` and `IDENTITY_PROVIDER_CONFIG` to required users to access the instance with Oracle Cloud Infrastructure Identity and Access Management authentication and authorization.

Examples

Enable `OCI_IAM` Authentication

```
BEGIN DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
     type => 'OCI_IAM',
     force=> TRUE );
END;
/

PL/SQL procedure successfully completed.
```

Enable `CMU` Authentication for Microsoft Active Directory

You pass in a directory name that contains the CMU configuration files through `params` JSON argument.

```
BEGIN DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
     type => 'CMU',
     force => TRUE,
     params => JSON_OBJECT('directory_name' value 'CMU_DIR'); // CMU_DIR
directory object already exists
END;
/

PL/SQL procedure successfully completed.
```

You pass in a location URI pointing to an Object Storage location that contains CMU configuration files through `params` JSON argument.

```
BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
       type     => 'CMU',
       params   => JSON_OBJECT('location_uri' value 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o',
                               'credential_name' value 'my_credential_name')
   );
```

```
END;
/

PL/SQL procedure successfully completed.
```

Enable Azure AD Authentication

```
BEGIN DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
     type => 'AZURE_AD',
     force => TRUE,
     params   => JSON_OBJECT( 'tenant_id' VALUE '....',
                              'application_id' VALUE '...',
                              'application_id_uri' VALUE '.....' ));
END;
/

PL/SQL procedure successfully completed.
```

Enable Kerberos Authentication

You pass in a directory name that contains Kerberos configuration files through `params` JSON argument.

```
BEGIN DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
     type => 'KERBEROS',
     force => TRUE,
     params => JSON_OBJECT('directory_name' value 'KERBEROS_DIR'); //
KERBEROS_DIR directory object already exists
END;
/

PL/SQL procedure successfully completed.
```

You pass in a location URI pointing to an Object Storage location that contains Kerberos configuration files through `params` JSON argument:

```
BEGIN DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
     type => 'KERBEROS',
     force => TRUE,
     params => JSON_OBJECT('location_uri' value 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o',
                           'credential_name' value 'my_credential_name');
END;
/
```

You pass in a service name with the `kerberos_service_name` in the `params` JSON argument:

```
BEGIN DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION(
     type => 'KERBEROS',
     force => TRUE,
     params => JSON_OBJECT('directory_name' value 'KERBEROS_DIR', //
KERBEROS_DIR directory object already exists
                           'kerberos_service_name' value 'oracle' ));
```

```
END;
/
```

After Kerberos is enabled on your Autonomous Database instance, use the following query to view the Kerberos service name:

```
SELECT SYS_CONTEXT('USERENV','KERBEROS_SERVICE_NAME') FROM DUAL;
```

ENABLE_FEATURE Procedure

This procedure enables the specified feature on the Autonomous Database instance.

Syntax

```
DBMS_CLOUD_ADMIN.ENABLE_FEATURE(
    feature_name     IN VARCHAR2,
    params           IN  CLOB    DEFAULT NULL);
```

Parameters

Parameter	Description
feature_name	Name of the feature to enable. The supported values are: • 'AUTO_DST_UPGRADE': Enable AUTO DST feature. • 'AUTO_DST_UPGRADE_EXCL_DATA': Enable AUTO DST EXCL DATA feature. • 'JAVAVM': Enable JAVAVM feature. • 'OLAP': Enable OLAP feature. OLAP requires that Java is enabled. When you enable OLAP, Java is automatically enabled along with the OLAP feature. • 'ORAMTS': Enable OraMTS feature. • 'OWM': Enable Oracle Workspace Manager. • 'WORKLOAD_AUTO_REPLAY': Enable the workload auto replay feature. This parameter is mandatory.

Parameter	Description
params	A JSON string that provides additional parameters for some features. For the OraMTS feature the params parameter is: • location_uri: the location_uri accepts a string value. The value specifies the HTTPS URL for the OraMTS server in a customer network. For the WORKLOAD_AUTO_REPLAY feature the params parameters are: • target_db_ocid: A string value. The value specifies the OCID of a target refreshable clone database on which the captured workload is replayed. The refreshable clone must have the **Early** patch level set. This parameter is mandatory. • capture_duration: A number value. The value specifies the duration in minutes for which the workload is captured on the production database. The value must be in the range between 1 and 720 minutes. This parameter is mandatory. • capture_day: A string value. The value specifies the day of the week the workload capture on the production database should begin. This parameter is optional. • capture_time: A value in the HH24:MM format. The value specifies the time of the day the workload capture on the production database should begin. This parameter is optional. By default the workload capture starts when you enable WORKLOAD_AUTO_REPLAY. If the optional capture_day and capture_time are specified, the capture and the replay happen at the specified timestamp. For example, if capture_day is Monday and capture_time is 15:00, the first capture happens at 3PM on the next Monday. The day of week and time are also used to schedule the later replay on the refreshable clone.

Example to Enable JAVAVM Feature:

```
BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_FEATURE (
        feature_name => 'JAVAVM' );
END;
/
```

Example to Enable Auto DST Feature:

```
BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_FEATURE (
        feature_name => 'AUTO_DST_UPGRADE' );
END;
/
```

Example to Enable Auto DST EXCL Data Feature:

```
BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_FEATURE (
        feature_name => 'AUTO_DST_UPGRADE_EXCL_DATA' );
END;
/
```

Example to Enable OraMTS Feature:

```
BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_FEATURE(
        feature_name => 'ORAMTS',
        params       => JSON_OBJECT('location_uri' VALUE 'https://
mymtsserver.mycorp.com')
   );
END;
/
```

Example to Enable OWM Feature:

```
BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_FEATURE(
        feature_name => 'OWM' );
END;
/
```

Example to Enable Workload Auto Replay Feature

```
BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_FEATURE(
        feature_name => 'WORKLOAD_AUTO_REPLAY',
        params       => JSON_OBJECT('target_db_ocid' VALUE
'OCID1.autonomousdatabase.REGION..ID1', 'capture_duration' VALUE 120,
'capture_day' VALUE 'MONDAY', 'capture_time' VALUE '15:00'));
END;
/
```

A an error value of `ORA-20000: Invalid argument for target_db_ocid` could indicate that the OCID you supplied is not a refreshable clone. In this case, you need to supply an OCID with a value for a refreshable clone.

Usage Notes

- You must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN` to run `DBMS_CLOUD_ADMIN.ENABLE_FEATURE`.
- After you run `DBMS_CLOUD_ADMIN.ENABLE_FEATURE` with `feature_name` value `'JAVAVM'`, you must restart the Autonomous Database instance to install Oracle Java.

  After you restart the Autonomous Database instance, the Java installation proceeds for some time after the database is opened.

  See Check Oracle Java Version for details on checking the status of the Java installation.
- After you run `DBMS_CLOUD_ADMIN.ENABLE_FEATURE` with `feature_name` value `'OWM'`, you must restart the Autonomous Database instance to enable Oracle Workspace Manager. Oracle.
- To enable `AUTO_DST_UPGRADE`, `AUTO_DST_UPGRADE_EXCL_DATA`, `ORAMTS`, `JAVAVM`, `OWM`, or `WORKLOAD_AUTO_REPLAY` features for your database, you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.

- By default, both `AUTO_DST_UPGRADE` and `AUTO_DST_UPGRADE_EXCL_DATA` are disabled. You can enable one or the other of these options, but not both.
- After you enable `AUTO_DST_UPGRADE`, the next time you restart, or stop and then start the Autonomous Database instance, the instance upgrades to use the latest available time zone files. After `AUTO_DST_UPGRADE` is enabled, when new time zone files are available, the instance continues to upgrade to the latest available version on every subsequent restart or stop and start, until the feature is disabled.

  Query `dba_cloud_config` to verify that `AUTO_DST_UPGRADE` is enabled.

```
SELECT param_name, param_value FROM dba_cloud_config WHERE
       LOWER(param_name) = 'auto_dst_upgrade';

PARAM_NAME       PARAM_VALUE
---------------- --------------
auto_dst_upgrade enabled
```

- After you enable `AUTO_DST_UPGRADE_EXCL_DATA` the Autonomous Database instance upgrades to use the latest available time zone files. After this feature is enabled, every subsequent maintenance window upgrades the instance to use the latest available time zone version. This feature assures that the time zone files are upgraded for the database (enabling `AUTO_DST_UPGRADE_EXCL_DATA` does not update any affected rows).

  Query `dba_cloud_config` to verify that `AUTO_DST_UPGRADE_EXCL_DATA` is enabled.

```
SELECT param_name, param_value FROM dba_cloud_config WHERE
       LOWER(param_name) = 'auto_dst_upgrade_excl_data';

PARAM_NAME                    PARAM_VALUE
----------------------------  -----------
auto_dst_upgrade_excl_data    enabled
```

ENABLE_OPERATOR_ACCESS Procedure

Oracle Cloud Operations does not access your Autonomous Database instance and access is disallowed by default. When access is required to troubleshoot or mitigate an issue, you can allow a cloud operator access to the database schemas for a limited time.

Syntax

```
DBMS_CLOUD_ADMIN.ENABLE_OPERATOR_ACCESS(
    auth_duration    IN  NUMBER  DEFAULT 1
);
```

Parameters

Parameter	Description
`auth_duration`	Specifies the number of Hours for which Cloud Operator is granted access. Valid values: must be whole numbers in the range of 1 to 24. Default value: 1

Example

```
BEGIN
   DBMS_CLOUD_ADMIN.ENABLE_OPERATOR_ACCESS(
       auth_duration => 12 );
END;
/
```

Usage Notes

- `ORA-20000: Operator access is already enabled` indicates that operator access was already granted. In this case you have two options:
  - Wait for operator access to expire, and then grant operator access again with `DBMS_CLOUD_ADMIN.ENABLE_OPERATOR_ACCESS`
  - Explicitly disable operator access with DISABLE_OPERATOR_ACCESS Procedure.
- `DBMS_CLOUD_ADMIN.ENABLE_OPERATOR_ACCESS` allows access only to the Cloud Operator and does not enable access for any other user. All operations performed by the Cloud Operator are stored in the view `DBA_OPERATOR_ACCESS`. See View Oracle Cloud Infrastructure Operations Actions for more information.
- You allow a cloud operator to access the database schemas by running the procedure `DBMS_CLOUD_ADMIN.ENABLE_OPERATOR_ACCESS`. This means if you file a service request with Oracle Cloud Support and Oracle Cloud Operators need to access your database schemas, you must also enable operator access by running`DBMS_CLOUD_ADMIN.ENABLE_OPERATOR_ACCESS`.

ENABLE_PRINCIPAL_AUTH Procedure

This procedure enables principal authentication on Autonomous Database for the specified provider and applies to the ADMIN user or the specified user.

Syntax

```
DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(
       provider    IN VARCHAR2,
       username    IN VARCHAR2 DEFAULT 'ADMIN',
       params      IN CLOB DEFAULT NULL);
```

Parameters

Parameter	Description
`provider`	Specifies the type of provider. Valid values: • `AWS`: Enable use of Amazon Resource Names (ARNs) • `AZURE`: Enable use of Azure Service Principal • `GCP`: Enable use of Google Service Account • `OCI`: Enable use of Resource Principal
`username`	Name of the user who has principal authentication usage enabled. A null value is valid for the `username`. If `username` is not specified, the procedure applies for the "`ADMIN`" user.

Parameter	Description
params	Specifies the configuration parameters. When the provider parameter is AWS, GCP, or OCI, params is not required. The default value is NULL. grant_option: This parameter is valid for all providers and is a Boolean value TRUE or FALSE. The default is FALSE. When TRUE and a username is specified, the specified user can use ENABLE_PRINCIPAL_AUTH to enable other users. **When the provider parameter is AWS, these options are also valid:** • aws_role_arn: Specifies the AWS role ARN. • external_id_type: Specifies the external ID type. Valid values are: "TENANT_OCID", "DATABASE_OCID", or "COMPARTMENT_OCID". The default value for external_id_type is "DATABASE_OCID". See Perform AWS Management Prerequisites to Use Amazon Resource Names (ARNs) for details on the values for external_id_type. **When the provider parameter is AZURE, this option is also valid:** • azure_tenantid: with the value of the Azure tenant ID.

Usage Notes

- When the provider parameter is AZURE, the params parameter must include the azure_tenantid in the following cases:
  - When DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH is called for the first time.
  - When DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH is called for the first time after DBMS_CLOUD_ADMIN.DISABLE_PRINCIPAL_AUTH is called with the provider parameter AZURE and the username ADMIN.
- When the provider parameter is AWS:
  - After you enable ARN on the Autonomous Database instance by running DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH, the credential named AWS$ARN is available to use with any DBMS_CLOUD API that takes a credential as the input.

Examples

```
BEGIN
     DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(
       provider => 'AZURE',
       username => 'SCOTT',
       params   => JSON_OBJECT('azure_tenantid' value 'azure_tenantid'));
END;
/


BEGIN
     DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(
     provider => 'GCP',
     username => 'SCOTT',
     params => JSON_OBJECT(
         'grant_option' value 'TRUE' ));
```

```
END;
/


BEGIN
     DBMS_CLOUD_ADMIN.ENABLE_PRINCIPAL_AUTH(
     provider => 'AWS',
     username => 'SCOTT',
     params => JSON_OBJECT(
         'aws_role_arn' value 'arn:aws:iam::123456:role/AWS_ROLE_ARN',
          'external_id_type' value 'TENANT_OCID'));
END;
/
```

ENABLE_RESOURCE_PRINCIPAL Procedure

Enable resource principal credential for the database or for the specified schema. This procedure creates the credential `OCI$RESOURCE_PRINCIPAL`.

Syntax

```
DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL(
   username         IN VARCHAR2,
   grant_option     IN BOOLEAN DEFAULT FALSE);
```

Parameter

Parameter	Description
`username`	Specifies an optional user name. The name of the database schema to be granted resource principal access. If you do not supply a `username`, the `username` is set to ADMIN.
`grant_option`	When `username` is supplied, if `grant_option` is `TRUE` the specified `username` can enable resource principal usage for other users.

Exceptions

Exception	Error	Description
`resource principal is already enabled`	`ORA-20031`	If you attempt to enable the resource principal when it is already enabled.

Usage Notes

- You must call `DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL` with the ADMIN `username` or with no arguments before you call `DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL` with a `username` for a database user schema.
- You must set up a dynamic group and policies for the dynamic group before you call `DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL`.

  See the following for more information on policies, creating a dynamic group, and creating rules:

  – Use Resource Principal to Access Oracle Cloud Infrastructure Resources

  - Managing Dynamic Groups
  - Managing Policies

- Enabling the resource principal with `DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL` is one-time operation. You do not need to enable the resource principal again, unless you run `DBMS_CLOUD_ADMIN.DISABLE_RESOURCE_PRINCIPAL` to disable the resource principal.
- Resource principal is not available with refreshable clones.

Example

```
EXEC DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL();

PL/SQL procedure successfully completed.

SQL> select owner, credential_name from dba_credentials where credential_name
= 'OCI$RESOURCE_PRINCIPAL';

OWNER    CREDENTIAL_NAME
-------  ---------------
ADMIN    OCI$RESOURCE_PRINCIPAL
```

FINISH_WORKLOAD_CAPTURE Procedure

This procedure finishes the current workload capture, stops any subsequent workload capture requests to the database, and uploads the capture files to Object Storage.

Example

```
BEGIN
    DBMS_CLOUD_ADMIN.FINISH_WORKLOAD_CAPTURE
END;
/
```

Usage Notes

- To run this procedure you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.
- When you pass the `duration` parameter to `START_WORKLOAD_CAPTURE`, the capture finishes when it reaches the specified time. However, if you call `FINISH_WORKLOAD_CAPTURE`, this stops the workload capture (possibly before the time specified with the duration parameter).

  You can query the `DBA_CAPTURE_REPLAY_STATUS` view to check the finish workload status. See DBA_CAPTURE_REPLAY_STATUS View for more information.

> **Note:**
>
> You must subscribe to the Information event `com.oraclecloud.databaseservice.autonomous.database.information` to be notified about the completion of `FINISH_WORKLOAD_CAPTURE` as well as the Object Storage link to download the capture file. This PAR URL is contained in the `captureDownloadURL` field of the event and is valid for 7 days from the date of generation. See Information Events on Autonomous Database for more information.

PREPARE_REPLAY Procedure

The `PREPARE_REPLAY` procedure prepares the refreshable clone for a replay.

Parameters

Parameter	Description
`capture_name`	Specifies the name of the workload capture. This parameter is mandatory.

Syntax

```
DBMS_CLOUD_ADMIN.PREPARE_REPLAY(
        capture_name IN VARCHAR2);
```

Example

```
BEGIN
    DBMS_CLOUD_ADMIN.PREPARE_REPLAY
      capture_name => 'cap_test1');
END;
/
```

This example prepares the refreshable clone to replay the workload indicated by the `capture_name` parameter, which involves bringing it up to the capture start time and then disconnecting it.

Usage Note

- To run this procedure you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.

PURGE_FLASHBACK_ARCHIVE Procedure

This procedure enables ADMIN users to purge historical data from Flashback Data Archive. You can either purge all historical data from Flashback Data Archive `flashback_archive` or selective data based on timestamps or System Change Number.

Syntax

```
DBMS_CLOUD_ADMIN.PURGE_FLASHBACK_ARCHIVE(
   scope       IN VARCHAR2,
```

```
    before_scn IN INTEGER DEFAULT NULL,
    before_ts  IN TIMESTAMP DEFAULT NULL);
```

Parameter	Description
scope	This specifies the scope to remove data from the flashback data archive. • all implies PURGE ALL;before_scn and before_timestamp must both be NULL. • scn implies PURGE BEFORE SCN;before_scn must be non-NULL and before_timestamp must be NULL. • TIMESTAMP implies PURGE BEFORE timestamp;before_scn must be NULL and before_timestamp must be non-NULL.
before_scn	This specifies the system change number before which all the data is removed from the flashback archive.
before_timestamp	This specifies the timestamp before which all the data is removed from the flashback archive.

Example

```
BEGIN
    DBMS_CLOUD_ADMIN.PURGE_FLASHBACK_ARCHIVE(
           scope => 'ALL'); // Purge all historical data from Flashback Data
Archive flashback_archive
END;
/
```

REPLAY_WORKLOAD Procedure

This procedure initiates a workload replay on your Autonomous Database instance. The overloaded form enables you to replay the capture files from an Autonomous Database instance, on-premises database, or other cloud service databases.

Syntax

```
DBMS_CLOUD_ADMIN.REPLAY_WORKLOAD(
        capture_name                 IN VARCHAR2,
        replay_name                  IN VARCHAR2 DEFAULT NULL,
        capture_source_tenancy_ocid  IN VARCHAR2 DEFAULT NULL,
        capture_source_db_name       IN VARCHAR2 DEFAULT NULL);

DBMS_CLOUD_ADMIN.REPLAY_WORKLOAD(
        location_uri      IN VARCHAR2,
        credential_name   IN VARCHAR2 DEFAULT NULL,
        synchronization   IN BOOLEAN  DEFAULT TRUE,
        process_capture   IN BOOLEAN  DEFAULT TRUE);
```

Parameters

Parameter	Description
CAPTURE_NAME	Specifies the name of the workload capture. This parameter is mandatory.

Parameter	Description
REPLAY_NAME	Specifies the replay name. If you do not supply a REPLAY_NAME value, the REPLAY_NAME is auto-generated with the format *REPLAY_RANDOMNUMBER*, for example, REPLAY_1678329506.
CAPTURE_SOURCE_TENANCY_OCID	Specifies the source tenancy OCID of the workload capture. If you do not supply a CAPTURE_SOURCE_TENANCY_OCID value, the CAPTURE_SOURCE_TENANCY_OCID is set to NULL. This parameter is only mandatory when running the workload capture in a full clone.
CAPTURE_SOURCE_DB_NAME	Specifies the source database name of the workload capture If you do not supply a CAPTURE_SOURCE_DB_NAME value, the CAPTURE_SOURCE_DB_NAME is set to NULL. This parameter is only mandatory when running the workload capture in a full clone.
LOCATION_URI	Specifies URI that points to an Object Storage location that contains the captured files. This parameter is mandatory.
CREDENTIAL_NAME	Specifies the credential to access the object storage bucket. If you do not supply a credential_name value, the database's default credentials are used.
SYNCHRONIZATION	Specifies the synchronization method used during workload replay. • TRUE specifies that the synchronization is based on SCN. • FALSE specifies that the synchronization is based on TIME. If you do not supply a synchronization value, the synchronization is set to TRUE.
PROCESS_CAPTURE	Specifies whether or not you need to specify process_capture value. It can be set to FALSE only when you replay the same workload on the target database repeatedly. If you do not supply a process_capture value, the process_capture is set to TRUE.

Example to replay the workload from an on-premises database on an Autonomous Database instance:

```
BEGIN
  DBMS_CLOUD_ADMIN.REPLAY_WORKLOAD(
        location_uri    => 'https://objectstorage.us-
phoenix-1.oraclecloud.com/n/namespace-string/b/bucketname/o',
        credential_name => 'CRED_TEST',
        synchronization => TRUE,
        process_capture => TRUE);
END;
/
```

When you run this example, it:

- Downloads the capture files from the Object Storage location specified in location_uri and processes the capture files based on the process_capture parameter value.
- Replays the captured workload based on the synchronization parameter value.

In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

See Navigate to Oracle Cloud Infrastructure Object Storage and Create Bucket for more information on Object Storage.

See Upload Files to Your Oracle Cloud Infrastructure Object Store Bucket for more information on uploading files to Object Storage.

The `credential_name` you use in this step is the credentials for the Object Store.

You don't need to create a credential to access Oracle Cloud Infrastructure Object Store if you enable resource principal credentials. See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

Example to replay the workload from an Autonomous Database instance on another Autonomous Database:

```
BEGIN
  DBMS_CLOUD_ADMIN.REPLAY_WORKLOAD(
        capture_name => 'CAP_TEST1');
END;
/
```

When you run this example, it:

- Disconnects the current Autonomous Database instance.
- Downloads the capture files from the Object Storage.
- Replays the captured workload.
- Uploads replay report after a replay.

Usage Notes for Replaying the Workload from an On-Premises or Other Cloud Service Database on another Autonomous Database

- To run this procedure you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.
- Before you start replay, you should upload the `cap` and `capfiles` subdirectories, which contain the workload capture files, to the object storage location.

Usage Notes for Replaying the Workload from an Autonomous Database instance on another Autonomous Database

- To run this procedure you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.
- The replay files are automatically uploaded to the Object Store as a zip file.
- You can query the `DBA_CAPTURE_REPLAY_STATUS` view to check the workload replay status.

  See DBA_CAPTURE_REPLAY_STATUS View for more information.

> **Note:**
>
> You must subscribe to the Information event `com.oraclecloud.databaseservice.autonomous.database.information` to be notified about the start and completion of the `REPLAY_WORKLOAD` as well as the Object Storage link to download the replay reports. This PAR URL is contained in the `replayDownloadURL` field of the event and is valid for 7 days from the date of generation. See Information Events on Autonomous Database for more information.

SET_FLASHBACK_ARCHIVE_RETENTION Procedure

This procedure allows ADMIN users to modify the retention period for Flashback Data Archive `flashback_archive`.

Syntax

```
DBMS_CLOUD_ADMIN.SET_FLASHBACK_ARCHIVE_RETENTION (
    retention_days INTEGER);
```

Parameter	Description
`retention_days`	This specifies the length of time in days that the archived data should be retained for. The value of `retention_days` must be greater than 0.

Example

```
BEGIN
     DBMS_CLOUD_ADMIN.SET_FLASHBACK_ARCHIVE_RETENTION(
           retention_days => 90); // sets the retention time to 90 days
END;
/
```

START_WORKLOAD_CAPTURE Procedure

This procedure initiates a workload capture on your Autonomous Database instance.

Syntax

```
DBMS_CLOUD_ADMIN.START_WORKLOAD_CAPTURE(
        capture_name  IN VARCHAR2,
        duration      IN NUMBER   DEFAULT NULL);
```

Parameters

Parameter	Description
`capture_name`	Specifies the name of the workload capture. This parameter is mandatory.
`duration`	Specifies the duration in minutes for which you want to run the workload capture. • If you do not supply a `duration` value, the `duration` is set to NULL. • If set to NULL, the workload will continue until you run the `FINISH_WORKLOAD_CAPTURE` procedure.

Example

```
BEGIN
  DBMS_CLOUD_ADMIN.START_WORKLOAD_CAPTURE(
        capture_name => 'test');
END;
/
```

Usage Notes

- To run this procedure you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_CLOUD_ADMIN`.
- To measure the impacts of a system change on a workload, you must ensure that the capture and replay systems are in the same logical state.
- Before initiating a workload capture, you should consider provisioning a refreshable clone to ensure the same start point for the replay.

> **Note:**
>
> You must subscribe to the Information event `com.oraclecloud.databaseservice.autonomous.database.information` to be notified at the start of `START_WORKLOAD_CAPTURE`. See Information Events on Autonomous Database for more information.

DBMS_CLOUD_ADMIN Exceptions

The following table describes exceptions for `DBMS_CLOUD_ADMIN`.

Exception	Code	Description
`invalid_service`	20001	An invalid service was specified.
`service_not_exist`	20002	A service specified does not exist.
`default_service`	20003	A service specified cannot be modified.
`invalid_char_set`	20029	Missing precondition or invalid (national) character set.
`invalid_enc_key_attr`	20030	Missing or invalid argument for key management.
`Already Using Oracle Managed Key`	000000	The Database is already using an Oracle managed key. You are trying to call the procedure while already using an Oracle managed key.
`Argument Provided for the procedure`	ORA-0000	An argument is provided for the procedure. Expected error message:No arguments required for this procedure.

DBMS_CLOUD_AI Package

The `DBMS_CLOUD_AI` package, with Select AI, facilitates and configures the translation of natural language prompts to generate, run, explain SQL statements. Also, enables retrieval augmented generation and natural language-based interactions, including chatting with LLMs.

- DBMS_CLOUD_AI Overview
  Describes the use of the `DBMS_CLOUD_AI` package.
- Summary of DBMS_CLOUD_AI Subprograms
  This section covers the `DBMS_CLOUD_AI` subprograms provided with Autonomous Database.

> **See Also:**
>
> Use Select AI for Natural Language Interaction with your Database

DBMS_CLOUD_AI Overview

Describes the use of the `DBMS_CLOUD_AI` package.

Use the `DBMS_CLOUD_AI` package to create AI profiles and configure them for access to a Large Language Model (LLM). Set the AI profile in the current database user session to perform tasks such as generating, running, and explaining SQL. Additionally, enable retrieval augmented generation and natural language-based interactions, including chatting with LLMs.

To know more about supported platforms and LLMs, see About Select AI.

Summary of DBMS_CLOUD_AI Subprograms

This section covers the `DBMS_CLOUD_AI` subprograms provided with Autonomous Database.

Subprogram	Description
CREATE_PROFILE Procedure	This procedure creates a new AI profile for translating natural language prompts to SQL statements.
Profile Attributes	Provides AI profile attributes that you can configure.
DISABLE_PROFILE Procedure	This procedure disables an AI profile in the current database.
DROP_PROFILE Procedure	This procedure drops an existing AI profile.
ENABLE_PROFILE Procedure	This procedure enables an AI profile to use in the current database.
GENERATE Function	This function enables you to use Select AI in a stateless manner with your existing profile.
GENERATE_SYNTHETIC_DATA Function	This function generates synthetic data.
GET_PROFILE Function	This function returns the profile name used in the current session.
GET_PROFILE Procedure	This procedure returns the profile name and the owner of the profile in the current session.
SET_ATTRIBUTE Procedure	This procedure sets AI profile attributes.
SET_PROFILE Procedure	This procedure sets AI profile for the current database.
ENABLE_DATA_ACCESS Procedure	Use this procedure to enable sending data to your LLM.
DISABLE_DATA_ACCESS Procedure	Use this procedure to disable sending data to your LLM.
CREATE_VECTOR_INDEX Procedure	This procedure creates a vector index in the specified vector database, and populates it with data from an object store using an asynchronous scheduler job.
DROP_VECTOR_INDEX Procedure	This procedure removes a vector store index. It normally removes the vector store index object and deletes the vector database.
DISABLE_VECTOR_INDEX Procedure	This procedure disables a vector index object in the current database. When disabled, an AI profile cannot use the vector index, and the system does not load data into the vector store.

Subprogram	Description
ENABLE_VECTOR_INDEX Procedure	This procedure enables or activates a previously disabled vector index object.
UPDATE_VECTOR_INDEX Procedure	This procedure updates an existing vector store index with a specified value of the vector index attribute.
Vector Index Attributes	Provides vector index profile attributes that you can configure.

- CREATE_PROFILE Procedure
  The procedure creates a new AI profile for translating natural language prompts to SQL statement.
- DROP_PROFILE Procedure
  The procedure drops an existing AI profile. If the profile does not exist, then the procedure throws an error.
- ENABLE_PROFILE Procedure
  This procedure enables the AI profile that the user specifies. The procedure changes the status of the AI profile to `ENABLED`.
- DISABLE_PROFILE Procedure
  This procedure disables the AI profile in the current database. The status of the AI profile is changed to `DISABLED` by this procedure.
- GET_PROFILE Function
  This function returns the AI profile name set in the current session.
- GET_PROFILE Procedure
  This procedure returns the AI profile name and the owner set in the current session.
- SET_ATTRIBUTE Procedure
  This procedure enables you to set AI profile attributes. It is overloaded to accept attribute values of various types.
- SET_PROFILE Procedure
  This procedure sets AI profile for current session.
- GENERATE Function
- GENERATE_SYNTHETIC_DATA Function
- Profile Attributes
- ENABLE_DATA_ACCESS Procedure
  This procedure enables sending data to LLM for applicable Select AI features, which is the default behavior. Only an administrator can run this procedure.
- DISABLE_DATA_ACCESS Procedure
  This procedure disables sending data to LLM for applicable Select AI features. Only an administrator can run this procedure.
- CREATE_VECTOR_INDEX Procedure
  This procedure creates a vector index in the specified vector database, and populates it with data from an object store using an asynchronous scheduler job.
- DROP_VECTOR_INDEX Procedure
  This procedure removes a vector store index. It normally removes the vector store index object and deletes the vector store. If set to `FALSE`, the argument `include_data` ensures the procedure only removes the vector store index object while retaining the vector store.

- DISABLE_VECTOR_INDEX Procedure
  This procedure disables a vector index object in the current database. When disabled, an AI profile cannot use the vector index, and the system does not load data into the vector store as new data is added to the object store and does not perform indexing, searching or querying based on the index.
- ENABLE_VECTOR_INDEX Procedure
  This procedure enables or activates a previously disabled vector index object. Generally, when you create a vector index, by default it is enabled such that the AI profile can use it to perform indexing and searching.
- UPDATE_VECTOR_INDEX Procedure
  This procedure updates an existing vector store index with a specified value of the vector index attribute. It is overloaded to accept attribute values of various types.
- Vector Index Attributes

CREATE_PROFILE Procedure

The procedure creates a new AI profile for translating natural language prompts to SQL statement.

Syntax

```
DBMS_CLOUD_AI.CREATE_PROFILE
   profile_name        IN  VARCHAR2,
   attributes          IN  CLOB      DEFAULT NULL,
   status              IN  VARCHAR2  DEFAULT NULL,
   description         IN  CLOB      DEFAULT NULL
);
```

Parameters

Parameter	Description
profile_name	A name for the AI profile. The profile name must follow the naming rules of Oracle SQL identifier. Maximum length of profile name is 125 characters. This is a mandatory parameter.
attributes	Profile attributes in JSON format. See AI Profile Attributes for more details. The default value is NULL.
status	Status of the profile. The default value is enabled.
description	Description for the AI profile. The default value is NULL.

Example

```
BEGIN
     DBMS_CLOUD_AI.CREATE_PROFILE(
          profile_name    => 'OpenAI',
          attributes      => JSON_OBJECT('provider' value 'openai',
                                         'credential_name' value
'openai_cred'),
```

```
        status     => 'enabled',
          description     => 'AI profile to use OpenAI for SQL translation'
    );
END;
/
```

Related Topics

- DBA_CLOUD_AI_PROFILE_ATTRIBUTES View
- USER_CLOUD_AI_PROFILE_ATTRIBUTES View

DROP_PROFILE Procedure

The procedure drops an existing AI profile. If the profile does not exist, then the procedure throws an error.

Syntax

```
DBMS_CLOUD_AI.DROP_PROFILE(
       profile_name        IN   VARCHAR2,
       force               IN   BOOLEAN DEFAULT FALSE
 );
```

Parameters

Parameter	Description
`profile_name`	Name of the AI profile
`force`	If `TRUE`, then the procedure ignores errors if AI profile does not exist. The default value for this parameter is `FALSE`.

Example

```
BEGIN
     DBMS_CLOUD_AI.DROP_PROFILE(profile_name => 'OPENAI');
END;
/
```

Usage Notes

Use `force` to drop a profile and ignore errors if AI profile does not exist.

ENABLE_PROFILE Procedure

This procedure enables the AI profile that the user specifies. The procedure changes the status of the AI profile to `ENABLED`.

Syntax

```
DBMS_CLOUD_AI.ENABLE_PROFILE(
     profile_name         IN   VARCHAR2
 );
```

Parameters

Parameter	Description
profile_name	Name for the AI profile to enable This parameter is mandatory.

Example to Enable AI Profile

```
BEGIN
     DBMS_CLOUD_AI.ENABLE_PROFILE(
         profile_name    => 'OPENAI'
     );
END;
/
```

DISABLE_PROFILE Procedure

This procedure disables the AI profile in the current database. The status of the AI profile is changed to `DISABLED` by this procedure.

Syntax

```
DBMS_CLOUD_AI.DISABLE_PROFILE(
      profile_name  IN  VARCHAR2
);
```

Parameters

Parameter	Description
profile_name	Name for the AI profile. This parameter is mandatory.

Example

```
BEGIN
     DBMS_CLOUD_AI.DISABLE_PROFILE(
         profile_name    => 'OPENAI'
     );
END;
/
```

GET_PROFILE Function

This function returns the AI profile name set in the current session.

Syntax

```
DBMS_CLOUD_AI.GET_PROFILE(
    profile_name      IN  VARCHAR2
);
```

Parameters

Parameter	Description
profile_name	A name for the AI profile in the current session. This parameter is mandatory.

Example

This example shows how you can display the name of the profile in the current session.

```
SELECT DBMS_CLOUD_AI.GET_PROFILE
from DUAL;
```

GET_PROFILE Procedure

This procedure returns the AI profile name and the owner set in the current session.

Syntax

```
DBMS_CLOUD_AI.GET_PROFILE(

      profile_name  OUT  VARCHAR2,
      profile_owner OUT  VARCHAR2
);
```

Parameters

Parameter	Description
profile_name	A name for the AI profile in the current session. This parameter is mandatory.
profile_owner	Identifies the owner of the AI profile in the current session.

Example

This example shows how you can display the name and owner of the profile in the current session.

```
DECLARE
       l_profile_name  DBMS_ID;
       l_profile_owner DBMS_ID;
     BEGIN
       DBMS_CLOUD_AI.GET_PROFILE(profile_name => l_profile_name,
```

```
                                        profile_owner => l_profile_owner);
    END;
```

Related Topics

- DBA_CLOUD_AI_PROFILE_ATTRIBUTES View
- USER_CLOUD_AI_PROFILE_ATTRIBUTES View

SET_ATTRIBUTE Procedure

This procedure enables you to set AI profile attributes. It is overloaded to accept attribute values of various types.

Syntax

```
DBMS_CLOUD_AI.SET_ATTRIBUTE(
      profile_name         IN   VARCHAR2,
      attribute_name       IN   VARCHAR2,
      attribute_value      IN   {BOOLEAN|VARCHAR2}
);

DBMS_CLOUD_AI.SET_ATTRIBUTE(
      profile_name         IN   VARCHAR2,
      attribute_name       IN   VARCHAR2,
      attribute_value      IN   CLOB DEFAULT NULL
);
```

Parameters

Only the owner can set or modify the attributes of the AI profile. For a list of supported attributes, see Profile Attributes.

Parameter	Description
`profile_name`	Name of the AI profile for which you want to set the attributes. This parameter is mandatory.
`attribute_name`	Name of the AI profile attribute This parameter is mandatory.
`attribute_value`	Value of the profile attribute. Value can be `BOOLEAN`, `CLOB`, `NUMBER` or `VARCHAR2` type. The default value is NULL.

Examples

```
BEGIN
 DBMS_CLOUD_AI.SET_ATTRIBUTE(
   profile_name    => 'OPENAI',
   attribute_name  => 'credential_name',
   attribute_value => 'OPENAI_CRED_NEW'
 );
END;
/
```

The following example accepts `NUMBER` type as the `attribute_value`.

```
BEGIN
 DBMS_CLOUD_AI.SET_ATTRIBUTE(
   profile_name    => 'OCI_PROFILE',
   attribute_name  => 'temperature',
   attribute_value => 0.5
 );
END;
/
```

The following example accepts `BOOLEAN` type as the `attribute_value`.

```
BEGIN
 DBMS_CLOUD_AI.SET_ATTRIBUTE(
   profile_name    => 'OCI_PROFILE',
   attribute_name  => 'comments',
   attribute_value => 'true'
 );
END;
/
```

The following example accepts `VARCHAR2` type as the `attribute_value`.

```
BEGIN
 DBMS_CLOUD_AI.SET_ATTRIBUTE(
   profile_name    => 'OCI_PROFILE',
   attribute_name  => 'model',
   attribute_value => 'meta.llama-3.3-70b-instruct'
 );
END;
/
```

SET_PROFILE Procedure

This procedure sets AI profile for current session.

After setting an AI profile for the database session, any SQL statement with the prefix `SELECT AI` is considered a natural language prompt. Depending on the action the you specify with the `AI` prefix, a response is generated using AI. To use the `AI` prefix, see Examples of Using Select AI and Use AI Keyword to Enter Prompts. Optionally, it is possible to override the profile attributes or modify attributes by specifying them in JSON format. See SET_ATTRIBUTE Procedure for setting the attributes.

The AI profile can only be set for current session if the owner of the AI profile is the session user.

To set an AI profile for all sessions of a specific database user or all user sessions in the database, consider using a database event trigger for `AFTER LOGON` event on the specific user or the entire database. See CREATE TRIGGER Statement for more details.

Syntax

```
DBMS_CLOUD_AI.SET_PROFILE(
    profile_name      IN  VARCHAR2,
);
```

Parameters

Parameter	Description
`profile_name`	A name for the AI profile in the current session. This parameter is mandatory.

Example

```
BEGIN
     DBMS_CLOUD_AI.SET_PROFILE(
       profile_name    => 'OPENAI'
     );
END;
/
```

GENERATE Function

This function provides AI translation in a stateless manner. With your existing AI profile, you can use this function to perform the supported actions such as `showsql`, `runsql`, `explainsql`, `narrate`, and `chat`. The default action is `showsql`.
Overriding some or all of the profile attributes is also possible using this function.

Syntax

```
DBMS_CLOUD_AI.GENERATE(
    prompt          IN  CLOB,
    profile_name    IN  VARCHAR2 DEFAULT NULL,
    action          IN  VARCHAR2 DEFAULT NULL,
    attributes      IN  CLOB     DEFAULT NULL
) RETURN CLOB;
```

Parameters

Parameter	Description
`prompt`	Natural language prompt to translate using AI. The prompt can include `SELECT AI <action>` as the prefix. The action can also be supplied separately as an "action" parameter. The *action* supplied in prompt overrides the "action" parameter. Default action is `showsql`. This parameter is mandatory.

Parameter	Description
profile_name	Name of the AI profile. This parameter is optional if an AI profile is already set in the session using DBMS_CLOUD_AI.SET_PROFILE. The default value is NULL. The following conditions apply: • If a profile is set in the current session, the user may omit profile_name argument in the DBMS_CLOUD_AI.GENERATE function. • If the profile_name argument is supplied in the DBMS_CLOUD_AI.GENERATE function, it overrides any value set in the session using the DBMS_CLOUD_AI.SET_PROFILE procedure. • If there is no profile set in the session using the DBMS_CLOUD_AI.SET_PROFILE procedure, the profile_name argument must be supplied in the DBMS_CLOUD_AI.GENERATE function.

Note:

For Database Actions, you can either specify profile_name argument in DBMS_CLOUD_AI.GENERATE or you can run two steps as a PL/SQL script: DBMS_CLOUD_AI.SET_PROFILE and DBMS_CLOUD_AI.GENERATE.

```
EXEC
DBMS_CLOUD_AI.set_profile('OPENAI')
;

-----------------------------------
------------
SELECT
DBMS_CLOUD_AI.GENERATE(prompt
 => 'how many customers',

profile_name => 'OPENAI',

action       => 'runsql')
FROM dual;
-----------------------------------
------------

SELECT
DBMS_CLOUD_AI.GENERATE(prompt
 => 'how many customers',

profile_name => 'OPENAI',

action       => 'showsql')
FROM dual;
-----------------------------------
------------
SELECT
```

Parameter	Description
	DBMS_CLOUD_AI.GENERATE(prompt => 'how many customers', profile_name => 'OPENAI', action => 'explainsql') FROM dual; ---------- ---------- SELECT DBMS_CLOUD_AI.GENERATE(prompt => 'how many customers', profile_name => 'OPENAI', action => 'narrate') FROM dual; ---------- -------- SELECT DBMS_CLOUD_AI.GENERATE(prompt => 'what is oracle autonomous database', profile_name => 'OPENAI', action => 'chat') FROM dual; See Executing SQL Statements in the Code Editor for more information.
action	Action for translating natural prompt using AI. The supported actions include runsql (default), showsql, explainsql, narrate, and chat. Descriptions of actions are included in Use AI Keyword to Enter Prompts.
attributes	Override specific AI profile attributes by supplying attributes in JSON format. See Profile Attributes for more details.

Examples

The following examples illustrate runsql, showsql, explainsql, narrate, and chat actions that can be used with the DBMS_CLOUD_AI.GENERATE function. See also Use AI Keyword to Enter Prompts for more details.

An example with runsql action is as follows:

```
SELECT DBMS_CLOUD_AI.GENERATE(prompt       => 'how many customers',
                              profile_name => 'OPENAI',
                              action       => 'runsql')
FROM dual;
```

An example with `showsql` action is as follows:

```
SELECT DBMS_CLOUD_AI.GENERATE(prompt       => 'how many customers',
                              profile_name => 'OPENAI',
                              action       => 'showsql')
FROM dual;
```

An example with `explainsql` action is as follows:

```
SELECT DBMS_CLOUD_AI.GENERATE(prompt       => 'how many customers',
                              profile_name => 'OPENAI',
                              action       => 'explainsql)
FROM dual;
```

An example with `narrate` action is as follows:

```
SELECT DBMS_CLOUD_AI.GENERATE(prompt       => 'how many customers',
                              profile_name => 'OPENAI',
                              action       => 'narrate')
FROM dual;
```

An example with `chat` action is as follows:

```
SELECT DBMS_CLOUD_AI.GENERATE(prompt       => 'what is oracle autonomous
database',
                              profile_name => 'OPENAI',
                              action       => 'chat')
FROM dual;
```

You can use `DBMS_CLOUD_AI.GENERATE` in a procedure and run the function. The following example takes an `ai_prompt`, `profile_name`, and `action` as input parameters and calls `DBMS_CLOUD_AI.GENERATE`

```
create or replace FUNCTION call_select_ai (ai_prompt  IN VARCHAR2,
                                           ai_profile IN VARCHAR2,
                                           ai_action  IN VARCHAR2) -- valid
for 'chat', 'narrate', 'showsql'
                                           RETURN CLOB AS sai_resp clob;
BEGIN
  sai_resp := DBMS_CLOUD_AI.GENERATE(prompt       => ai_prompt,
                                     profile_name => ai_profile,
                                     action       => ai_action);
  return(sai_resp);
END call_select_ai;
```

GENERATE_SYNTHETIC_DATA Function

Use this procedure to generate synthetic data for a single table, multiple tables or a full schema.
The following is the syntax to generate synthetic data for a single table.

Syntax

```
DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA(
  profile_name         IN  VARCHAR2,
  object_name          IN  DBMS_ID,
  owner_name           IN  DBMS_ID,
  record_count         IN  NUMBER,
  user_prompt          IN  CLOB DEFAULT NULL,
  params               IN  CLOB DEFAULT NULL
);
```

The following is the syntax to generate synthetic data for multiple tables.

```
DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA(
  profile_name         IN  VARCHAR2,
  object_list          IN  CLOB,
  params               IN  CLOB DEFAULT NULL
);
```

If you do not want table data or vector search documents to be sent to an LLM, a user with administrator privileges can disable such access for all users of the given database. This, in effect, disables the `narrate` action.

Parameters

Parameter	Mandatory	Description
`profile_name`	Yes	The AI profile containing necessary LLM service information. This can be created by CREATE_PROFILE Procedure.
`object_name`	Yes	Specify a table name to populate synthetic data. • `SELECT` and `INSERT` privilege on the table objects are needed for the user using it. • The table is either empty or have records in it.
`owner_name`	No	Specify the database user who owns the referenced object. If no specific owner is provided, the procedure defaults to using the schema of the user running it.
`record_count`	No	The number of records to be synthetically generated.
`user_prompt`	No	Additional information that a user can mention to generate synthetic data. For example, to generate a record for a table called `MOVIE` with a `release_date` column, the `user_prompt` can be: *the release date for the movies should be in 2019*
`params`	No	Optional attributes provided in JSON object string format to modify the behavior of an API. See Optional Parameters.

Parameter	Mandatory	Description
`object_list`	Yes	Use this parameter for generating synthetic data on multiple tables. This parameter takes in table object information along with it's arguments and contains the same arguments provided in the single table. See object_list Parameters.

Optional Parameters

Parameter	Value Datatype	Value	Description
`sample_rows`	Number	`0 <= sample_rows <= 100`	Specify the number of rows from the table to use as a sample to guide the LLM in data generation. A value of 0 means no sample rows will be used. The default value is `0`.
`table_statistics`	Boolean	• `True` • `False`	Enable or disable the use of table statistics information. The default value is `False`.

Parameter	Value Datatype	Value	Description
`priority`	String	Valid values: • `HIGH` • `MEDIUM` • `LOW`	Assign a priority value that defines the number of parallel requests sent to the LLM for generating synthetic data. Tasks with a higher priority will consume more database resources and complete faster. The default value is `HIGH` • `HIGH`: Specifies the number of parallel LLM requests based on the database's ECPU count (or OCPU count if your database uses OCPUs). • `MEDIUM`: Sets the number of concurrent processes according to the concurrency limit for Medium service. The default value is 4. • `LOW`: Runs the pipeline job in a serial order, without parallel processing. The maximum number of concurrent parallel processes used for synthetic data generation is limited to 64.
`comments`	Boolean	• `True` • `False`	Enable or disable sending comments to the LLM to guide data generation. The default value is `False`.

object_list Parameters

Parameter	Value Datatype	Mandatory	Description
`owner`	String	Yes	Specifies the database user who owns the object being referenced. If no specific owner is provided, the procedure will default to using the schema of the user running it.
`name`	String	No	Specify a table name to populate synthetic data. `SELECT` and `INSERT` privilege on the table objects are needed for the user using it.The table is either empty or have records in it.
`record_count`	Number	No	The number of records to be synthetically generated. Provide a number greater than 0. Supply `record_count` or `record_count_percentage`.

Parameter	Value Datatype	Mandatory	Description
record_count_percentage	Number	No	The percentage of number of records to be synthetically generated. Provide a number greater than 0. For a Metadata Clone database, where the table metadata including statistics is preserved, the record_count_percentage parameter is supported. Supply record_count or record_count_percentage. When using the record_count_percentage parameter, the final record count in the table is calculated as: Original_Num_Rows * record_count_percentage
user_prompt	String	No	Same as user_prompt in Parameters. The user_prompt is associated with a specific table object.

Examples

The following examples show the DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA function for generating synthetic data for a single table and multiple tables. For a complete example and to view more examples, see Example: Generate Synthetic Data.

```
BEGIN
    DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA(
        profile_name => 'GENAI',
        object_name  => 'Director',
        owner_name   => 'ADB_USER',
        record_count => 5
    );
END;
```

```
/
PL/SQL procedure successfully completed.


BEGIN
    DBMS_CLOUD_AI.GENERATE_SYNTHETIC_DATA(
        profile_name => 'GENAI',
        object_list => '[{"owner": "ADB_USER", "name":
"Director","record_count":5},
                        {"owner": "ADB_USER", "name":
"Movie_Actor","record_count":5},
                        {"owner": "ADB_USER", "name":
"Actor","record_count":10},
                        {"owner": "ADB_USER", "name":
"Movie","record_count":5,"user_prompt":"all movies are released in 2009"}]'
    );
END;
/
PL/SQL procedure successfully completed.
```

Profile Attributes

Attributes of an AI profile help to manage and configure the behavior of the AI profile. Some attributes are optional and have a default value.

Attributes

Attribute Name	Description
`annotations`	Provides additional metadata about your database tables and columns using the 23ai annotations feature. The valid values are `TRUE` or `FALSE`. The default value is `FALSE`. The values are not case sensitive. `TRUE`: Augments the user prompt with table-level and column-level annotations in addition to other metadata to the LLM. `FALSE`: Does not include annotations as part of the augmented prompt metadata.
`azure_deployment_name`	Name of the Azure OpenAI Service deployed model. The name can only include alphanumeric characters, underscore character (_) and a hyphen (-) character. The name cannot end with an underscore (_) or a hyphen (-). To know how to get the `azure_deployment_name`, see Create and deploy an Azure OpenAI Service resource.
`azure_embedding_deployment_name`	Name of the Azure OpenAI deployed embedding model. The name can only include alphanumeric characters, underscore, and hyphen. The name can't start or end with a hyphen or underscore.
`azure_resource_name`	Name of the Azure OpenAI Service resource. The resource name can only include alphanumeric characters and hyphens, and can't start or end with a hyphen. To know how to get the `azure_resource_name`, see Create and deploy an Azure OpenAI Service resource.

Attribute Name	Description
case_sensitive_values	Specifies whether Select AI should instruct the LLM to generate case-insensitive SQL for columns containing string or text values. The supported values are: • true: Indicates the LLM to generate SQL queries with case-sensitive string values. • false: Indicates the LLM to generate SQL queries that ignore case for string values. The default value is false.
comments	Includes table-level comments and column comments in the metadata used for translating natural language prompts using AI. BOOLEAN data type is supported. The valid values are TRUE or FALSE for a string with VARCHAR2 data type. The values are not case sensitive. **Note:** For DBMS_CLOUD_AI.SET_ATTRIBUTE procedure, when setting a single attribute, supply the attribute_value parameter as a string ('true' or 'false') because the attribute_value parameter is of CLOB data type.
constraints	Manages whether to include referential integrity constraints such as primary and foreign keys in the metadata sent to the LLM. The valid values are TRUE or FALSE. The default value is FALSE. The values are not case sensitive. TRUE: Includes referential constraints in the augmented prompt metadata that is sent to the LLM. FALSE: Does not include referential integrity constraints.
conversation	A VARCHAR2 attribute that indicates if conversation history is enabled for a profile. Valid values are true or false. The default value is false. The values are not case sensitive.
credential_name	The name of the credential to access the AI provider APIs. Credential using bearer tokens can be created by using the provider name as the user name and bearer token as the password. Vault Secret credentials are also supported. Principal authentication, for example, Azure service principal, is also supported. For more information on how to configure it, see Use Azure Service Principal to Access Azure Resources. This is a mandatory attribute. See CREATE_CREDENTIAL Procedure.

Attribute Name	Description
`embedding_model`	The embedding model defined in the AI profile. The following are the supported AI providers for the embedding models with their default values: • OCI GenAI: `cohere.embed-english-v3.0` • OpenAI: `text-embedding-ada-002` • Azure OpenAI: `text-embedding-ada-002` • Cohere: `embed-english-v2.0` • Google: `text-embedding-004` **Note:** • The `embedding_model` parameter is not applicable for Anthropic and Hugging Face. • For in-database models, specify the model name imported into your database if you want to use Select AI with in-database transformer model. You can get the model name by querying `USER_MINING_MODELS`. **Syntax:** The following is the syntax for using in-database transformer model for `embedding_model`: `"embedding_model": "[<model_schema>.]<model_name>"`
`enforce_object_list`	Specifies whether to restrict the LLM to generate SQL that uses only tables covered by the object list. The supported values are: • `true`: Enforces the use of only the tables specified in the object list. • `false`: Allows the LLM to use tables accessible to the user. The values are not case sensitive. The default value is `false`.
`max_tokens`	Denotes the number of tokens to predict per generation. Default is 1024. See Tokens and Tokenizers for more details.

Attribute Name	Description
`model`	The name of the AI model being used to generate responses. See Select your AI Provider and LLMs and specify the model name. **Note:** • **Cohere**: Smaller, "light" models are faster, while larger models perform better. Custom models can also be supplied with their full ID. • **OCI Generative AI**: The Chat Models are supported for all Select AI actions such as `runsql`, `showsql`, `explainsql`, `narrate`, and `chat`. Select AI supports pretrained models for OCI Generative AI. Custom models can also be supplied with their full OCIDs. If you are supplying OCID or `oci_endpoint_id`, be sure to provide either `oci_runtimetype` or `oci_apiformat` depending on the OCI Chat models. To know more about supported models in OCI Generative AI, see Pretrained Foundational Models in Generative AI. • This parameter is not used for Azure as the model is determined when you create your deployment in the Azure OpenAI Service portal.

Attribute Name	Description
object_list	Array of JSON objects specifying the owner and object names that are eligible for natural language translation to SQL. To include all objects of a given user, omit the "name" and only specify the "owner" key in the JSON object. The following types of objects can be used: • tables • views • materialized views • global temporary tables • external tables • synonyms on the above object types For translation natural language to SQL, the object name, object owner, object columns and comments are sent to the AI provider using HTTPS requests. Avoid specifying objects with sensitive object name, column names or comments in the object list. AI providers may have limit on the size of metadata allowed in translation requests. Consider limiting the list of objects suitable for the natural language prompts by your application users. **Format:** `[ {"owner": "SH", "name": "SALES", {"owner": "TEST_USER"} ]` External tables created using sync of OCI Data Catalog or AWS Glue can also be used the object list. This helps in managing metadata in central Data Catalogs and use the metadata directly for translating natural language prompts using AI.
object_list_mode	Specifies whether to send metadata for the most relevant tables or all tables to the LLM. Supported values are: • automated: Sends metadata only for the specific tables identified as most relevant to the query. • all: Sends metadata for all tables accessible to the user. When object_list_mode is set to automated, Select AI automatically creates a vector index, named <profile_name>_OBJECT_LIST_VECINDEX with default attributes. See

Attribute Name	Description
oci_apiformat	Specifies the format in which the API expects data to be sent and received. Use this attribute to generate text responses. This attribute applies to OCI Generative AI Chat Models in a dedicated AI cluster. Specify this attribute when you specify a model OCID in the model attribute or provide an endpoint in the oci_endpoint_id attribute. Supported values are: • COHERE • GENERIC **Note:** Use this attribute for OCI Generative AI Chat Models
oci_compartment_id	Specifies the OCID of the compartment you are permitted to access when calling the OCI Generative AI service. The compartment ID can contain alphanumeric characters, hyphens and dots. The default is the compartment ID of the PDB.
oci_endpoint_id	This attributes indicates the endpoint OCID of the Oracle dedicated AI hosting cluster. The endpoint ID can contain alphanumeric characters, hyphens and dots. To find the endpoint OCID, see Getting an Endpoint's Details in Generative AI. When you want to use the Oracle dedicated AI cluster, you must provide the endpoint OCID of the hosting cluster. By default, the endpoint ID is empty and the model is on-demand on a shared infrastructure.
oci_runtimetype	This attribute indicates the runtime type of the provided model. This attribute is applicable to OCI Generate Text models in a dedicated AI cluster. Specify this attribute when you specify a model OCID in the model attribute or provide an endpoint in the oci_endpoint_id attribute. All permitted values can be found in OCI Generative AI runtimeType. See LlmInferenceRequest Reference. The supported values are: • COHERE • LLAMA **Note:** This attribute is deprecated. Use oci_apiformat.

Attribute Name	Description
`provider`	AI provider for the AI profile. Supported providers: • openai • cohere • azure • database • oci • google • anthropic • huggingface This is a mandatory attribute.
`region`	This attribute indicates the location of the Generative AI cluster that you want to use. The region can contain alphanumeric characters and hyphen characters. **Note:** The Oracle Generative AI cluster is available in Chicago, Frankfurt, and London regions. See Pretrained Foundational Models in Generative AI. The default region is `us-chicago-1`.
`stop_tokens`	The generated text will be terminated at the beginning of the earliest stop sequence. Sequence will be incorporated into the text. The attribute value must be a valid array of string values in JSON format. `stop_tokens` takes a JSON array as input. To learn more about stop tokens or stop sequences, see OpenAI or Cohere documentation.
`temperature`	Sampling from Generate Text models incorporates randomness, so that the same prompt may yield different outputs each time you hit "generate". Temperature is a non-negative float number used to tune the degree of randomness. Lower temperatures mean less random generations. See Temperature for more details. This parameter is applicable to all the supported service providers.
`vector_index_name`	Name of the vector index. The vector index name must follow the naming rules of Oracle SQL identifier. Maximum length of the vector store name is 125 characters.

The following example is using Cohere as the provider and displays custom profile attributes:

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
     profile_name => 'COHERE',
     attributes =>
      '{"provider": "cohere",
        "credential_name": "COHERE_CRED",
        "object_list": [{"owner": "ADB_USER"}],
        "max_tokens":512,
        "stop_tokens": [";"],
        "model": "command-nightly",
        "temperature": 0.5,
```

```
        "comments": true
       }');
END;
/
```

The following example shows custom profile attributes using OCI Generative AI:

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
      profile_name =>
'GENAI',
      attributes => '{"provider":
"oci",
        "credential_name": "GENAI_CRED",
        "object_list": [{"owner": "SH", "name": "customers"},
                        {"owner": "SH", "name": "countries"},
                        {"owner": "SH", "name": "supplementary_demographics"},
                        {"owner": "SH", "name": "profits"},
                        {"owner": "SH", "name": "promotions"},
                        {"owner": "SH", "name": "products"}],
        "oci_compartment_id": "ocid1.compartment.oc1...",
    "oci_endpoint_id": "ocid1.generativeaiendpoint.oc1.us-chicago-1....",
    "region": "us-chicago-1",
    "model": "cohere.command-light",
    "oci_apiformat": "COHERE"
       }');
END;
/
```

ENABLE_DATA_ACCESS Procedure

This procedure enables sending data to LLM for applicable Select AI features, which is the default behavior. Only an administrator can run this procedure.

This procedure controls data access for the following Select AI capabilities:

- `narrate` action
- Retrieval Augmented Generation (RAG)
- Synthetic Data Generation

Syntax

```
DBMS_CLOUD_AI.ENABLE_DATA_ACCESS();
```

Parameters

This procedure does not require any parameters.

Example to Enable Data Access

```
BEGIN
  DBMS_CLOUD_AI.ENABLE_DATA_ACCESS();
```

```
END;
/
```

DISABLE_DATA_ACCESS Procedure

This procedure disables sending data to LLM for applicable Select AI features. Only an administrator can run this procedure.

This procedure limits the following Select AI capabilities:

- `narrate` action
- Retrieval Augmented Generation (RAG)
- Synthetic Data Generation

Syntax

```
DBMS_CLOUD_AI.DISABLE_DATA_ACCESS();
```

Parameters

This procedure does not require any parameters.

Example to Disable Data Access

```
BEGIN
  DBMS_CLOUD_AI.DISABLE_DATA_ACCESS();
END;
/
```

CREATE_VECTOR_INDEX Procedure

This procedure creates a vector index in the specified vector database, and populates it with data from an object store using an asynchronous scheduler job.

Syntax

```
PROCEDURE CREATE_VECTOR_INDEX(
   index_name          IN  VARCHAR2,
   attributes          IN  CLOB      DEFAULT NULL,
   status              IN  VARCHAR2  DEFAULT NULL,
   description         IN  CLOB      DEFAULT NULL
);
```

Parameters

Parameter	Description
`index_name`	Name of the vector index. The vector index name must follow the naming rules of Oracle SQL identifier. Maximum length of vector store name is 125 characters. This is a mandatory parameter.

Parameter	Description
attributes	Custom attributes for the vector index in JSON. To see a list of configurable parameters, see Vector Index Attributes. The default value is NULL.
status	Status of the vector index. The possible values are: • Enabled • Disabled The default value is Disabled.
description	Description for the vector index. The default value is NULL.

Example

The following example demonstrates how to create a vector index and configure the attributes as JSON parameters.

```
BEGIN
       DBMS_CLOUD_AI.CREATE_VECTOR_INDEX(
            index_name    => 'MY_INDEX'
            attributes    => JSON_OBJECT(
                       'vector_db_provider' value 'oracle',
                       'vector_table_name'  value
'oracle_mycollection',
                       'profile_name'      value 'OCIGENAI',
                       'location'          value
                         'https://objectstorage.us-phoenix-1.' ||
                         'oraclecloud.com/n/mynamespace/b/mybucket',
                       'object_store_credential_name'   value
'OS_CRED',
                       'chunk_size'        value 2048,
                       'chunk_overlap'     value 256,
                       'refresh_rate'      value 720)
       );
END;
/
```

Related Topics

- DBA_CLOUD_VECTOR _INDEXES View
- USER_CLOUD_VECTOR _INDEXES View

DROP_VECTOR_INDEX Procedure

This procedure removes a vector store index. It normally removes the vector store index object and deletes the vector store. If set to FALSE, the argument include_data ensures the procedure only removes the vector store index object while retaining the vector store.

Syntax

```
PROCEDURE DROP_VECTOR_INDEX(
   index_name          IN  VARCHAR2,
   include_data        IN  BOOLEAN DEFAULT TRUE,
```

```
   force               IN  BOOLEAN DEFAULT FALSE
);
```

Parameters

Parameter	Description
index_name	Name of the vector index. The vector index name must follow the naming rules of Oracle SQL identifier. Maximum length of vector store name is 125 characters. This is a mandatory parameter.
include_data	Indicates whether to delete both the customer's vector store and vector index along with the vector index object. Possible values: • TRUE • FALSE The default value is TRUE.
force	Indicates whether to ignore errors that occur if the vector index does not exist. Possible values: • TRUE • FALSE If set to TRUE, this parameter bypasses any errors that occur if the vector index does not exist. The default value is FALSE.

Example

```
BEGIN
DBMS_CLOUD_AI.DROP_VECTOR_INDEX(
        index_name     => 'MY_INDEX',
        include_data   => FALSE,
        force          => TRUE
    );
END;
/
```

DISABLE_VECTOR_INDEX Procedure

This procedure disables a vector index object in the current database. When disabled, an AI profile cannot use the vector index, and the system does not load data into the vector store as new data is added to the object store and does not perform indexing, searching or querying based on the index.

Syntax

```
DBMS_CLOUD_AI.DISABLE_VECTOR_INDEX(

   index_name       IN  VARCHAR2
);
```

Parameters

Parameter	Description
index_name	Name of the vector index. The vector index name must follow the naming rules of Oracle SQL identifier. Maximum length of vector store name is 125 characters. This is a mandatory parameter.

Example

You can disable a vector index by providing the name of the vector index.

```
BEGIN
   DBMS_CLOUD_AI.DISABLE_VECTOR_INDEX(index_name => 'MY_INDEX');
END;
/
```

ENABLE_VECTOR_INDEX Procedure

This procedure enables or activates a previously disabled vector index object. Generally, when you create a vector index, by default it is enabled such that the AI profile can use it to perform indexing and searching.

When enabled, a vector index allows an AI profile to use it for loading new data from an object store into a vector store at a user-specified refresh rate. You can specify the `refresh_rate` parameter through the JSON object list. To configure the JSON attributes, see Vector Index Attributes.

Syntax

```
DBMS_CLOUD_AI.ENABLE_VECTOR_INDEX(

   index_name       IN  VARCHAR2
);
```

Parameters

Parameter	Description
index_name	Name of the vector index. The vector index name must follow the naming rules of Oracle SQL identifier. Maximum length of vector store name is 125 characters. This is a mandatory parameter.

Example

You can enable or activate a vector index by specifying the vector index name as follows:

```
BEGIN
   DBMS_CLOUD_AI.ENABLE_VECTOR_INDEX(index_name => 'MY_INDEX');
END;
/
```

UPDATE_VECTOR_INDEX Procedure

This procedure updates an existing vector store index with a specified value of the vector index attribute. It is overloaded to accept attribute values of various types.

Syntax

```
DBMS_CLOUD_AI.UPDATE_VECTOR_INDEX(
     index_name          IN  VARCHAR2,
     attribute_name      IN  VARCHAR2,
     attribute_value     IN  VARCHAR2
);

DBMS_CLOUD_AI.UPDATE_VECTOR_INDEX(

     index_name          IN  VARCHAR2,
     attribute_name      IN  VARCHAR2,
     attribute_value     IN  CLOB     DEFAULT NULL
  );
```

Parameters

Parameter	Description
`index_name`	Name of the vector index. The vector index name must follow the naming rules of Oracle SQL identifier. Maximum length of vector store name is 125 characters. This is a mandatory parameter.
`attribute_name`	Name of the custom attributes specified as JSON parameters in `CREATE_VECTOR_INDEX` procedure. You cannot modify the following attributes: • `location` • `vector_table_name` • `chunk_size` • `chunk_overlap` • `vector_distance_metric` • `vector_dimension` This is a mandatory parameter.
`attribute_value`	User specified value for the custom `attribute_name`. Value can be `CLOB`, `NUMBER` or `VARCHAR2` type. The default value is NULL.

Example

```
BEGIN
   DBMS_CLOUD_AI.UPDATE_VECTOR_INDEX(
      index_name       => 'MY_INDEX',
      attribute_name   => 'object_storage_credential_name',
      attribute_value  => 'NEW_CRED'
   );
END;
/
```

The following example accepts `NUMBER` type as the `attribute_value`.

```
BEGIN
   DBMS_CLOUD_AI.UPDATE_VECTOR_INDEX(
      index_name       => 'MY_INDEX',
      attribute_name   => 'match_limit',
      attribute_value  => 10
   );
END;
/
```

The following example accepts `VARCHAR2` type as the `attribute_value`.

```
BEGIN
   DBMS_CLOUD_AI.UPDATE_VECTOR_INDEX(
      index_name       => 'MY_INDEX',
      attribute_name   => 'profile_name',
      attribute_value  => 'AI_PROF2'
   );
END;
/
```

Vector Index Attributes

Attributes of a vector index help to manage and configure the behavior of the vector index. You can add custom index attributes as necessary. Some attributes are optional and have a default value.

Attributes

Attribute Name	Value	Mandatory	Description
`chunk_size`	`1024` (default)	N	Text size of chunking the input data. For text data, this means the number of characters.
`chunk_overlap`	`128` (default)	N	Specifies the amount of overlapping characters between adjacent chunks of text. This attribute is useful for ensuring contextual continuity and accuracy in text processing by allowing overlaps between segments, which helps prevent loss of contextual information at chunk boundaries.
`location`	NA	Y	Location of the object store. The files in this location can be documents in formats such as PDF, DOC, JSON, XML, or HTML. See Supported Document Formats.
`match_limit`	`5` (default)	N	Specifies the maximum number of results to return in a vector search query, controlling the output size and improving the efficiency of data retrieval operations.

Attribute Name	Value	Mandatory	Description
object_storage_credential_name	NA	Yes	Specifies the name of the credentials for accessing an object storage.
pipeline_name	<vector_index_name>$VECPIPELINE	No	Specifies the name of the vector index data load pipeline. This attribute is automatically set for the vector index, you cannot specify or modify. The pipeline name can be used to monitor the vector index data load using Monitor and Troubleshoot Pipelines.
profile_name	NA	Yes	Name of the AI profile which is used for embedding resource data and user prompts.
refresh_rate	1440 minutes (default)	No	Interval of updating data in the vector store. The unit is minutes.
similarity_threshold	0 (default)	No	Defines the minimum level of similarity required for two items to be considered a match, useful for filtering results in matching algorithms to ensure relevance.
vector_distance_metric	A string corresponding to one of the values specified in the description.	No	Specifies the type of distance calculation used to compare vectors in a database, determining how similarity between items is quantified. Valid values for Oracle 23ai: • EUCLIDEAN • L2_SQUARED (EUCLIDEAN_SQUARED) • COSINE (default) • DOT • MANHATTAN • HAMMING
vector_db_provider	oracle	Yes	Specifies the provider name that manages and serves as the vector store.
vector_dimension	NA	No	Specifies the number of elements in each vector within the vector store, defining the size and structure of the data representation.
vector_table_name	<vector_index_name>$VECTAB (default)	No	Specifies the name of the table or collection to store vector embeddings and chunked data.

The following example demonstrates creating a vector index with OCI Generative AI vector store.

```
BEGIN
       DBMS_CLOUD_AI.CREATE_VECTOR_INDEX(
         index_name  => 'MY_INDEX',
         attributes  => '{"vector_db_provider": "oracle",
```

```
                          "location": "https://swiftobjectstorage.us-
phoenix-1.oraclecloud.com/v1/my_namespace/my_bucket/my_data_folder",
                          "object_storage_credential_name": "OCI_CRED",
                          "profile_name": "OPENAI_ORACLE",
                          "vector_dimension": 1024,
                          "vector_distance_metric": "cosine",
                          "chunk_overlap":128,
                          "chunk_size":1024
      }');
END;
/
/
```

DBMS_CLOUD_FUNCTION Package

The `DBMS_CLOUD_FUNCTION` package allows you invoke OCI and AWS Lambda remote functions in your Autonomous Database as SQL functions.

- Summary of DBMS_CLOUD_FUNCTION Subprograms
  This table summarizes the subprograms included in the `DBMS_CLOUD_FUNCTION` package.

Summary of DBMS_CLOUD_FUNCTION Subprograms

This table summarizes the subprograms included in the `DBMS_CLOUD_FUNCTION` package.

Subprogram	Description
CREATE_CATALOG Procedure	This procedure creates a catalog.
CREATE_FUNCTION Procedure	This procedure creates functions in a catalog.
DROP_CATALOG Procedure	This procedure drops a catalog and functions created using the catalog
DROP_FUNCTION Procedure	This procedure drops functions from a catalog.
LIST_FUNCTIONS Procedure	This procedure lists all the functions in a catalog.
SYNC_FUNCTIONS Procedure	This procedure creates a PL/SQL wrapper for adding new functions to the catalog and removing wrappers for functions that have been deleted from the catalog.

- CREATE_CATALOG Procedure
  This procedure creates a catalog in the database. The `DBMS_CLOUD_FUNCTION.CREATE_CATALOG` procedure creates a catalog. A catalog is a set of functions that creates the required infrastructure to execute subroutines. This procedure is overloaded.
- CREATE_FUNCTION Procedure
  This procedure creates functions in a catalog. There are two overloaded `DBMS_CLOUD_FUNCTION.CREATE_FUNCTION` procedures.
- DROP_CATALOG Procedure
  The `DBMS_CLOUD_FUNCTION.DROP_CATALOG` procedure drops the catalog and functions created using the catalog. This procedure is overloaded.
- DROP_FUNCTION Procedure
  The `DBMS_CLOUD_FUNCTION.DROP_FUNCTION` procedure drops the function. This procedure is overloaded.

- LIST_FUNCTIONS Procedure
  This procedure lists all the functions in a catalog.
- SYNC_FUNCTIONS Procedure
  This procedure creates a PL/SQL wrapper for adding new functions to the catalog and removing wrappers for functions that have been deleted from the catalog.

CREATE_CATALOG Procedure

This procedure creates a catalog in the database. The `DBMS_CLOUD_FUNCTION.CREATE_CATALOG` procedure creates a catalog. A catalog is a set of functions that creates the required infrastructure to execute subroutines. This procedure is overloaded.

Syntax

```
DBMS_CLOUD_FUNCTION.CREATE_CATALOG (
    credential_name             IN VARCHAR2,
    catalog_name                IN VARCHAR2,
    service_provider            IN VARCHAR2,
    cloud_params                IN CLOB
);

DBMS_CLOUD_FUNCTION.CREATE_CATALOG (
    library_name                IN VARCHAR2,
    library_listener_url        IN VARCHAR2,
    library_wallet_dir_name     IN VARCHAR2,
    library_ssl_server_cert_dn  IN VARCHAR2,
    library_remote_path         IN VARCHAR2
);
```

Parameters

Parameter	Description
`credential_name`	Specifies the name of the credential for authentication. This parameter is mandatory.
`service_provider`	Specifies the type of the service provider. This parameter can have `OCI` or `AWS` or `AZURE` or `GCP` as a parameter value. This parameter is mandatory.
`catalog_name`	Specifies the catalog name. This parameter is mandatory.
`cloud_params`	Provides parameter to the function. For example, Compartment OCID, Regions and Azure subscription_id. This parameter is mandatory.
`library_name`	Specifies the name of the library when creating a remote library. This parameter is mandatory.
`library_listener_url`	Specifies the remote location of the library. The parameter accepts a String value in *`host_name:port_number`* format. For example: `EHRPMZ_DBDOMAIN.adb-us-phoenix1.com:16000` This parameter is mandatory.

Parameter	Description
library_remote_path	Specifies the remote library path. You must provide the full absolute path to the remote library. For example:/u01/app/oracle/product/21.0.0.0/client_1/lib/libst_shape.so This parameter is mandatory.
library_wallet_dir_name	Specifies the directory where the self-signed wallet is stored. This parameter is mandatory.
library_ssl_server_cert_dn	Specifies the server certificate Distinguished Name (DN). This parameter is mandatory.

Errors

Error Code	Description
ORA-20000	This error is raised in either of the following conditions: • cloud_params value is missing or incorrect parameter values are passed. • library_name value is not unique when creating a library.
ORA-20001	This error is raised in either of the following conditions: • The credential referenced in the credential_name does not exist. • The Listener specified at the library_listener_url is not reachable when creating a library.
ORA-20002	This error is raised in either of the following conditions: • This error is raised when the catalog already exists. • The specified Server certificate directory is empty when creating a library.
ORA-20009	This error is raised when the service provider doesn't exist.

Examples

```
BEGIN
    DBMS_CLOUD_FUNCTION.CREATE_CATALOG (
        credential_name  => 'DEFAULT_CREDENTIAL',
        catalog_name     => 'OCI_DEMO_CATALOG',
        service_provider => 'OCI',
        cloud_params     => ("region_id":"us-phoenix-1",
"compartment_id":"compartment_id"
 );
END;
/


BEGIN
    DBMS_CLOUD_FUNCTION.CREATE_CATALOG (
        credential_name  => 'AZURE$PA',
        catalog_name     => 'AZURE_DEMO_CATALOG',
        service_provider => 'AZURE',
        cloud_params     => '{"subscription_id":"44495e6a-8ff1-4161-
b387-0e14e675b878"}'
 );
```

```
END;
/


BEGIN
    DBMS_CLOUD_FUNCTION.CREATE_CATALOG (
    library_name               => 'EXT_DEMOLIB',
    library_listener_url       => 'remote_extproc_hostname:16000',
    library_wallet_dir_name    => 'WALLET_DIR',
    library_ssl_server_cert_dn => 'CN=VM Hostname',
    library_remote_path        => '/u01/app/oracle/extproc_libs/library name'
 );
END;
/


BEGIN
    DBMS_CLOUD_FUNCTION.CREATE_CATALOG (
        credential_name  => 'GCP$PA',
        catalog_name     => 'GCP_DEMO_CATALOG',
        service_provider => 'GCP',
        cloud_params     => '{"project_id":"example_XXXXXX"}'
 );
END;
/
```

Usage Note

- To create a catalog you must be logged in as the `ADMIN` user or have privileges on the following:
  - `DBMS_CLOUD_OCI_FNC_FUNCTIONS_INVOKE`
  - `DBMS_CLOUD_OCI_FNC_FUNCTIONS_INVOKE_INVOKE_FUNCTION_RESPONSE_T`
  - `DBMS_CLOUD`
  - Read privilege on `USER_CLOUD_FUNCTION`
  - Read privilege on `USER_CLOUD_FUNCTION_CATALOG`

CREATE_FUNCTION Procedure

This procedure creates functions in a catalog. There are two overloaded `DBMS_CLOUD_FUNCTION.CREATE_FUNCTION` procedures.

The `DBMS_CLOUD_FUNCTION.CREATE_FUNCTION` procedure is only supported for cloud functions.

`CREATE_FUNCTION` Syntax

```
DBMS_CLOUD_FUNCTION.CREATE_FUNCTION (
    credential_name   IN VARCHAR2,
    catalog_name      IN VARCHAR2,
    function_name     IN VARCHAR2,
    function_id       IN VARCHAR2,
    input_args        IN CLOB      DEFAULT NULL,
    return_type       IN VARCHAR2  DEFAULT 'CLOB',
```

```
    response_handler  IN VARCHAR2  DEFAULT NULL
);
```

Response Handler signature

```
<USER DEFINED TYPE> response_handler_name(function_response in CLOB)RETURNS
CLOB;
```

The return type of this is user defined type or PL/SQL type. The `function_response` is of JSON with fields.

```
'{
"STATUS":"<RESPONCE STATUS>",
"RESPONSE_BODY":"<FUNCTION RESPONSE>"
}'
```

`CREATE_FUNCTION` Parameters

Parameter	Description
`credential_name`	Specifies the name of the credential for authentication. This parameter is mandatory.
`catalog_name`	Specifies the catalog name. This parameter is mandatory.
`function_name`	Specifies the PL/SQL function name. This parameter is mandatory.
`function_id`	The `function_id` parameter value refers to the OCI function, AWS Lambda or Azure function. This parameter is mandatory.
`input_args`	Specifies the key value JSON pair accepting input arguments and their types.
`return_type`	Defines the return type of the function. The return type is of `CLOB` data type.
`response_handler`	Specifies the user defined callback to handle response.

`CREATE_FUNCTION` Exceptions

Error Code	Description
`ORA-20001`	This error is raised when the credential referenced in the `credential_name` does not exist.
`ORA-20003`	This error is raised when the specified catalog does not exist.
`ORA-20004`	This error is raised when the specified function already exists.
`ORA-20005`	This error is raised when the function ID or function Amazon Resource Names (ARN) does not exist.
`ORA-20006`	This error is raised when the input arguments are invalid.
`ORA-20007`	This error is raised when the return type is missing or invalid.
`ORA-20008`	This error is raised when the response handler is missing or invalid.

CREATE_FUNCTION Example

```
VAR function_args CLOB;
EXEC :function_args := TO_CLOB('{"command": "VARCHAR2", "value":
"VARCHAR2"}');
BEGIN
    DBMS_CLOUD_FUNCTION.CREATE_FUNCTION (
    credential_name  => 'DEFAULT_CREDENTIAL',
       catalog_name     => 'OCI_DEMO_CATALOG',
       function_name    => 'demo_function',
       function_id      =>
'ocid1.fnfunc.oc1.phx.aaaaaaaazkrbjv6ntowwxlbbp5ct4otsj4o2hdw4ayosyosv4sthmya2
lyza',
       input_args       => :function_args);
 );
END;
/
```

CREATE_FUNCTION Syntax

```
DBMS_CLOUD_FUNCTION.CREATE_FUNCTION (
  library_name          IN VARCHAR2,
  function_name         IN VARCHAR2,
  function_id           IN VARCHAR2 DEFAULT NULL,
  plsql_params          IN CLOB DEFAULT NULL,
  external_params       IN CLOB DEFAULT NULL,
  api_type              IN VARCHAR2 DEFAULT 'FUNCTION',
  with_context          IN BOOLEAN DEFAULT FALSE,
  return_type           IN VARCHAR2 DEFAULT NULL
);
```

CREATE_FUNCTION Parameters

Parameter	Description
library_name	Specifies the remote library name. This parameter is mandatory.
function_name	Specifies the PL/SQL function name. This parameter is mandatory.
function_id	The function_id parameter value refers to the external procedures (extproc). If the value for function_id is not provided, the value in the function_name is used.

Parameter	Description
plsql_params	Specifies the key value JSON pair accepting the parameters for the PL/SQL wrapper. The values must be provided in the "var_name":"modetype, datatype" format. • var_name : is the name of the variable. This parameter is mandatory. • modetype : specifies the variable mode. The variable mode can be one of the following: – IN – OUT – IN OUT • datatype : specifies the variable datatype. This parameter is mandatory. The default value for plsql_params is NULL.
external_params	Specifies the parameters that need to be provided to the external C function. If value is not provided for external_params, the PL/SQL parameters are used.
api_type	Specifies the type of API (function or procedure). The default value for api_type is function.
with_context	Specifies that a context pointer is passed to the external procedure. This context is used by the external C library for connecting back to the database. The default value for with_context is FALSE.
return_type	Specifies the return type of the function created.

CREATE_FUNCTION Example

```
DECLARE
    plsql_params clob    := TO_CLOB('{"sal": "IN, FLOAT", "comm" :"IN,
FLOAT"}');
    external_params clob := TO_CLOB('sal FLOAT, sal INDICATOR SHORT, comm
FLOAT, comm INDICATOR SHORT,
    RETURN INDICATOR SHORT, RETURN FLOAT');
BEGIN
DBMS_CLOUD_FUNCTION.CREATE_FUNCTION (
    LIBRARY_NAME    => 'demolib',
    FUNCTION_NAME   => '"PercentComm"',
    PLSQL_PARAMS    => plsql_params,
    EXTERNAL_PARAMS => external_params,
    API_TYPE        => 'FUNCTION',
    WITH_CONTEXT    => FALSE,
    RETURN_TYPE     => 'FLOAT'
);
END;
/
```

DROP_CATALOG Procedure

The `DBMS_CLOUD_FUNCTION.DROP_CATALOG` procedure drops the catalog and functions created using the catalog. This procedure is overloaded.

Syntax

```
DBMS_CLOUD_FUNCTION.DROP_CATALOG (
    catalog_name      IN VARCHAR2
 );

DBMS_CLOUD_FUNCTION.DROP_CATALOG (
    library_name      IN VARCHAR2
 );
```

Parameters

Parameter	Description
`catalog_name`	Specifies the catalog name. This parameter is mandatory.
`library_name`	Specifies the library name. This parameter is mandatory.

Errors

Error Code	Description
`ORA-20003`	This error is raised when the specified catalog doesn't exist.

Example:

```
BEGIN
    DBMS_CLOUD_FUNCTION.DROP_CATALOG (
      catalog_name     => 'OCI_DEMO_CATALOG'
  );
END;
/
```

Example:

```
BEGIN
    DBMS_CLOUD_FUNCTION.DROP_CATALOG (
      library_name     => 'library_name'
  );
END;
/
```

DROP_FUNCTION Procedure

The `DBMS_CLOUD_FUNCTION.DROP_FUNCTION` procedure drops the function. This procedure is overloaded.

The `DBMS_CLOUD_FUNCTION.DROP_FUNCTION` procedure is only supported for cloud functions.

Syntax

```
DBMS_CLOUD_FUNCTION.DROP_FUNCTION (
    catalog_name      IN VARCHAR2,
    function_name     IN VARCHAR2
);

DBMS_CLOUD_FUNCTION.DROP_FUNCTION (
    library_name     IN VARCHAR2,
    function_name    IN VARCHAR2
);
```

Parameters

Parameter	Description
`catalog_name`	Specifies the catalog name. This parameter is mandatory.
`function_name`	Specifies the name of the function to be dropped. This parameter is mandatory.
`library_name`	Specifies the library name. This parameter is mandatory.

Examples

```
BEGIN
    DBMS_CLOUD_FUNCTION.DROP_FUNCTION (
       catalog_name     => 'OCI_DEMO_CATALOG',
       function_name    => 'demo_function');
END;
/


BEGIN
    DBMS_CLOUD_FUNCTION.DROP_FUNCTION (
       library_name     => 'EXTPROC_DEMO_LIBRARY',
       function_name    => 'demo_function');
END;
/
```

LIST_FUNCTIONS Procedure

This procedure lists all the functions in a catalog.

Syntax

```
DBMS_CLOUD_FUNCTION.LIST_FUNCTIONS (
    credential_name   IN VARCHAR2,
    catalog_name      IN VARCHAR2,
    function_list     OUT VARCHAR2
);
```

Parameters

Parameter	Description
`credential_name`	Specifies the name of the credential for authentication. This parameter is mandatory.
`function_list`	Returns the list of functions in JSON format. This parameter is mandatory.
`catalog_name`	Specifies the catalog name. This parameter is mandatory.

Errors

Error Code	Description
`ORA-20000`	This error is raised when `cloud_params` value is missing or incorrect parameter values are passed.
`ORA-20001`	This error is raised when the credential referenced in the `credential_name` does not exist.
`ORA-20003`	This error is raised when the specified catalog does not exist.

Example:

```
VAR function_list CLOB;
BEGIN
    DBMS_CLOUD_FUNCTION.LIST_FUNCTIONS (
    credential_name  => 'DEFAULT_CREDENTIAL',
    catalog_name     => 'OCI_DEMO_CATALOG',
    function_list    => :function_list);
 );
END;
/
SELECT JSON_QUERY(:function_list, '$' RETURNING VARCHAR2(32676) pretty) AS
search_results FROM dual;
```

SYNC_FUNCTIONS Procedure

This procedure creates a PL/SQL wrapper for adding new functions to the catalog and removing wrappers for functions that have been deleted from the catalog.

Syntax

```
DBMS_CLOUD_FUNCTION.SYNC_FUNCTIONS (
    catalog_name      IN VARCHAR2,
    refresh_rate     IN VARCHAR2 DEFAULT 'DAILY'
);
```

Parameters

Parameter	Description
`catalog_name`	Specifies the catalog name. This parameter is mandatory.
`refresh_rate`	Specifies the refresh rate of the function. `refresh_rate` can accept the following values: • `HOURLY` • `DAILY` • `WEEKLY` • `MONTHLY` The default value for this parameter is `DAILY`.

Errors

Error Code	Description
`ORA-20003`	This error is raised when the specified catalog does not exist.
`ORA-20004`	This error is raised when an invalid value is passed for the `refresh_rate` parameter.

Example:

```
BEGIN
    DBMS_CLOUD_FUNCTION.SYNC_FUNCTIONS (
       catalog_name     => 'OCI_DEMO_CATALOG'
 );
END;
/
```

DBMS_CLOUD_LINK Package

The `DBMS_CLOUD_LINK` package allows a user to register a table or a view as a data set for read only access with Cloud Links.

- DBMS_CLOUD_LINK Overview
  Describes the use of the `DBMS_CLOUD_LINK` package.

- Summary of DBMS_CLOUD_LINK Subprograms
  Shows a table with a summary of the subprograms included in the DBMS_CLOUD_LINK package.

DBMS_CLOUD_LINK Overview

Describes the use of the DBMS_CLOUD_LINK package.

The DBMS_CLOUD_LINK package provides the REGISTER procedure that allows you to register a table or a view as a data set for use with Cloud Links. Before you can register a data set, the ADMIN user must grant a user permission to register a data set using the DBMS_CLOUD_LINK_ADMIN.GRANT_REGISTER procedure. After the ADMIN runs GRANT_REGISTER, a user can register a table or a view they own as a registered data set (or register an object in another schema if the user has READ WITH GRANT OPTION privilege on the object). Registered data sets provide remote access to the registered object with Cloud Links, subject to the scope specified with the REGISTER procedure.

To run DBMS_CLOUD_LINK.REGISTER, DBMS_CLOUD_LINK.UPDATE_REGISTRATION, or DBMS_CLOUD_LINK.UNREGISTER you have to have execute privilege on DBMS_CLOUD_LINK package, in addition to having previously run DBMS_CLOUD_LINK_ADMIN.GRANT_REGISTER. By default only the ADMIN user and schemas with PDB_DBA role have execute privilege on DBMS_CLOUD_LINK.

Summary of DBMS_CLOUD_LINK Subprograms

Shows a table with a summary of the subprograms included in the DBMS_CLOUD_LINK package.

Subprogram	Description
DESCRIBE Function	This function retrieves the description for a data set. The description is provided when a data set is registered with DBMS_CLOUD_LINK.REGISTER.
FIND Procedure	Retrieves the namespace, name, and description for data sets that match the search string. Matching data sets are only shown if they are accessible to the user, based on access restrictions.
GET_DATABASE_ID Function	Returns a unique identifier for the Autonomous Database instance. Repeated calls to DBMS_CLOUD_LINK.GET_DATABASE_ID on the same instance always return the same value.
GRANT_AUTHORIZATION Procedure	Grants authorization to a specified database to access the specified data set.
REGISTER Procedure	Registers a table or view as data set.
REVOKE_AUTHORIZATION Procedure	Revokes authorization for a specified database to access the specified data set.
UNREGISTER Procedure	Removes a registered data set.
UPDATE_REGISTRATION Procedure	Updates attributes of a data set that has been registered using DBMS_CLOUD_LINK.REGISTER.

- DESCRIBE Function
- FIND Procedure
- GET_DATABASE_ID Function
- GRANT_AUTHORIZATION Procedure
- REGISTER Procedure

- REVOKE_AUTHORIZATION Procedure
- UNREGISTER Procedure
- UPDATE_REGISTRATION Procedure

DESCRIBE Function

This function retrieves the description for a data set. The description is provided when a data set is registered with `DBMS_CLOUD_LINK.REGISTER`.

Syntax

```
DBMS_CLOUD_LINK.DESCRIBE(
      namespace         IN   VARCHAR2,
      name              IN   VARCHAR2
) return CLOB;
```

Parameters

Parameter	Description
`namespace`	Specifies the namespace of registered data set.
`name`	Specifies the name of a registered data set.

Usage Note

You can use this function subject to the access restrictions imposed at registration time with `DBMS_CLOUD_LINK.REGISTER`. If a data set is not accessible to a database, then its description will not be retrieved.

FIND Procedure

This procedure retrieves the namespace, name, and description for data sets that match the search string. Matching data sets are only shown if they are accessible to the user, based on access restrictions.

Syntax

```
DBMS_CLOUD_LINK.FIND(
      search_string         IN   VARCHAR2,
      search_result         OUT   CLOB
);
```

Parameters

Parameter	Description
`search_string`	Specifies the search string. The search string is not case sensitive.
`search_result`	A JSON document that includes the namespace, name, and description values for the data set.

Usage Note

The search string is not case sensitive and the package leverages free text search using Oracle Text.

GET_DATABASE_ID Function

The function returns a unique identifier for the Autonomous Database instance. Repeated calls to `DBMS_CLOUD_LINK.GET_DATABASE_ID` on the same instance always return the same value.

You can call this function on a database that is accessing a registered data set remotely to obtain the database ID. This allows you to provide the database iD so that a data set owner can leverage a more fine-grained data access control, for example with VPD, based on a specified database IDs from remote sites.

A database ID identifies each remote database that accesses a registered data set, to track and audit access in the V$CLOUD_LINK_ACCESS_STATS and GV$CLOUD_LINK_ACCESS_STATS Views on the database that owns a registered data set.

Syntax

```
DBMS_CLOUD_LINK.GET_DATABASE_ID()
  RETURN VARCHAR2;
```

Usage Notes

Cloud links use the unique identifier that `DBMS_CLOUD_LINK.GET_DATABASE_ID` returns to identify individual databases that are accessing a data set remotely. The database owning the registered data set tracks and audits the database ID as a record of the origin for data set access in the V$CLOUD_LINK_ACCESS_STATS and GV$CLOUD_LINK_ACCESS_STATS Views.

The `DBMS_CLOUD_LINK.GET_DATABASE_ID` identifier is available as a `SYS_CONTEXT` value so that you can programmatically obtain this information about a connecting remote session using `SYS_CONTEXT`, to further restrict and control what specific data can be accessed remotely by individual Autonomous Database instances with Virtual Private Databases (VPD)s.

Return Values

A unique identifier for the Autonomous Database instance of `VARCHAR2`.

GRANT_AUTHORIZATION Procedure

This procedure grants authorization to a specified database to access the specified data set.

Syntax

```
DBMS_CLOUD_LINK.GRANT_AUTHORIZATION(
      database_id         IN   VARCHAR2,
      namespace           IN   VARCHAR2 DEFAULT,
      name                IN   VARCHAR2
);
```

Parameters

Parameter	Description
database_id	Specifies the database ID for an Autonomous Database instance. Use DBMS_CLOUD_LINK.GET_DATABASE_ID to obtain the database ID.
namespace	Specifies the data set namespace to grant access authorization for the specified database_id.
name	Specifies the data set name to grant access authorization for the specified database_id.

REGISTER Procedure

The procedure registers a table or a view as a data set to allow remote read only access, subject to restrictions imposed by the scope parameter.

Syntax

```
DBMS_CLOUD_LINK.REGISTER(
     schema_name      IN VARCHAR2,
     schema_object    IN VARCHAR2,
     namespace        IN VARCHAR2,
     name             IN VARCHAR2,
     description      IN CLOB,
     scope            IN CLOB,
     auth_required    IN BOOLEAN DEFAULT,
     data_set_owner   IN VARCHAR2 DEFAULT,
     offload_targets  IN CLOB DEFAULT

);
```

Parameters

Parameter	Description
schema_name	Specifies the owner of the table or view specified with the schema_object parameter.
schema_object	Specifies the name of a table or a view. Valid objects are: • Tables (including Heap, External or Hybrid) • Views • Materialized Views • Cloud Tables Other objects such as analytic views or synonyms are not supported.
namespace	Specifies the namespace for the data set. A NULL value specifies a system-generated namespace value, unique to the Autonomous Database instance.
name	Specifies the name for the data set.
description	Specifies text to describe the data.

Parameter	Description
scope	Describes who is allowed to access the data set. The value is a comma separated list consisting of one or more of the following: • Database OCID: Access to the data set is allowed for the specific Autonomous Database instances identified by OCID. • Compartment OCID: Access to the data set is allowed for databases in the compartments identified by compartment OCID. • Tenancy OCID: Access to the data set is allowed for databases in the tenancies identified by tenancy OCID. • Region name: Access to the data set is allowed for databases in the region identified by the named region. By scope, Cloud Links access is limited to within a single region and is not cross-region. A consumer in a different region can access a data set only when a cross-region Refreshable Clone of the database that is the data set owner exists in the consumer database's region. See Register or Unregister a Data Set in a Different Region for more information. • MY$COMPARTMENT: Access to the data set is allowed for databases in the same compartment as the data set owner. • MY$TENANCY: Access to the data set is allowed for databases in the same tenancy as the data set owner. • MY$REGION: Access to the data set is allowed for databases in the same region as the data set owner. • List of regions: Access to the data set is allowed for databases in the specified regions For example: `scope => 'us-phoenix-1,us-ashburn-1',` A consumer in a different region can access a data set only when a cross-region Refreshable Clone of the database that is the data set owner exists in the consumer database's region. The scope values, MY$REGION, MY$TENANCY, and MY$COMPARTMENT are variables that act as convenience macros and resolve to OCIDs.
auth_required	Specifies that an additional authorization is required for databases to read from the data set. The following cases are possible: • Databases that are within the SCOPE specified and have been authorized with DBMS_CLOUD_LINK.GRANT_AUTHORIZATION can view rows from the data set. • Any databases that are within the specified SCOPE but have not been authorized with DBMS_CLOUD_LINK.GRANT_AUTHORIZATION cannot view data set rows. In this case, consumers without authorization see the data set as empty. • Databases that are not within the SCOPE specified see an error when attempting to access the data set.
data_set_owner	Specifies the data set owner. This is indicates who the data set belongs to or who is responsible for updating and maintaining the data set. For example, you can set the data_set_owner to the email address of the person who registered the data set.

Parameter	Description
offload_targets	Specifies one or more Autonomous Database OCIDs of refreshable clones where access to data sets is offloaded, from the Autonomous Database where the data set is registered. The offload_targets value is a JSON document that defines one or more CLOUD_LINK_DATABASE_ID and OFFLOAD_TARGET key value pairs: • CLOUD_LINK_DATABASE_ID is one of: – A database ID: This specifies a database ID for the data set consumer whose request is offloaded to the corresponding refreshable clone specified with the OFFLOAD_TARGET value. Obtain the database ID by running DBMS_CLOUD_LINK.GET_DATABASE_ID. See GET_DATABASE_ID Function for more information. – ANY: This specifies that any data set consumer's request is offloaded to the corresponding offload target. A consumer's data set request will be routed to the corresponding offload target. If you specify ANY without specifying database IDs, all data set requests from consumers are offloaded to the refreshable clone specified with the OFFLOAD_TARGET value. If you specify both database IDs and ANY, data set requests from consumers that do not match a database ID are offloaded to the refreshable clone specified with the OFFLOAD_TARGET value. • The OFFLOAD_TARGET is the OCID for an Autonomous Database instance that is a refreshable clone. For example, the following shows a JSON sample with three OFFLOAD_TARGET/CLOUD_LINK_DATABASE_ID value pairs: `{ "OFFLOAD_TARGETS": [{ "CLOUD_LINK_DATABASE_ID": "34xxxxx69708978", "OFFLOAD_TARGET": "ocid1.autonomousdatabase.oc1..xxxxx3pv6wkcr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqstifsfabc" }, { "CLOUD_LINK_DATABASE_ID": "34xxxxx89898978", "OFFLOAD_TARGET": "ocid1.autonomousdatabase.oc1..xxxxx3pv6wkcr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqstifsfdef" }, { "CLOUD_LINK_DATABASE_ID": "34xxxxx4755680", "OFFLOAD_TARGET": "ocid1.autonomousdatabase.oc1..xxxxx3pv6wkcr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqstifsfghi" }] }` When a data set consumer requests access to a data set that you register with offload_targets with the ANY keyword, then any request for access is offloaded to the refreshable clone identified with OFFLOAD_TARGET in the

Parameter	Description

	supplied JSON (except requests that have an explicit entry in the supplied JSON).
	For example, the following shows a JSON sample with one explicit `OFFLOAD_TARGET`/`CLOUD_LINK_DATABASE_ID` value pair, and one `ANY` entry:
	
	`{`
	`  "OFFLOAD_TARGETS": [`
	`    {`
	`      "CLOUD_LINK_DATABASE_ID": "ANY",`
	`      "OFFLOAD_TARGET":`
	`"ocid1.autonomousdatabase.oc1..xxxxx3pv6wkcr4jqae5f44n2b2`
	`m2yt2j6rx32uzr4h25vqstifsfdef"`
	`    },`
	`    {`
	`      "CLOUD_LINK_DATABASE_ID": "34xxxxx4755680",`
	`      "OFFLOAD_TARGET":`
	`"ocid1.autonomousdatabase.oc1..xxxxx3pv6wkcr4jqae5f44n2b2`
	`m2yt2j6rx32uzr4h25vqstifsfghi"`
	`    }`
	`  ]`
	`}`
	
	`DBMS_CLOUD_LINK.REGISTER` reports an error if the OCID supplied as an `OFFLOAD_TARGET` value is not an OCID of a refreshable clone in the same region.
	See Use Refreshable Clones with Autonomous Database for information on using refreshable clones.

Usage Notes

- After you register an object, users may need to wait up to ten (10) minutes to access the object with Cloud Links.
- Use the procedure `DBMS_CLOUD_LINK.UPDATE_REGISTRATION` to change the attributes for an existing data set.

  The wait time for the update to complete can be up to 10 minutes for a registration change to be propagated and accessible through Cloud Links. This delay can impact the accuracy of the data in the both the `DBA_CLOUD_LINK_REGISTRATIONS` and `DBA_CLOUD_LINK_ACCESS` views.
- You can register a table or view that resides in another user's schema when the you have `READ WITH GRANT OPTION` privileges for the table or view.
- The scope you set when you register a data set is only honored when it matches or is more restrictive than the value set with `DBMS_CLOUD_LINK_ADMIN.GRANT_REGISTER`. For example, assume the ADMIN granted the scope `'MY$TENANCY'` with `GRANT_REGISTER`, and the user specifies `'MY$REGION'` when they register a data set with `DBMS_CLOUD_LINK.REGISTER`. In this case they would see an error such as the following:

  ```
  ORA-20001: Share privileges are not enabled for current user or it is
  enabled but not for scope MY$REGION
  ```

- Certain hierarchical validity checks for registration cannot happen at registration time. Invalid registrations will not be visible, discoverable, or even accessible to anyone.
- To use `DBMS_CLOUD_LINK.REGISTER`, you must have execute privilege on the `DBMS_CLOUD_LINK` package, in addition to the register privilege assigned with `DBMS_CLOUD_LINK_ADMIN.GRANT_REGISTER`. Only the ADMIN user and schemas with `PDB_DBA` have this privilege by default.
- When you register a data set on a refreshable clone in a remote region, the invocation of `DBMS_CLOUD_LINK.REGISTER` on the remote region clone must use the same parameters with the same values as on the source database, with the exception of the `offload_targets` parameter.

  For example, when you run `DBMS_CLOUD_LINK.REGISTER` with scope set to `MY$COMPARTMENT` on the source Autonomous Database instance, run the procedure again on the cross-region refreshable clone with the same scope parameter value (`MY$COMPARTMENT`).
- If you specify the `offload_targets` parameter with `DBMS_CLOUD_LINK.REGISTER` on source, you should omit this parameter when you register the data set on a cross-region refreshable clone.

REVOKE_AUTHORIZATION Procedure

This procedure revokes authorization for a specified database to access the specified data set.

Syntax

```
DBMS_CLOUD_LINK.REVOKE_AUTHORIZATION(
      database_id         IN   VARCHAR2,
      namespace           IN   VARCHAR2 DEFAULT,
      name                IN   VARCHAR2
);
```

Parameters

Parameter	Description
`database_id`	Specifies the database ID for an Autonomous Database instance. Use `DBMS_CLOUD_LINK.GET_DATABASE_ID` to obtain the database ID.
`namespace`	Specifies the data set namespace to revoke access authorization for the specified `database_id`.
`name`	Specifies the data set name to revoke access authorization for the specified `database_id`.

UNREGISTER Procedure

The procedure allows a user who previously registered a table or a view with the `REGISTER` procedure to unregister the table or view so that it no longer is available for remote access.

Syntax

```
DBMS_CLOUD_LINK.UNREGISTER(
      namespace         IN   VARCHAR2,
      name              IN   VARCHAR2
);
```

Parameters

Parameter	Description
namespace	Specifies a username.
name	Specifies the name for the data set.

Usage Note

DBMS_CLOUD_LINK.UNREGISTER may also take up to ten (10) minutes to fully propagate, after which the data can longer be accessed remotely.

UPDATE_REGISTRATION Procedure

The procedure updates one or more of the attributes for a data set that was registered using DBMS_CLOUD_LINK.REGISTER.

Syntax

```
DBMS_CLOUD_LINK.UPDATE_REGISTRATION(
     namespace         IN VARCHAR2,
     name              IN VARCHAR2,
     description       IN CLOB      DEFAULT,
     scope             IN CLOB      DEFAULT,
     auth_required     IN BOOLEAN   DEFAULT,
     data_set_owner    IN VARCHAR2  DEFAULT,
     offload_targets   IN CLOB      DEFAULT

);
```

Parameters

Parameter	Description
namespace	Specifies the namespace of the data set to be updated.
name	Specifies the name of the data set to be updated.
description	Specifies the updated text to describe the data. If NULL is passed in for this parameter value, the attribute is not modified. By default this attribute is not updated.

Parameter	Description
scope	Update the scope with the specified value. The scope describes who is allowed to access the data set. The value is a comma separated list consisting of one or more of the following: • Database OCID: Access to the data set is allowed for the specific Autonomous Database instances identified by OCID. • Compartment OCID: Access to the data set is allowed for databases in the compartments identified by compartment OCID. • Tenancy OCID: Access to the data set is allowed for databases in the tenancies identified by tenancy OCID. • Region name: Access to the data set is allowed for databases in the region identified by the named region. By scope, Cloud Links access is limited to within a single region and is not cross-region. A consumer in a different region can access a data set only when a cross-region Refreshable Clone of the database that is the data set owner exists in the consumer database's region. See Register or Unregister a Data Set in a Different Region for more information. • MY$COMPARTMENT: Access to the data set is allowed for databases in the same compartment as the data set owner. • MY$TENANCY: Access to the data set is allowed for databases in the same tenancy as the data set owner. • MY$REGION: Access to the data set is allowed for databases in the same region as the data set owner. • List of regions: Access to the data set is allowed for databases in the specified regions For example: `scope => 'us-phoenix-1,us-ashburn-1',` A consumer in a different region can access a data set only when a cross-region Refreshable Clone of the database that is the data set owner exists in the consumer database's region. The scope values, MY$REGION, MY$TENANCY, and MY$COMPARTMENT are variables that act as convenience macros and resolve to OCIDs. If NULL is passed in for this parameter value, the attribute is not modified. By default this attribute is not updated.
auth_required	Specifies that an additional authorization is required for databases to read from the data set. The following cases are possible: • Databases that are within the SCOPE specified and have been authorized with DBMS_CLOUD_LINK.GRANT_AUTHORIZATION can view rows from the data set. • Any databases that are within the specified SCOPE but have not been authorized with DBMS_CLOUD_LINK.GRANT_AUTHORIZATION cannot view data set rows. In this case, consumers without authorization see the data set as empty. • Databases that are not within the SCOPE specified see an error when attempting to access the data set. If NULL is passed in for this parameter value, the attribute is not modified. By default this attribute is not updated.

Parameter	Description
`data_set_owner`	Specifies the data set owner. This is indicates who the data set belongs to or who is responsible for updating and maintaining the data set. For example, you can set the `data_set_owner` to the email address of the person who registered the data set. If `NULL` is passed in for this parameter value, the attribute is not modified. By default this attribute is not updated.

Parameter	Description
offload_targets	Specifies one or more Autonomous Database OCIDs of refreshable clones where access to data sets is offloaded, from the Autonomous Database where the data set is registered. The offload_targets value is a JSON document that defines one or more CLOUD_LINK_DATABASE_ID and OFFLOAD_TARGET key value pairs: • CLOUD_LINK_DATABASE_ID is one of: – A database ID: This specifies a database ID for the data set consumer whose request is offloaded to the corresponding refreshable clone specified with the OFFLOAD_TARGET value. Obtain the database ID by running DBMS_CLOUD_LINK.GET_DATABASE_ID. See GET_DATABASE_ID Function for more information. – ANY: This specifies that any data set consumer's request is offloaded to the corresponding offload target. A consumer's data set request will be routed to the corresponding offload target. If you specify ANY without specifying database IDs, all data set requests from consumers are offloaded to the refreshable clone specified with the OFFLOAD_TARGET value. If you specify both database IDs and ANY, data set requests from consumers that do not match a database ID are offloaded to the refreshable clone specified with the OFFLOAD_TARGET value. • The OFFLOAD_TARGET is the OCID for an Autonomous Database instance that is a refreshable clone. For example, the following shows a JSON sample with three OFFLOAD_TARGET/CLOUD_LINK_DATABASE_ID value pairs: `{ "OFFLOAD_TARGETS": [{ "CLOUD_LINK_DATABASE_ID": "34xxxxx69708978", "OFFLOAD_TARGET": "ocid1.autonomousdatabase.oc1..xxxxx3pv6wkcr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqstifsfabc" }, { "CLOUD_LINK_DATABASE_ID": "34xxxxx89898978", "OFFLOAD_TARGET": "ocid1.autonomousdatabase.oc1..xxxxx3pv6wkcr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqstifsfdef" }, { "CLOUD_LINK_DATABASE_ID": "34xxxxx4755680", "OFFLOAD_TARGET": "ocid1.autonomousdatabase.oc1..xxxxx3pv6wkcr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqstifsfghi" }] }` When a data set consumer requests access to a data set that you register with offload_targets with the ANY keyword, then any request for access is offloaded to the refreshable clone identified with OFFLOAD_TARGET in the

Parameter	Description

	supplied JSON (except requests that have an explicit entry in the supplied JSON).
	For example, the following shows a JSON sample with one explicit `OFFLOAD_TARGET`/`CLOUD_LINK_DATABASE_ID` value pair, and one `ANY` entry:
	
	`{`
	`  "OFFLOAD_TARGETS": [`
	`    {`
	`      "CLOUD_LINK_DATABASE_ID": "ANY",`
	`      "OFFLOAD_TARGET":`
	`"ocid1.autonomousdatabase.oc1..xxxxx3pv6wkcr4jqae5f44n2b2`
	`m2yt2j6rx32uzr4h25vqstifsfdef"`
	`    },`
	`    {`
	`      "CLOUD_LINK_DATABASE_ID": "34xxxxx4755680",`
	`      "OFFLOAD_TARGET":`
	`"ocid1.autonomousdatabase.oc1..xxxxx3pv6wkcr4jqae5f44n2b2`
	`m2yt2j6rx32uzr4h25vqstifsfghi"`
	`    }`
	`  ]`
	`}`
	
	`DBMS_CLOUD_LINK.REGISTER` reports an error if the OCID supplied as an `OFFLOAD_TARGET` value is not an OCID of a refreshable clone in the same region.
	See Use Refreshable Clones with Autonomous Database for information on using refreshable clones.
	If `NULL` is passed in for this parameter value, the attribute is not modified.
	By default this attribute is not updated.

Usage Notes

- The `schema_name` and `schema_object` attributes for a data set cannot be updated.
- After you update the registration for an object, users may need to wait up to ten (10) minutes to access the object using the updated attributes with Cloud Links.
- You can update the registration for a table or view that resides in another user's schema when the you have `READ WITH GRANT OPTION` privileges for the table or view.
- The scope you set when you update the registration of a data set is only honored when it matches or is more restrictive than the value set with `DBMS_CLOUD_LINK_ADMIN.GRANT_REGISTER`. For example, assume the ADMIN granted the scope `'MY$TENANCY'` with `GRANT_REGISTER`, and the user specifies `'MY$REGION'` when they update the registration for the data set with `DBMS_CLOUD_LINK.UPDATE_REGISTRATION`. In this case they would see an error such as the following:

  ```
  ORA-20001: Share privileges are not enabled for current user or it is
  enabled but not for scope MY$REGION
  ```

- Certain hierarchical validity checks for an updated registration cannot happen when at the time the registration is updated. Invalid registrations will not be visible, discoverable, or even accessible to anyone.

- The user who registered a data set can update its attributes with `DBMS_CLOUD_LINK.UPDATE_REGISTRATION`. In addition, you must have execute privilege on the `DBMS_CLOUD_LINK` package and the register privilege assigned with `DBMS_CLOUD_LINK_ADMIN.GRANT_REGISTER`. Only the ADMIN user and schemas with `PDB_DBA` have this privilege by default.
- When you update the registration for a data set on a refreshable clone in a remote region, the invocation of `DBMS_CLOUD_LINK.UPDATE_REGISTRATION` on the remote region clone must use the same parameters with the same values as on the source database, with the exception of the `offload_targets` parameter.

  For example, when you run `DBMS_CLOUD_LINK.UPDATE_REGISTRATION` with scope set to `MY$COMPARTMENT` on the source Autonomous Database instance, run the procedure again on the cross-region refreshable clone with the same scope parameter value (`MY$COMPARTMENT`).
- If you specify the `offload_targets` parameter with `DBMS_CLOUD_LINK.REGISTER` on the source, you should omit this parameter when you register the data set on a cross-region refreshable clone.

DBMS_CLOUD_LINK_ADMIN Package

The `DBMS_CLOUD_LINK_ADMIN` package allows the ADMIN user to enable a database user to register data sets or to access registered data sets for a given Autonomous Database instance, subject to the access restrictions as defined with the granted scope.

Privileges can also be disabled for a user that has the privileges set to register data sets or access registered data sets.

- DBMS_CLOUD_LINK_ADMIN Overview
  Describes use of the `DBMS_CLOUD_LINK_ADMIN` package.
- Summary of DBMS_CLOUD_LINK_ADMIN Subprograms
  This table summarizes the subprograms included in the `DBMS_CLOUD_LINK_ADMIN` package.

DBMS_CLOUD_LINK_ADMIN Overview

Describes use of the `DBMS_CLOUD_LINK_ADMIN` package.

Cloud Links provide a cloud-based method to remotely access read only data on an Autonomous Database instance. The `DBMS_CLOUD_LINK_ADMIN` package leverages Oracle Cloud Infrastructure access mechanisms to make data sets accessible within a specific scope and in addition there is a optional authorization step.

Summary of DBMS_CLOUD_LINK_ADMIN Subprograms

This table summarizes the subprograms included in the `DBMS_CLOUD_LINK_ADMIN` package.

Subprogram	Description
ADD_SERVICE_MAPPING Procedure	The procedure associates a Cloud Links consumer database with a database service.
GRANT_AUTHORIZE Procedure	Grants a user the permission to invoke `DBMS_CLOUD_LINK.GRANT_AUTHORIZATION` and `DBMS_CLOUD_LINK.REVOKE_AUTHORIZATION` procedures.

Subprogram	Description
GRANT_READ Procedure	Allows a user to read registered data sets, subject to access restrictions imposed on data sets at registration.
GRANT_REGISTER Procedure	Allows a user to register a data set for remote access.
REMOVE_SERVICE_MAPPING Procedure	Remove a service mapping for a specified database.
REVOKE_AUTHORIZE Procedure	Revokes a user's permission to invoke `DBMS_CLOUD_LINK.GRANT_AUTHORIZATION` and `DBMS_CLOUD_LINK.REVOKE_AUTHORIZATION` procedures.
REVOKE_READ Procedure	Disallows a user from accessing registered data sets of the Autonomous Database instance.
REVOKE_REGISTER Procedure	Disallows a user from registering data sets for remote access. Data sets that were already registered by the user are unaffected.

- ADD_SERVICE_MAPPING Procedure
- GRANT_AUTHORIZE Procedure
- GRANT_READ Procedure
- GRANT_REGISTER Procedure
- REMOVE_SERVICE_MAPPING Procedure
- REVOKE_AUTHORIZE Procedure
- REVOKE_READ Procedure
- REVOKE_REGISTER Procedure

ADD_SERVICE_MAPPING Procedure

The procedure `DBMS_CLOUD_LINK_ADMIN.ADD_SERVICE_MAPPING` associates a consumer database with a database service.

Syntax

```
DBMS_CLOUD_LINK_ADMIN.ADD_SERVICE_MAPPING(
      database_id         IN   VARCHAR2,
      service_name        IN   VARCHAR2
);
```

Parameters

Parameter	Description
`database_id`	Specifies the database ID for an Autonomous Database instance that is a Cloud Link consumer. Use `DBMS_CLOUD_LINK.GET_DATABASE_ID` to obtain the database ID. The value "`ANY`" associates the specified `service_name` value with all consumer databases that do not have an entry that matches their database identifier. That is, any `database_id` whose `service_name` has not been set with `DBMS_CLOUD_LINK_ADMIN.ADD_SERVICE_MAPPING`. Valid values: a database ID or "`ANY`".

Parameter	Description
service_name	Specifies the database service name. Valid values depend on the workload type: • **Data Warehouse**: Valid values are: HIGH, MEDIUM, LOW • **Transaction Processing**, **JSON Database**, or **APEX**: Valid values are: TP, TPURGENT, HIGH, MEDIUM, LOW

Usage Notes

- Only the ADMIN user and schemas with PDB_DBA role can run DBMS_CLOUD_LINK_ADMIN.ADD_SERVICE_MAPPING.
- You must run DBMS_CLOUD_LINK_ADMIN.ADD_SERVICE_MAPPING on the data set owner's Autonomous Database instance.

GRANT_AUTHORIZE Procedure

The procedure grants a user permission to invoke the DBMS_CLOUD_LINK.GRANT_AUTHORIZATION and DBMS_CLOUD_LINK.REVOKE_AUTHORIZATION procedures.

Syntax

```
DBMS_CLOUD_LINK_ADMIN.GRANT_AUTHORIZE(
      username         IN    VARCHAR2
);
```

Parameters

Parameter	Description
username	Specifies a username.

Usage Notes

- To enable authorization for a data set with DBMS_CLOUD_LINK.GRANT_AUTHORIZATION, you have to have granted the privilege with DBMS_CLOUD_LINK_ADMIN.GRANT_AUTHORIZE. This is true for ADMIN user as well; however ADMIN user can grant this privilege to themself.

GRANT_READ Procedure

The procedure allows a user to read registered data sets, subject to the access restrictions imposed on data sets when a data set is registered using DBMS_CLOUD_LINK_ADMIN.GRANT_REGISTER.

Syntax

```
DBMS_CLOUD_LINK_ADMIN.GRANT_READ(
      username         IN    VARCHAR2
);
```

Parameters

Parameter	Description
username	Specifies a username.

Usage Notes

- To read data sets, you have to have granted the privilege with DBMS_CLOUD_LINK_ADMIN.GRANT_READ. This is true for ADMIN user as well; however ADMIN user can grant this privilege to themself.
- A user can query SYS_CONTEXT('USERENV', 'CLOUD_LINK_READ_ENABLED') to check if they are enabled for READ access to a data set.

  For example, the following query:

  ```
  SELECT SYS_CONTEXT('USERENV', 'CLOUD_LINK_READ_ENABLED') FROM DUAL;
  ```

  Returns 'YES' or 'NO'.

GRANT_REGISTER Procedure

The procedure allows a user to register a data set for remote access.

Syntax

```
DBMS_CLOUD_LINK_ADMIN.GRANT_REGISTER(
      username        IN   VARCHAR2,
      scope           IN   CLOB
);
```

Parameters

Parameter	Description
username	Specifies a user name.
scope	Specifies the scope in which permissions to publish are to be granted to the specified user. Valid values are: • 'MY$REGION' • 'MY$TENANCY' • 'MY$COMPARTMENT'

Usage Notes

- To register data sets, you have to have granted the privilege with DBMS_CLOUD_LINK_ADMIN.GRANT_REGISTER. This is true for ADMIN user as well; however ADMIN user can grant this privilege to themself.
- A user can query SYS_CONTEXT('USERENV', 'CLOUD_LINK_REGISTER_ENABLED') to check if they are enabled for registering data sets.

For example, the following query:

```
SELECT SYS_CONTEXT('USERENV', 'CLOUD_LINK_REGISTER_ENABLED') FROM DUAL;
```

Returns 'YES' or 'NO'.

REMOVE_SERVICE_MAPPING Procedure

The procedure DBMS_CLOUD_LINK_ADMIN.REMOVE_SERVICE_MAPPING removes a service mapping for a specified database.

Syntax

```
DBMS_CLOUD_LINK_ADMIN.REMOVE_SERVICE_MAPPING(
      database_id         IN    VARCHAR2
);
```

Parameters

Parameter	Description
database_id	Specifies the database ID for an Autonomous Database instance. Use DBMS_CLOUD_LINK.GET_DATABASE_ID to obtain the database ID. The value "ANY" removes any specified service_name value with all consumer databases that do not have an entry that matches their database identifier. That is, any database_id whose service_name has not been set with DBMS_CLOUD_LINK_ADMIN.ADD_SERVICE_MAPPING. Valid values: a database ID or "ANY".

Usage Note

Only the ADMIN user and schemas with PDB_DBA role can run DBMS_CLOUD_LINK_ADMIN.REMOVE_SERVICE_MAPPING.

REVOKE_AUTHORIZE Procedure

This procedure disallows a user from invoking DBMS_CLOUD_LINK.GRANT_AUTHORIZATION and DBMS_CLOUD_LINK.REVOKE_AUTHORIZATION procedures.

Syntax

```
DBMS_CLOUD_LINK_ADMIN.REVOKE_AUTHORIZE(
      username         IN    VARCHAR2
);
```

Parameters

Parameter	Description
username	Specifies a username.

REVOKE_READ Procedure

This procedure disallows a user from accessing registered data sets on the Autonomous Database instance.

Syntax

```
DBMS_CLOUD_LINK_ADMIN.REVOKE_READ(
      username         IN    VARCHAR2
);
```

Parameters

Parameter	Description
`username`	Specifies a username.

Usage Note

- A user can query `SYS_CONTEXT('USERENV', 'CLOUD_LINK_READ_ENABLED')` to check if they are enabled for `READ` access to a data set.

  For example, the following query:

  ```
  SELECT SYS_CONTEXT('USERENV', 'CLOUD_LINK_READ_ENABLED') FROM DUAL;
  ```

  Returns '`YES`' or '`NO`'.

REVOKE_REGISTER Procedure

The procedure disallows a user from registering data sets for remote access. Data sets that were already registered by the user are unaffected.

Syntax

```
DBMS_CLOUD_LINK_ADMIN.REVOKE_REGISTER(
      username         IN    VARCHAR2
);
```

Parameters

Parameter	Description
`username`	Specifies a user name.

Usage Note

- A user can query `SYS_CONTEXT('USERENV', 'CLOUD_LINK_REGISTER_ENABLED')` to check if they are enabled for registering data sets.

  For example, the following query:

  ```
  SELECT SYS_CONTEXT('USERENV', 'CLOUD_LINK_REGISTER_ENABLED') FROM DUAL;
  ```

Returns 'YES' or 'NO'.

DBMS_CLOUD_MACADM Package

This section covers the DBMS_CLOUD_MACADM subprograms provided with Autonomous Database.

- CONFIGURE_DATABASE_VAULT Procedure
  This procedure configures the initial two Oracle Database user accounts, which are granted the DV_OWNER and DV_ACCTMGR roles, respectively for Autonomous Database.
- DISABLE_DATABASE_VAULT Procedure
  This procedure disables Oracle Database Vault on Autonomous Database. To use this procedure you must have the DV_OWNER role.
- DISABLE_USERMGMT_DATABASE_VAULT Procedure
  This procedure disallows user management related operations for specified components on Autonomous Database with Oracle Database Vault enabled.
- ENABLE_DATABASE_VAULT Procedure
  This procedure enables Oracle Database Vault on Autonomous Database. To use this procedure you must have the DV_OWNER role.
- ENABLE_USERMGMT_DATABASE_VAULT Procedure
  This procedure allows user management with Oracle Database Vault enabled for specified components on Autonomous Database.

CONFIGURE_DATABASE_VAULT Procedure

This procedure configures the initial two Oracle Database user accounts, which are granted the DV_OWNER and DV_ACCTMGR roles, respectively for Autonomous Database.

Syntax

```
DBMS_CLOUD_MACADM.CONFIGURE_DATABASE_VAULT(
       dvowner_uname    IN VARCHAR2,
       dvacctmgr_uname  IN VARCHAR2);
```

Parameters

Parameter	Description
dvowner_uname	Name of the user who will be the Database Vault Owner. This user will be granted the DV_OWNER role.
dvacctmgr_uname	Name of the user who will be the Database Vault Account Manager. This user will be granted the DV_ACCTMGR role. If you omit this setting, the user specified by the dvowner_uname parameter is made the Database Vault Account Manager and granted the DV_ACCTMGR role.

Usage Notes

- Only the ADMIN user can run the DBMS_CLOUD_MACADM.CONFIGURE_DATABASE_VAULT procedure.
- The DBMS_CLOUD_MACADM.CONFIGURE_DATABASE_VAULT procedure does not allow the ADMIN user to be specified as an input for the dvowner_uname or dvacctmgr_uname arguments.

Example

```
BEGIN
   DBMS_CLOUD_MACADM.CONFIGURE_DATABASE_VAULT(
       dvowner_uname         => 'adb_dbv_owner',
       dvacctmgr_uname       => 'adb_dbv_acctmgr');
END;
/
```

DISABLE_DATABASE_VAULT Procedure

This procedure disables Oracle Database Vault on Autonomous Database. To use this procedure you must have the `DV_OWNER` role.

Syntax

```
DBMS_CLOUD_MACADM.DISABLE_DATABASE_VAULT;
```

Usage Notes

After you run `DBMS_CLOUD_MACADM.DISABLE_DATABASE_VAULT` you must restart the Autonomous Database instance.

To use this procedure you must have the `DV_OWNER` role.

Example

```
EXEC DBMS_CLOUD_MACADM.DISABLE_DATABASE_VAULT;
```

DISABLE_USERMGMT_DATABASE_VAULT Procedure

This procedure disallows user management related operations for specified components on Autonomous Database with Oracle Database Vault enabled.

Syntax

```
DBMS_CLOUD_MACADM.DISABLE_USERMGMT_DATABASE_VAULT('component_name');
```

Where: *component_name* is the component name. Valid value is: `APEX`.

APEX is the Oracle APEX component.

Usage Notes

If you enable Oracle Database Vault with Autonomous Database and you want to enforce strict separation of duty to disallow user management related operations for the APEX, use the `DBMS_CLOUD_MACADM.DISABLE_USERMGMT_DATABASE_VAULT` procedure.

To use this procedure you must have the `DV_ACCTMGR` and `DV_ADMIN` roles.

Example

The following example disables user management for the APEX component:

```
EXEC DBMS_CLOUD_MACADM.DISABLE_USERMGMT_DATABASE_VAULT('APEX');
```

ENABLE_DATABASE_VAULT Procedure

This procedure enables Oracle Database Vault on Autonomous Database. To use this procedure you must have the `DV_OWNER` role.

Syntax

```
DBMS_CLOUD_MACADM.ENABLE_DATABASE_VAULT;
```

Usage Notes

After you run `DBMS_CLOUD_MACADM.ENABLE_DATABASE_VAULT` you must restart the Autonomous Database instance.

To use this procedure you must have the `DV_OWNER` role.

Example

The following example enables Oracle Database Vault:

```
BEGIN
   DBMS_CLOUD_MACADM.ENABLE_DATABASE_VAULT;
END;
/
```

ENABLE_USERMGMT_DATABASE_VAULT Procedure

This procedure allows user management with Oracle Database Vault enabled for specified components on Autonomous Database.

Syntax

```
DBMS_CLOUD_MACADM.ENABLE_USERMGMT_DATABASE_VAULT('component_name');
```

Where: *`component_name`* is the component name. Valid value is: `APEX`.

`APEX` is the Oracle APEX component.

Usage Notes

To use this procedure you must have the `DV_ACCTMGR` and `DV_ADMIN` roles.

Example

The following example enables user management for the APEX component:

```
EXEC DBMS_CLOUD_MACADM.ENABLE_USERMGMT_DATABASE_VAULT('APEX');
```

DBMS_CLOUD_MIGRATION Package

The `DBMS_CLOUD_MIGRATION` package facilitates translation of the SQL code from a non-Oracle SQL flavor to Oracle SQL.

- Summary of DBMS_CLOUD_MIGRATION Subprograms
  This table summarizes the subprograms included in the `DBMS_CLOUD_MIGRATION` package.

Summary of DBMS_CLOUD_MIGRATION Subprograms

This table summarizes the subprograms included in the `DBMS_CLOUD_MIGRATION` package.

Subprogram	Description
MIGRATE_SQL Procedure and Function	The `MIGRATE_SQL` facilitates the translation of SQL statements written in non-Oracle SQL to Oracle SQL. The supported SQL flavors are PostgreSQL, MySQL, and SQL Server.
MIGRATE_FILE Procedure	This procedure translates a file containing SQL statements from a specified source database as an input and translates the SQL statements to Oracle SQL.
ENABLE_TRANSLATION Procedure	This procedure enables real-time SQL translation and running of non-Oracle SQL statements in an Autonomous Database. The supported SQL flavors are PostgreSQL, MySQL, and SQL Server.
DISABLE_TRANSLATION Procedure	This procedure disables the SQL language translation.

- MIGRATE_SQL Procedure and Function
  The MIGRATE_SQL facilitates the translation of SQL statements written in non-Oracle SQL to Oracle SQL. This procedure is overloaded and has function and procedure variants.
- MIGRATE_FILE Procedure
  `MIGRATE_FILE` takes a file containing SQL statements from a specified source database as an input and translates the SQL statements to Oracle SQL.
- ENABLE_TRANSLATION Procedure
  The `ENABLE_TRANSLATION` procedure enables real-time SQL translation and the running of non-Oracle SQL statements in your Autonomous Database.
- DISABLE_TRANSLATION Procedure
  The `DISABLE_TRANSLATION` procedure disables the SQL language translation. An error is returned if the SQL language translation is not enabled for your session.

MIGRATE_SQL Procedure and Function

The MIGRATE_SQL facilitates the translation of SQL statements written in non-Oracle SQL to Oracle SQL. This procedure is overloaded and has function and procedure variants.

Syntax

```
DBMS_CLOUD_MIGRATION.MIGRATE_SQL (
   original_sql  IN  CLOB,
   output_sql    OUT CLOB,
   source_db     IN  VARCHAR2);
```

```
DBMS_CLOUD_MIGRATION.MIGRATE_SQL (
   original_sql   IN  CLOB,
   source_db      IN  VARCHAR2);
RETURN CLOB;
```

Parameters

Parameter	Description
original_sql	Specifies the original SQL statement in the source database that needs to be translated into Oracle SQL. This parameter is mandatory.
output_sql	Returns the translated SQL statement. This parameter is mandatory.
source_db	Specifies the name of the source database. The valid values for this parameter are POSTGRES, MYSQL, SQLSERVER. This parameter is mandatory.

Return Values

The function form of DBMS_CLOUD_MIGRATION.MIGRATE_SQL returns a CLOB that contains the translated Oracle SQL version of the ORIGINAL_SQL input.

Usage Notes

- You may encounter an error during the translation if the input SQL statement is not supported in Oracle SQL. See the following for more information:
  - Limitations for Migration and Translation of PostgreSQL **Statements** to Oracle SQL
  - Limitations for Migration and Translation of MySQL Statements to Oracle SQL
  - Limitations for Migration and Translation of SQL Server Statements to Oracle SQL
- The DBMS_CLOUD_MIGRATION.MIGRATE_SQL subprograms only accept one SQL statement as input. So, only a single SQL statement can be translated per call.
- You must be logged in as the ADMIN user or have the EXECUTE privilege on the DBMS_CLOUD_MIGRATION package to run the DBMS_CLOUD_MIGRATION.MIGRATE_SQL procedure.

Examples

Procedure form of DBMS_CLOUD_MIGRATION.MIGRATE_SQL to translate the PostgreSQL statement to Oracle SQL:

```
BEGIN
  DBMS_CLOUD_MIGRATION.MIGRATE_SQL(
    original_sql => 'SELECT e.employee_id, e.last_name, e.salary FROM
employees AS e;',
    output_sql   => output_variable,
    source_db    => 'POSTGRES');
END;
/
```

Function form of DBMS_CLOUD_MIGRATION.MIGRATE_SQL to translate the PostgreSQL statement to Oracle SQL:

```
SELECT DBMS_CLOUD_MIGRATION.MIGRATE_SQL
    ('CREATE TABLE IF NOT EXISTS cars (brand VARCHAR(255), model
VARCHAR(255), year INT)', 'POSTGRES') AS output
    FROM DUAL;

OUTPUT
--------------------------------------------------------------------------------
create table cars (brand VARCHAR2(255), model VARCHAR2(255), year NUMBER(10);
```

MIGRATE_FILE Procedure

MIGRATE_FILE takes a file containing SQL statements from a specified source database as an input and translates the SQL statements to Oracle SQL.

Syntax

```
DBMS_CLOUD_MIGRATION.MIGRATE_FILE (
    credential_name  IN VARCHAR2,
    location_uri     IN VARCHAR2,
    target_uri       IN VARCHAR2 DEFAULT NULL,
    source_db        IN VARCHAR2,
    params           IN CLOB
```

Parameters

Parameter	Description
credential_name	The name of the credential to access the Cloud Object Storage. This parameter is mandatory. Optionally, the credential you use to access your Cloud Object Store can be defined by configuring policies and roles. See Configure Policies and Roles to Access Resources for more information.
location_uri	The source file URI. The format of the URI depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD URI Formats. This parameter is mandatory.
target_uri	The target file URI, where the translated file should be stored. If you do not provide a value for this parameter, the translated file is stored at the same location as the source file. The format of the URI depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD URI Formats.
source_db	Specifies the source database language. The valid values for this parameter are POSTGRES, MYSQL, SQLSERVER. This parameter is mandatory.
params	Specifies the additional parameters that can be passed in JSON format.

Usage Notes

- To run DBMS_CLOUD_MIGRATION.MIGRATE_FILE you must be logged in as the ADMIN user or have the EXECUTE privilege on DBMS_CLOUD_MIGRATION.

- The credential_name parameter specifies the credential to access the Cloud Object Storage URI. The user that runs `DBMS_CLOUD_MIGRATION.MIGRATE_FILE`.must have `EXECUTE` privilege for the credential object used to access the Object Storage URI. That is, the credential you specify with the `credential_name` parameter.

Example

```
BEGIN
 DBMS_CLOUD_MIGRATION.MIGRATE_FILE (
     credential_name => 'OBJ_STORE_CRED',
     location_uri    => 'https://objectstorage.region.oraclecloud.com/n/
namespace/b/bucket/o/files/postgrestest.sql',
     source_db       => 'POSTGRES'
    );
END;
/
```

The `credential_name` parameter specifies the credential to access the Cloud Object Storage URI.

The `location_uri` is the the source file URI. The format of the URI depends on the Cloud Object Storage service you are using. See DBMS_CLOUD URI Formats for more information.

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

ENABLE_TRANSLATION Procedure

The `ENABLE_TRANSLATION` procedure enables real-time SQL translation and the running of non-Oracle SQL statements in your Autonomous Database.

Syntax

```
DBMS_CLOUD_MIGRATION.ENABLE_TRANSLATION (
   source_db  IN  VARCHAR2);
```

Parameters

Parameter	Description
`source_db`	Specifies the source database language. The valid values for this parameter are `POSTGRES`, `MYSQL`, `SQLSERVER`. This parameter is mandatory.

Usage Note

- To run `DBMS_CLOUD_MIGRATION.ENABLE_TRANSLATION` you must be logged in as the `ADMIN` user or have the `EXECUTE` privilege on `DBMS_CLOUD_MIGRATION`.

Example

```
BEGIN
  DBMS_CLOUD_MIGRATION.ENABLE_TRANSLATION (source_db => 'POSTGRES');
```

```
END;
/
```

DISABLE_TRANSLATION Procedure

The `DISABLE_TRANSLATION` procedure disables the SQL language translation. An error is returned if the SQL language translation is not enabled for your session.

Syntax

```
DBMS_CLOUD_MIGRATION.DISABLE_TRANSLATION();
```

Usage Note

- To run `DBMS_CLOUD_MIGRATION.DISABLE_TRANSLATION` you must be logged in as the `ADMIN` user or have the `EXECUTE` privilege on `DBMS_CLOUD_MIGRATION`.

Example

```
BEGIN
  DBMS_CLOUD_MIGRATION.DISABLE_TRANSLATION();
END;
/
```

DBMS_CLOUD_NOTIFICATION Package

The `DBMS_CLOUD_NOTIFICATION` package allows you to send messages or the output of a SQL query to a provider.

- DBMS_CLOUD_NOTIFICATION Overview
  Describes the `DBMS_CLOUD_NOTIFICATION` package.
- Summary of DBMS_CLOUD_NOTIFICATION Subprograms
  This table summarizes the subprograms included in the package.

DBMS_CLOUD_NOTIFICATION Overview

Describes the `DBMS_CLOUD_NOTIFICATION` package.

The supported providers with `DBMS_CLOUD_NOTIFICATION` are: Email, Microsoft Teams, OCI Notifications (ONS), and Slack.

When you use the OCI Notifications service with `DBMS_CLOUD_NOTIFICATION` package, only the `DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE` procedure is supported (the `DBMS_CLOUD_NOTIFICATION.SEND_DATA` procedure is not supported)

Summary of DBMS_CLOUD_NOTIFICATION Subprograms

This table summarizes the subprograms included in the package.

Subprogram	Description
SEND_DATA Procedure	Send SQL query output to a provider.
SEND_MESSAGE Procedure	Send a text message to a provider.

- SEND_DATA Procedure
  The SEND_DATA procedure sends the results of the specified query to a provider.
- SEND_MESSAGE Procedure
  The procedure sends a text message to a provider.

SEND_DATA Procedure

The SEND_DATA procedure sends the results of the specified query to a provider.

Syntax

```
DBMS_CLOUD_NOTIFICATION.SEND_DATA(
       provider           IN  VARCHAR2,
       credential_name    IN  VARCHAR2,
       query              IN  CLOB,
       params             IN  CLOB
 );
```

Parameters

Parameter	Description
provider	Specifies the provider. Valid values are: 'email', 'msteams', and 'slack' This parameter is mandatory.
credential_name	The name of the credential object to access the provider. For the email provider, the credential is the name of the credential of the SMTP approved sender that contains its username and password. For the msteams provider, the credential is the name of the credential. For the Slack provider, the credential's username must be "SLACK_TOKEN" and the password is the Slack Bot Token. This parameter is mandatory.
query	Specifies the SQL query to run. Results are sent to the provider. This parameter is mandatory.

Parameter	Description
params	Specifies specific parameters for the provider in JSON format. For the provider type email, the valid parameters are: • sender. This specifies the Email ID of the approved sender in a String value. • smtp_host. This specifies the SMTP host name in a String value. • subject. This specifies the subject of the email in a String value The maximum size is 100 characters. • recipient. This specifies the email IDs of recipients in a String value. Use a comma between email IDs when there are multiple recipients. • to_cc. This specifies the email IDs that are receiving a CC of the email. It is a String value. Use a comma between email IDs when there are multiple CC recipients. • to_bcc. This specifies the email IDs that are receiving a BCC of the email. It is a String value. Use a comma between email IDs when there are multiple BCC recipients. • message. This specifies the message text in a String value. • type. This specifies the output format in a String value as either CSV or JSON. • title. This specifies the title of the attachment of SQL output in a String value. The title should only contain letters, digits, underscores, hyphens, or dots as characters in its value due to it being used to generate a file name. For the provider type msteams, the valid parameters are: • tenant. This specifies the tenant ID in a String format. • team. This specifies the team ID in a String format. • channel. This specifies the channel ID in a String format. • title. This specifies the title of the file in a String format. The maximum size is 50 characters. • type. This specifies the output format in a String value. Valid values are CSV or JSON. For the provider type slack, the valid parameters are: • channel. This specifies the Channel ID in a String value. The Channel ID is a unique ID for a channel and is different from the channel name. In Slack, when you view channel details you can find the Channel ID on the “About” tab. • type. This specifies the output format in a String value. Valid values are CSV or JSON. This parameter is mandatory.

Usage Notes

- Use the procedure DBMS_CLOUD.CREATE_CREDENTIAL(credential_name, username, password) to create the credential object.

  See CREATE_CREDENTIAL Procedure for more information.

- For the email provider, note the following:
  - The user requires an SMTP connection endpoint for the Email Delivery server to obtain smtp_host. The user also requires an SMTP approved sender and its credentials to authenticate with the Email Delivery server to obtain the credential_name. The SMTP connection must be configured and the SMTP credentials must be generated and approved.

  - The maximum message size supported when using `DBMS_CLOUD_NOTIFICATION.SEND_DATA` with the `email` provider is 32k bytes.
  - You can only use `DBMS_CLOUD_NOTIFICATION` for mail notifications with Autonomous Database version 19.21 and above.

- For the `msteams` provider, the user requires the Microsoft Teams App and a bot configured in it. The app needs to be published to the org and installed after obtaining approval from the admin from the admin center. The user also requires access for `Files.ReadWrite.All` and `ChannelSettings.Read.All` permission for Graph API from their Azure Portal. To generate the required token, the user requires `bot_id` in the username and `bot_secret` in the password.

  The maximum file size supported when using `DBMS_CLOUD_NOTIFICATION.SEND_DATA` for Microsoft Teams is 4MB.

- For the `slack` provider, the `username` value can be any valid string and the `password` is the Slack bot token.

Examples

Send SQL output with the `email` provider:

```
BEGIN
    DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'EMAIL_CRED',
    username        => 'username',
    password        => 'password');
END;
/

BEGIN
   DBMS_CLOUD_NOTIFICATION.SEND_DATA(
        provider        => 'email',
        credential_name => 'EMAIL_CRED',
        query           => 'SELECT tablespace_name FROM dba_tablespaces',
        params          => json_object('recipient' value  'mark@oracle.com,
suresh@oracle.com',
                                       'to_cc'  value 'nicole@oracle.com1,
jordan@oracle.com',
                                       'to_bcc' value 'manisha@oracle.com',
                                       'subject' value 'Test subject',
                                       'type' value 'json',
                                       'title' value 'mytitle',
                                       'message' value 'This is the message',
                                       'smtp_host' value
'smtp.email.example.com',
                                       'sender'    value
'approver_sender@oracle.com' )
   );
   END;
/
```

Send SQL output with the `msteams` provider:

```
BEGIN
    DBMS_CLOUD.CREATE_CREDENTIAL(credential_name => 'TEAMS_CRED',
```

```
        username        => 'bot_id',
        password        => 'bot_secret');
END;
/

BEGIN
     DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE(provider => 'msteams',
        credential_name => 'TEAMS_CRED',
        query           => 'SELECT tablespace_name FROM dba_tablespaces',
        params          => json_object('tenant'value '5b743bc******c0286',
                                       'team'value '0ae401*********5d2bd',
                                       'channel'value
'19%3a94be023*****%40thread.tacv2',
                                       'title'value 'today',
                                       'type'value 'csv'));
END;
/
```

Send SQL output with the `slack` provider:

```
BEGIN
   DBMS_CLOUD.CREATE_CREDENTIAL(
     credential_name => 'SLACK_CRED',
     username        => 'SLACK_TOKEN',
     password        => 'password');
END;
/
BEGIN
  DBMS_CLOUD_NOTIFICATION.SEND_DATA(
      provider        => 'slack',
      credential_name => 'SLACK_CRED',
      query           => 'SELECT username, account_status, expiry_date FROM
USER_USERS WHERE rownum < 5',
      params          => json_object('channel' value 'C0....08','type' value
'csv'));
END;
/
```

SEND_MESSAGE Procedure

The procedure sends a text message to a provider.

Syntax

```
DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE(
       provider           IN  VARCHAR2,
       credential_name    IN  VARCHAR2,
       message            IN  CLOB,
       params             IN  CLOB
 );
```

Parameters

Parameter	Description
`provider`	Specifies the provider. Valid values are: '`email`', '`msteams`', '`slack`', and '`oci`'. This parameter is mandatory.
`credential_name`	The name of the credential object to access the provider. • `email` provider: the credential is the name of the credential of the SMTP approved sender that contains its username and password. • `msteams` provider: the credential must contain the `bot_id` and the `bot_secret` key in both the `username` and the `password`. • `slack` provider: the credential's `username` must be "`SLACK_TOKEN`" and the `password` is the Slack Bot Token. • `oci` provider: the credential is either a valid Oracle Cloud Infrastructure signing key based credential or resource principal credential. This parameter is mandatory. Use the procedure `DBMS_CLOUD.CREATE_CREDENTIAL` to create the credential object. See CREATE_CREDENTIAL Procedure for more information.
`message`	Specifies the message text. This parameter is mandatory.
`params`	Specifies specific parameters for the `provider` in JSON format. For the provider type `email`, the valid parameters are: • `sender`: Specifies the Email ID of the approved sender in a `String` value. • `smtp_host`: Specifies the SMTP host name in a `String` value. • `subject`: Specifies the subject of the email in a `String` value. The maximum size is 100 characters. • `recipient`: Specifies the email IDs of recipients in a `String` value. Use a comma between email IDs when there are multiple recipients. • `to_cc`: Specifies the email IDs that are receiving a CC of the email. Use a comma between email IDs when there are multiple CC recipients. • `to_bcc`: Specifies the email IDs that are receiving a BCC of the email. Use a comma between email IDs when there are multiple BCC recipients. For the provider type `msteams`, the valid parameter is: • `channel`: Specifies the Channel ID in a `String` value. For the provider type `slack`, the valid parameter is: • `channel`. Specifies the Channel ID in a `String` value. This parameter is mandatory. The Channel ID is a unique ID for a channel and is different from the channel name. In Slack, when you view channel details you can find the Channel ID on the “About” tab. For the provider OCI Notifications with the type `oci`, the valid parameters are: • `topic_ocid`: Specifies • `title`: Specifies for the subject line of the message. With the `oci` provider, there are different possible endpoints for a message, based on the subscriptions. The `title` parameter is treated differently for different subscriptions: • email: The `title` specifies the subject line of the message. • Slack: The `title` value is not used. If it is included the value is ignored.

Usage Note for Email Provider

For the `email` provider, the user requires an SMTP connection endpoint for the Email Delivery server to obtain `smtp_host`. The user also requires an SMTP approved sender and its credentials to authenticate with the Email Delivery server to obtain the `credential_name`. SMTP connection must be configured, and SMTP credentials must be generated and approved.

Usage Note for the Msteams Provider

For the `msteams` provider, the user requires the Microsoft Teams App and a bot configured in it. The app needs to be published to the org and installed after obtaining approval from the admin from the admin center. The user also requires access for `Files.ReadWrite.All` and `ChannelSettings.Read.All` permission for Graph API from their Azure Portal. To generate the required token, the user requires `bot_id` in the username and `bot_secret` in the password.

Usage Note for the Slack Provider

For the `slack` provider, the `username` value can be any valid string and the `password` is the Slack bot token.

Usage Notes for the OCI Provider OCI Notifications

The `DBMS_CLOUD_NOTIFICATION` supports the OCI Notifications provider. See Overview of Notifications for details on the OCI Notifications provider.

- To use the OCI Notifications provider, you must do the following:
  - Create a topic. See Creating a Topic for details on creating a topic
  - Create the required subscriptions. See Creating a Subscription for details on creating a subscription.
  - After you create the subscriptions, confirm the subscriptions. Messages are only published to confirmed subscriptions. See Confirming a Subscription for details on confirming subscriptions.
- The the Autonomous Database instance, the credential, if you are using user credentials, and the subscription topic must all be located in the same region.
- Note the following limitations:
  - Message Size Limit: The message size limit per request is 64KB. Messages exceeding this size cannot be sent or processed within a single request.
  - Message Delivery Rate Limit: There are different message delivery rate limits depending on the protocol used.
    * `email` protocol: the limit is 10 messages per minute per endpoint.

    This restriction may impact the speed at which messages can be transmitted and received.
  - Transactions Per Minute (TPM) per-Tenancy Limit: There is a TPM limit of 60 per topic for this operation. This limit represents the maximum number of messages per minute that can be processed or handled for a given topic. Exceeding this limit may result in messages not being processed or processed at a slower rate.

Examples

Send a text message with the `email` provider:

```
BEGIN
    DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'EMAIL_CRED',
    username        => 'username',
    password        => 'password');
END;
/

BEGIN
   DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE(
        provider        => 'email',
        credential_name => 'EMAIL_CRED',
        message         => 'Subject content',
        params          => json_object('recipient' value  'mark@oracle.com,
suresh@oracle.com',
                                       'to_cc'  value 'nicole@oracle.com,
jordan@oracle.com',
                                       'to_bcc' value 'manisha@oracle.com',
                                       'subject' value 'Test subject',
                                       'smtp_host' value
'smtp.email.example.com',
                                       'sender'    value
'approver_sender@oracle.com' )
   );
   END;
/
```

Send a text message with the `msteams` provider:

```
BEGIN
    DBMS_CLOUD.CREATE_CREDENTIAL(credential_name => 'TEAMS_CRED',
        username        => 'bot_id',
        password        => 'bot_secret');
END;
/

BEGIN
     DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE(
        provider        => 'msteams',
        credential_name => 'TEAMS_CRED',
        message         => 'text from new teams api',
        params          => json_object('channel'value
'C0....08'),'region'value 'india');
END;
/
```

Send a text message with the `slack` provider:

```
BEGIN
   DBMS_CLOUD.CREATE_CREDENTIAL(
     credential_name => 'SLACK_CRED',
     username        => 'SLACK_TOKEN',
     password        => 'password');
END;
/
BEGIN
   DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE(
      provider          => 'slack',
      credential_name   => 'SLACK_CRED',
      message           => 'Send text from Autonomous Database.',
      params            => json_object('channel' value 'C0....08'));
END;
/
```

Send a text message with the OCI Notifications, type `oci` provider:

```
BEGIN
   DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE(
      provider          => 'oci',
      credential_name  => 'OCI_CRED',
      message          => 'Text message for you.',
      params           => json_object('topic_ocid'  value 'oci********pa',
      'title'  value 'Title for message subject' ));
END;
/
```

```
BEGIN
   DBMS_CLOUD_NOTIFICATION.SEND_MESSAGE(
      provider          => 'oci',
      credential_name  => 'OCI$RESOURCE_PRINCIPAL',
      message          => 'Text message for you.',
      params           => json_object('topic_ocid'  value 'oci********pa',
      'title'  value 'Title for message subject' ));
END;
/
```

DBMS_CLOUD_PIPELINE Package

The `DBMS_CLOUD_PIPELINE` package allows you to create data pipelines for loading and exporting data in the cloud. This package supports continuous incremental data load of files in object store into the database. `DBMS_CLOUD_PIPELINE` also supports continuous incremental export of table data or query results from the database to object store based on a timestamp column.

- Summary of DBMS_CLOUD_PIPELINE Subprograms
  This table summarizes the subprograms included in the `DBMS_CLOUD_PIPELINE` package.

- DBMS_CLOUD_PIPELINE Attributes
  Attributes help to control and configure the behavior of a data pipeline.
- DBMS_CLOUD_PIPELINE Views
  The `DBMS_CLOUD_PIPELINE` package uses the following views.

Summary of DBMS_CLOUD_PIPELINE Subprograms

This table summarizes the subprograms included in the `DBMS_CLOUD_PIPELINE` package.

Subprogram	Description
CREATE_PIPELINE Procedure	Creates a new data pipeline.
DROP_PIPELINE Procedure	Drops an existing data pipeline.
RESET_PIPELINE Procedure	Resets the tracking state of a data pipeline. Use reset pipeline to restart the pipeline from the initial state of data load or export. Optionally reset pipeline can purge data in the database or in object store, depending on the type of pipeline.
RUN_PIPELINE_ONCE Procedure	Performs an on-demand run of the pipeline in the current foreground session, instead of a scheduled job.
SET_ATTRIBUTE Procedure	Sets pipeline attributes. There are two overloaded procedures, one to set a single attribute and another to set multiple attributes using a JSON document of attribute name/value pairs
START_PIPELINE Procedure	Starts the data pipeline. When a pipeline is started, the pipeline operation will continuously run in a scheduled job according to the "interval" configured in pipeline attributes.
STOP_PIPELINE Procedure	Stops the data pipeline. When a pipeline is stopped, no future jobs are scheduled for the pipeline.

- CREATE_PIPELINE Procedure
  The procedure creates a new data pipeline.
- DROP_PIPELINE Procedure
  The procedure drops an existing data pipeline. If a pipeline has been started, then it must be stopped before it can be dropped.
- RESET_PIPELINE Procedure
  Resets the tracking state of a data pipeline. Use reset pipeline to restart the pipeline from the initial state of data load or export. Optionally reset pipeline can purge data in the database or in object store, depending on the type of pipeline. A data pipeline must be in stopped state to reset it.
- RUN_PIPELINE_ONCE Procedure
  This procedure performs an on-demand run of the pipeline in the current foreground session, instead of a running in a scheduled job. Use `DBMS_CLOUD_PIPELINE.RUN_PIPELINE_ONCE` to test a pipeline before you start the pipeline as a continuous job.
- SET_ATTRIBUTE Procedure
  This procedure sets pipeline attributes. There are two overloaded procedures, one to set a single attribute and another to set multiple attributes using a JSON document of attribute name/value pairs.
- START_PIPELINE Procedure
- STOP_PIPELINE Procedure
  The procedure stops the data pipeline. When a pipeline is stopped, no future jobs are scheduled for the pipeline.

CREATE_PIPELINE Procedure

The procedure creates a new data pipeline.

Syntax

```
DBMS_CLOUD_PIPELINE.CREATE_PIPELINE(
      pipeline_name        IN   VARCHAR2,
      pipeline_type        IN   VARCHAR2,
      attributes           IN   CLOB        DEFAULT NULL,
      description          IN   VARCHAR2    DEFAULT NULL
);
```

Parameters

Parameter	Description
`pipeline_name`	Specifies a name for the pipeline. The pipeline name must follow the naming rules of Oracle SQL identifiers. See Identifiers for more information. This parameter is mandatory.
`pipeline_type`	Specifies the pipeline type. Valid values: `LOAD`, `EXPORT` This parameter is mandatory.
`attributes`	Pipeline attributes in JSON format. Default value: `NULL` See DBMS_CLOUD_PIPELINE Attributes for more information.
`description`	Description for the pipeline. Default value: `NULL`

DROP_PIPELINE Procedure

The procedure drops an existing data pipeline. If a pipeline has been started, then it must be stopped before it can be dropped.

Syntax

```
DBMS_CLOUD_PIPELINE.DROP_PIPELINE(
       pipeline_name        IN   VARCHAR2,
       force                IN   BOOLEAN DEFAULT FALSE
 );
```

Parameters

Parameter	Description
`pipeline_name`	Specifies a pipeline name. This parameter is mandatory.
`force`	Forcibly drop a pipeline, even if it is in started state. Valid values: `TRUE`, `FALSE` Default value: `FALSE`

Usage Note

- In order to drop a pipeline that is in started state, set the `force` parameter to `TRUE`.

RESET_PIPELINE Procedure

Resets the tracking state of a data pipeline. Use reset pipeline to restart the pipeline from the initial state of data load or export. Optionally reset pipeline can purge data in the database or in object store, depending on the type of pipeline. A data pipeline must be in stopped state to reset it.

Syntax

```
DBMS_CLOUD_PIPELINE.RESET_PIPELINE(
       pipeline_name        IN   VARCHAR2,
       purge_data           IN   BOOLEAN DEFAULT FALSE
 );
```

Parameters

Parameter	Description
`pipeline_name`	Specifies a name for the pipeline. This parameter is mandatory.
`purge_data`	Purge data applies for either a load pipeline or an export pipeline: • For a load pipeline, when `TRUE`, truncate the data in database table. • For an export pipeline, when `TRUE`, delete the files in the object store location. Valid values: `TRUE`, `FALSE` Default value: `FALSE`

Usage Notes

- A data pipeline must be in stopped state to reset it. See STOP_PIPELINE Procedure for more information.
- For a load pipeline, resetting the pipeline clears the record of the files being loaded by the pipeline. When you call either `START_PIPELINE` or `RUN_PIPELINE_ONCE` after resetting a load pipeline, the pipeline repeats the data load and includes all the files present in the object store location.

  When `purge_data` is set to `TRUE`, `DBMS_CLOUD_PIPELINE.RESET_PIPELINE` does the following:

  – Truncates the data in the pipeline's database table you specify with `table_name` attribute.

  – Drops the pipeline's status table, and the pipeline's bad file table and error table (if they exist).
- For an export pipeline, resetting the pipeline clears the last tracked data in the database table. When you call either `START_PIPELINE` or `RUN_PIPELINE_ONCE` after resetting an export pipeline, the pipeline repeats exporting data from the table or query.

  When `purge_data` set to `TRUE`, `DBMS_CLOUD_PIPELINE.RESET_PIPELINE` deletes existing files in the object store location specified with the `location` attribute.

RUN_PIPELINE_ONCE Procedure

This procedure performs an on-demand run of the pipeline in the current foreground session, instead of a running in a scheduled job. Use `DBMS_CLOUD_PIPELINE.RUN_PIPELINE_ONCE` to test a pipeline before you start the pipeline as a continuous job.

Syntax

```
DBMS_CLOUD_PIPELINE.RUN_PIPELINE_ONCE(
       pipeline_name   IN  VARCHAR2
);
```

Parameters

Parameter	Description
`pipeline_name`	Specifies a name for the pipeline to run. This parameter is mandatory.

Usage Notes

- After you perform a test run of a pipeline you can reset the pipeline state using `DBMS_CLOUD_PIPELINE.RESET_PIPELINE`. This allows you to reset the pipeline state before you start the pipeline in a scheduled job.
- If a pipeline is in the started state, then it cannot be run in the foreground session.

SET_ATTRIBUTE Procedure

This procedure sets pipeline attributes. There are two overloaded procedures, one to set a single attribute and another to set multiple attributes using a JSON document of attribute name/value pairs.

Syntax

```
PROCEDURE DBMS_CLOUD_PIPELINE.SET_ATTRIBUTE(
      pipeline_name         IN   VARCHAR2,
      attribute_name        IN   VARCHAR2,
      attribute_value       IN   CLOB
);

PROCEDURE DBMS_CLOUD_PIPELINE.SET_ATTRIBUTE(
      pipeline_name         IN   VARCHAR2,
      attributes            IN   CLOB
);
```

Parameters

Parameter	Description
`pipeline_name`	Specifies a name for the pipeline to set attributes. This parameter is mandatory.

Parameter	Description
attribute_name	Specifies the attribute name for the attribute to be set. See DBMS_CLOUD_PIPELINE Attributes for more information.
attribute_value	Specifies the value for the pipeline attribute to set. See DBMS_CLOUD_PIPELINE Attributes for more information.
attributes	Specifies a JSON document containing attribute names and values. See DBMS_CLOUD_PIPELINE Attributes for more information.

Usage Note

- When you use DBMS_CLOUD_PIPELINE.SET_ATTRIBUTE to set multiple attributes with the attributes parameter, all the existing attributes are deleted and overwritten with the specified attributes from the JSON document.

START_PIPELINE Procedure

The procedure starts the data pipeline. When a pipeline is started, the pipeline operation runs continuously in a scheduled job according to the interval configured with the pipeline attributes.

Syntax

```
DBMS_CLOUD_PIPELINE.START_PIPELINE(
      pipeline_name         IN   VARCHAR2,
      start_date            IN   TIMESTAMP WITH TIME ZONE DEFAULT NULL
);
```

Parameters

Parameter	Description
pipeline_name	Specifies a name for the pipeline. This parameter is mandatory.
start_date	Specifies the starting date for the pipeline job. Default value: NULL

Usage Notes

- By default, a pipeline job begins immediately, as soon as the pipeline is started. To start a pipeline job at a later time specify a valid date or timestamp using the start_date parameter.
- See DBMS_CLOUD_PIPELINE Attributes for information on the pipeline interval and other pipeline attributes.

STOP_PIPELINE Procedure

The procedure stops the data pipeline. When a pipeline is stopped, no future jobs are scheduled for the pipeline.

Syntax

```
DBMS_CLOUD_PIPELINE.STOP_PIPELINE(
      pipeline_name        IN   VARCHAR2,
      force                IN   BOOLEAN  DEFAULT  FALSE
);
```

Parameters

Parameter	Description
`pipeline_name`	Specifies a name for the pipeline. This parameter is mandatory.
`force`	If force parameter is passed as `TRUE`, then it will terminate any running jobs for the pipeline. Valid values: `TRUE, FALSE` Default value: `FALSE`

DBMS_CLOUD_PIPELINE Attributes

Attributes help to control and configure the behavior of a data pipeline.

Attributes

> **Note:**
>
> As indicated in the **Pipeline Type** column, depending on the pipeline type `LOAD` or `EXPORT`, a pipeline supports a different set of attributes.

Attribute Name	Description	Pipeline Type	Modifiable After Pipeline Starts
`credential_name`	The name of the credential to access the source Cloud Object Storage. You can use `'OCI$RESOURCE_PRINCIPAL'` as the `credential_name` when resource principal is enabled. Credentials specified with the following are also supported: Amazon Resource Names, Azure Service Principal or Google Service Account. See Configure Policies and Roles to Access Resources for more information. Default value: `NULL`. If you do not supply a `credential_name`, `credential_name` is set to `NULL`.	`LOAD`, `EXPORT`	Yes

Attribute Name	Description	Pipeline Type	Modifiable After Pipeline Starts
`field_list`	Identifies the fields in the source files and their data types. This argument's syntax is the same as the `field_list` clause in regular Oracle external tables. For more information about `field_list` see *Oracle® Database Utilities*. Default value: `NULL` The default value specifies the fields and their data types are determined by the columns in the table specified in `table_name` attribute.	`LOAD`	Yes
`format`	The options describing the format for the type of pipeline. • For a load pipeline, see DBMS_CLOUD Package Format Options. • For an export pipeline, see DBMS_CLOUD Package Format Options for EXPORT_DATA Datapump `format` is not supported for an export pipeline. This attribute is mandatory for both `LOAD` and `EXPORT` pipelines.	`LOAD`, `EXPORT`	Yes
`interval`	The time interval in minutes between consecutive executions of scheduled pipeline job. Default value: 15 minutes	`LOAD`, `EXPORT`	Yes
`key_column`	A timestamp or date column in the specified `table` or `query` for exporting newer data continuously to object store. The last execution timestamp or date is tracked by Export pipeline and compared with the value in the `key_column` to identify new data for exporting to object store. Default value: NULL If `key_column` is not specified for an export pipeline, the entire contents of the `table` or `query` are uploaded to object store in each pipeline job execution.	`EXPORT`	No
`location`	Specifies a URI that points to an Object Storage location. The format of the URI depends on the Cloud Object Storage service. See DBMS_CLOUD URI Formats for more information. This attribute is mandatory for both `LOAD` and `EXPORT` pipelines.	`LOAD`, `EXPORT`	No

Attribute Name	Description	Pipeline Type	Modifiable After Pipeline Starts
priority	Specifies a string value that determines the number of parallel operations performed for the pipeline. • In a load pipeline, it determines the number of files loaded in parallel. • In an export pipeline, it determines the degree of parallelism for fetching data from the database. An operation with a higher priority consumes more database resources and is completed sooner. Valid values: • HIGH: Determines the number of parallel files handled using the database's **ECPU count** (**OCPU count** if your database uses OCPUs). • MEDIUM: Determines the number of simultaneous processes using the concurrency limit for Medium service. The default value is 4. • LOW: Run the pipeline job in serial order. Default value: MEDIUM The maximum number of concurrent file operations is limited to 64.	LOAD, EXPORT	Yes
query	Specifies a SELECT statement so that only the required data is exported. The query determines the contents of the files you export as text files (CSV, JSON, Parquet, or XML) or dump files. For example: `SELECT warehouse_id, quantity FROM inventories` Default value: NULL For an export pipeline, either table_name or query is mandatory.	EXPORT	No
table_name	Specifies the name of the target table for loading or exporting data. For a load pipeline table_name is mandatory. For an export pipeline, either table_name or query is mandatory.	LOAD, EXPORT	No
table_owner	The name of the schema where the target table resides for loading or exporting data. Default value: NULL With a NULL value the target table is in the same schema as the user running the procedure.	LOAD, EXPORT	No

DBMS_CLOUD_PIPELINE Views

The DBMS_CLOUD_PIPELINE package uses the following views.

- DBA_CLOUD_PIPELINES View

- USER_CLOUD_PIPELINES View
- DBA_CLOUD_PIPELINE_HISTORY View
- USER_CLOUD_PIPELINE_HISTORY View
- DBA_CLOUD_PIPELINE_ATTRIBUTES View
- USER_CLOUD_PIPELINE_ATTRIBUTES View

DBMS_CLOUD_REPO Package

The `DBMS_CLOUD_REPO` package provides for use of and management of cloud hosted code repositories from Oracle Database. Supported cloud code repositories include GitHub, AWS CodeCommit , and Azure Repos.

- DBMS_CLOUD_REPO Overview
  The `DBMS_CLOUD_REPO` package provides easy access to files in Cloud Code (Git) Repositories, including: GitHub, AWS CodeCommit, and Azure Repos.
- DBMS_CLOUD_REPO Data Structures
  The `DBMS_CLOUD_REPO` package defines record types and a generic JSON object type `repo`.
- DBMS_CLOUD_REPO Subprogram Groups
  The `DBMS_CLOUD_REPO` package subprograms can be grouped into four categories: Initialization Operations, Repository Management Operations, File Operations, and SQL Install Operations.
- Summary of DBMS_CLOUD_REPO Subprograms
  This section covers the `DBMS_CLOUD_REPO` subprograms provided with Autonomous Database.

DBMS_CLOUD_REPO Overview

The `DBMS_CLOUD_REPO` package provides easy access to files in Cloud Code (Git) Repositories, including: GitHub, AWS CodeCommit, and Azure Repos.

This package is a single interface for access to Multicloud Code repositories and allows you to upload SQL files to Git repositories or install SQL scripts directly from Cloud Code Repositories. This package also allows you to use a Cloud Code Repository to manage code versions for SQL scripts and to install or patch application code from Git repositories.

Concepts

- Git Version Control System: Git is software for tracking changes in any set of files, usually used for coordinating work among programmers collaboratively developing source code during software development. Its goals include speed, data integrity, and support for distributed, non-linear workflows.
- Git Repository: A Git repository is a virtual storage of your project. It allows you to save versions of your code, which you can access when needed.

Architecture

`DBMS_CLOUD_REPO` package provides four feature areas:

- Repository Initialization with Generic Cloud Code Repository Handle
  - Initialize a GitHub Code Repository
  - Initialize an AWS CodeCommit Code Repository

  - Initialize an Azure Repos Code Repository
- Repository Management Operations
  - Create a repository
  - Update a repository
  - List repositories
  - Delete a repository
- Repository File Management Operations
  - Upload a file to Code Repository from Oracle Database.
  - Download a file from Code Repository to Oracle Database.
  - Delete files from Code Repository.
  - List files from Code Repository.
- SQL Install Operations
  - Export Database object metadata DDL to repository.
  - Install SQL statements from a file in the Code Repository in Oracle Database.
  - Install SQL statements from a buffer.

DBMS_CLOUD_REPO Data Structures

The `DBMS_CLOUD_REPO` package defines record types and a generic JSON object type `repo`.

REPO JSON Object

A `DBMS_CLOUD_REPO REPO` is an opaque JSON object to represent a Cloud Code Repository of a specific cloud provider. A REPO object can be passed to different `DBMS_CLOUD_REPO` APIs. This opaque object ensures that `DBMS_CLOUD_REPO` procedures and functions are multicloud compatible; you do not have to change any code when you migrate from one Cloud Code Repository provider to another Cloud Code Repository.

DBMS_CLOUD_REPO Subprogram Groups

The `DBMS_CLOUD_REPO` package subprograms can be grouped into four categories: Initialization Operations, Repository Management Operations, File Operations, and SQL Install Operations.

- DBMS_CLOUD_REPO Initialization Operations
  Lists the subprograms for initialization operations within the `DBMS_CLOUD_REPO` package.
- DBMS_CLOUD_REPO Repository Management Operations
  Shows the subprograms for repository management operations within the `DBMS_CLOUD_REPO` package.
- DBMS_CLOUD_REPO Repository Branch Management Operations
  Lists the subprograms for repository branch management operations within the `DBMS_CLOUD_REPO` package.
- DBMS_CLOUD_REPO File Operations
  Lists the subprograms for file operations within the `DBMS_CLOUD_REPO` package.
- DBMS_CLOUD_REPO SQL Install Operations
  Lists the subprograms for SQL install operations within the `DBMS_CLOUD_REPO` package.

DBMS_CLOUD_REPO Initialization Operations

Lists the subprograms for initialization operations within the `DBMS_CLOUD_REPO` package.

Subprogram	Description
INIT_AWS_REPO Function	This function initializes an AWS repository handle and returns an opaque type.
INIT_AZURE_REPO Function	This function initializes an Azure repository handle and returns an opaque type.
INIT_GITHUB_REPO Function	This function initializes a GitHub repository handle and returns an opaque type.
INIT_REPO Function	This function initializes a Cloud Code Repository handle and returns an opaque JSON object.

DBMS_CLOUD_REPO Repository Management Operations

Shows the subprograms for repository management operations within the `DBMS_CLOUD_REPO` package.

Subprogram	Description
CREATE_REPOSITORY Procedure	This procedure creates a Cloud Code Repository identified by the `repo` handle argument.
DELETE_REPOSITORY Procedure	This procedure deletes the Cloud Code Repository identified by the `repo` handle argument.
LIST_REPOSITORIES Function	This function lists all the Cloud Code Repositories identified by the `repo` handle argument.
UPDATE_REPOSITORY Procedure	This procedure updates a Cloud Code repository identified by the `repo` handle argument. The procedure supports updating the name, description, or private visibility status, as supported by the Cloud Code repository.

DBMS_CLOUD_REPO Repository Branch Management Operations

Lists the subprograms for repository branch management operations within the `DBMS_CLOUD_REPO` package.

Subprogram	Description
CREATE_BRANCH Procedure	This procedure creates a branch in a Cloud Code Repository identified by the `repo` handle argument.
DELETE_BRANCH Procedure	This procedure deletes a branch in a Cloud Code Repository identified by the `repo` handle argument.
LIST_BRANCHES Function	This function lists all the Cloud Code Repository branches identified by the `repo` handle argument.
LIST_COMMITS Function	This function lists all the commits in a Cloud Code Repository branch identified by the `repo` handle argument.
MERGE_BRANCH Procedure	This procedure merges a Cloud Code Repository branch into another specified branch in a Cloud Code Repository identified by the `repo` handle argument.

DBMS_CLOUD_REPO File Operations

Lists the subprograms for file operations within the `DBMS_CLOUD_REPO` package.

Subprogram	Description
DELETE_FILE Procedure	This procedure deletes a file from the Cloud Code repository identified by the `repo` handle argument.
GET_FILE Procedure and Function	The function downloads the contents of a file from the Cloud Code repository. The procedure allows you to download the contents of a file from the Cloud Code repository and save the file in a directory.
LIST_FILES Function	This function downloads a file from Cloud Code repository. Optionally, file content can be accessed from either a specific branch, tag or commit name. By default, the file is accessed from the default repository branch.
PUT_FILE Procedure	This procedure uploads a file to the Cloud Code repository identified by the `repo` handle argument. The procedure is overloaded to support either uploading a file from a directory object or uploading the contents from a CLOB to the repository file.

DBMS_CLOUD_REPO SQL Install Operations

Lists the subprograms for SQL install operations within the `DBMS_CLOUD_REPO` package.

Subprogram	Description
EXPORT_OBJECT Procedure	This procedure uploads the DDL metadata of a database object to the Cloud Code repository identified by the `repo` handle argument.
EXPORT_SCHEMA Procedure	This procedure exports metadata of all objects in a schema to a Cloud Code Repository branch identified by the `repo` handle argument.
INSTALL_FILE Procedure	This procedure installs SQL statements from a file in the Cloud Code repository identified by the `repo` handle argument.
INSTALL_SQL Procedure	This procedure installs SQL statements from a buffer given as input.

Summary of DBMS_CLOUD_REPO Subprograms

This section covers the `DBMS_CLOUD_REPO` subprograms provided with Autonomous Database.

The `DBMS_CLOUD_REPO` package is made up of the following:

- DBMS_CLOUD_REPO Initialization Operations
- DBMS_CLOUD_REPO Repository Management Operations
- DBMS_CLOUD_REPO File Operations
- DBMS_CLOUD_REPO SQL Install Operations
- CREATE_BRANCH Procedure
  This procedure creates a branch in the Cloud Code Repository identified by the `repo` handle argument.
- CREATE_REPOSITORY Procedure
  This procedure creates a Cloud Code Repository identified by the `repo` handle argument.
- DELETE_BRANCH Procedure
  This procedure deletes a branch in the Cloud Code repository identified by the `repo` handle argument.
- DELETE_FILE Procedure
  This procedure deletes a file from the Cloud Code repository identified by the `repo` handle argument.

- DELETE_REPOSITORY Procedure
  This procedure deletes the Cloud Code Repository identified by the `repo` handle argument.
- EXPORT_OBJECT Procedure
  This procedure uploads the DDL metadata of a database object to the Cloud Code repository identified by the `repo` handle argument. This procedure is an easy way to upload the metadata definition of a database object in single step.
- EXPORT_SCHEMA Procedure
  This procedure exports metadata of all objects in a schema to the Cloud Code Repository branch identified by the `repo` handle argument.
- GET_FILE Procedure and Function
  The function downloads the contents of a file from the Cloud Code repository. The procedure allows you to download the contents of a file from the Cloud Code repository and save the file in a directory.
- INIT_AWS_REPO Function
  This function initializes an AWS repository handle and returns an opaque type.
- INIT_AZURE_REPO Function
  This function initializes an Azure repository handle and returns an opaque type. This function is only supported for Azure cloud provider.
- INIT_GITHUB_REPO Function
  This function initializes a GitHub repository handle and returns an opaque type.
- INIT_REPO Function
  This function initializes a Cloud Code Repository handle and returns an opaque JSON object. This function is a generic interface to accept a JSON document, and avoids having to change code, you only need to change a JSON document, when moving a code repository from one Cloud Code repository to another Cloud Code repository.
- INSTALL_FILE Procedure
  This procedure installs SQL statements from a file in the Cloud Code repository identified by the `repo` handle argument.
- INSTALL_SQL Procedure
  This procedure installs SQL statements from a buffer given as input.
- LIST_BRANCHES Function
  This function lists branches in the Cloud Code Repository branch identified by the `repo` handle argument.
- LIST_COMMITS Function
  This function lists commits in the Cloud Code Repository branch identified by the `repo` handle argument.
- LIST_FILES Function
  This function downloads a file from Cloud Code repository. Optionally, file content can be accessed from either a specific branch, tag or commit name. By default, the file is accessed from the default repository branch. The results include the file names and additional metadata about the files.
- LIST_REPOSITORIES Function
  This function lists all Cloud Code Repositories identified by the `repo` handle argument. If the `repo` handle has a repository name specified, the function does not restrict the listing to the specified repository name; it lists all repositories accessible to the user.

- MERGE_BRANCH Procedure
  This procedure merges a repository branch into another specified branch in the Cloud Code Repository identified by the `repo` handle argument. The `MERGE_BRANCH` procedure is currently not supported in Azure.
- PUT_FILE Procedure
  This procedure uploads a file to the Cloud Code repository identified by the `repo` handle argument. The procedure is overloaded to support either uploading a file from a directory object or uploading the contents from a BLOB to the repository file.
- UPDATE_REPOSITORY Procedure
  This procedure updates a Cloud Code repository identified by the `repo` handle argument. UPDATE_REPOSITORY supports updating the name, description, or private visibility status, as supported by the Cloud Code repository.

CREATE_BRANCH Procedure

This procedure creates a branch in the Cloud Code Repository identified by the `repo` handle argument.

Syntax

```
PROCEDURE DBMS_CLOUD_REPO.CREATE_BRANCH(
    repo               IN   CLOB,
    branch_name        IN   VARCHAR2,
    parent_branch_name IN   VARCHAR2 DEFAULT NULL,
    parent_commit_id   IN   VARCHAR2 DEFAULT NULL
 );
```

Parameters

Parameter	Description
`repo`	Specifies the repository handle. This parameter is mandatory and supported for all cloud providers.
`branch_name`	Specifies the repository branch name. This parameter is mandatory and supported for all cloud providers.
`parent_branch_name`	Creates the new branch using the head commit of the specified parent branch. This parameter is supported for all cloud providers. If you do not supply a `parent_branch_name` value, the `parent_branch_name` is set to `main`.
`parent_commit_id`	Creates the new branch using the specified repository commit. This parameter is supported for all cloud providers. If you do not supply a `parent_commit_id` value, the `parent_commit_id` is set to a NULL value.

> **Note:**
>
> To create a branch in a Cloud Code repository, you must specify the parent branch or the parent commit id.

Example

```
BEGIN
  DBMS_CLOUD_REPO.CREATE_BRANCH (
    repo                 => l_repo,
    branch_name          => 'test_branch',
    parent_branch_name   => 'main'
    );
END;
/
```

Usage Note

To run the `DBMS_CLOUD_REPO.CREATE_BRANCH` procedure, you must be logged in as the ADMIN user or have `EXECUTE` privilege on `DBMS_CLOUD_REPO`.

CREATE_REPOSITORY Procedure

This procedure creates a Cloud Code Repository identified by the `repo` handle argument.

Syntax

```
PROCEDURE DBMS_CLOUD_REPO.CREATE_REPOSITORY(
      repo                 IN   CLOB,
      description          IN   CLOB     DEFAULT NULL,
      private              IN   BOOLEAN  DEFAULT TRUE
);
```

Parameters

Parameter	Description
`repo`	Specifies the repository handle. This parameter is supported for all cloud providers.
`description`	A short text description for the repository. This parameter is supported for GITHUB and AWS cloud provider.
`private`	Repository is private and only accessible with valid credentials This parameter is only supported for the GITHUB cloud provider.

Example

```
BEGIN
  DBMS_CLOUD_REPO.CREATE_REPOSITORY(
    repo        => l_repo,
    description => 'My test repo',
    private => TRUE
    );
END;
/
```

DELETE_BRANCH Procedure

This procedure deletes a branch in the Cloud Code repository identified by the `repo` handle argument.

Syntax

```
PROCEDURE DBMS_CLOUD_REPO.DELETE_BRANCH (
      repo              IN  CLOB,
      branch_name       IN  VARCHAR2  DEFAULT NULL
 );
```

Parameters

Parameter	Description
`repo`	Specifies the repository handle. This parameter is mandatory and supported for all cloud providers.
`branch_name`	Delete branch from a specific repository. This parameter is mandatory and supported for all cloud providers.

Example

```
BEGIN
  DBMS_CLOUD_REPO.DELETE_BRANCH (
      repo        => l_repo,
      branch_name => 'test_branch'
  );
END;
/
```

Usage Note

To run the `DBMS_CLOUD_REPO.DELETE_BRANCH` procedure, you must be logged in as the ADMIN user or have `EXECUTE` privilege on `DBMS_CLOUD_REPO`.

DELETE_FILE Procedure

This procedure deletes a file from the Cloud Code repository identified by the `repo` handle argument.

Syntax

```
PROCEDURE DBMS_CLOUD_REPO.DELETE_FILE(
      repo              IN  CLOB,
      file_path         IN  VARCHAR2,
      branch_name       IN  VARCHAR2  DEFAULT NULL,
      commit_details    IN  CLOB      DEFAULT NULL
);
```

Parameters

Parameter	Description
repo	Specifies the repository handle.
file_path	File path to delete file in the repository.
branch_name	Delete file from a specific branch.
commit_details	Commit Details as a JSON document {"message": "Commit message", "author": {"name": "Committing user name", "email": "Email of committing user" } }

Example

```
BEGIN
  DBMS_CLOUD_REPO.DELETE_FILE(
      repo        => l_repo,
      file_path   => 'scripts/test3.sql',
      branch_name => 'test_branch'
  );
END;
/
```

DELETE_REPOSITORY Procedure

This procedure deletes the Cloud Code Repository identified by the `repo` handle argument.

Syntax

```
PROCEDURE DBMS_CLOUD_REPO.DELETE_REPOSITORY(
      repo           IN   CLOB
);
```

Parameters

Parameter	Description
repo	Specifies the repository handle.

Example

```
BEGIN
  DBMS_CLOUD_REPO.DELETE_REPOSITORY(
       repo => l_repo
  );
END;
/
```

EXPORT_OBJECT Procedure

This procedure uploads the DDL metadata of a database object to the Cloud Code repository identified by the `repo` handle argument. This procedure is an easy way to upload the metadata definition of a database object in single step.

Syntax

```
PROCEDURE DBMS_CLOUD_REPO.EXPORT_OBJECT(
      repo              IN  CLOB,
      file_path         IN  VARCHAR2,
      object_type       IN  VARCHAR2,
      object_name       IN  VARCHAR2 DEFAULT NULL,
      object_schema     IN  VARCHAR2 DEFAULT NULL,
      branch_name       IN  VARCHAR2 DEFAULT NULL,
      commit_details    IN  CLOB     DEFAULT NULL,
      append            IN  BOOLEAN  DEFAULT FALSE
);
```

Parameters

Parameter	Description
`repo`	Specifies the repository handle.
`file_path`	File path to upload object metadata in the repository.
`object_type`	Object type supported by DBMS_METADATA. See DBMS_METADATA: Object Types table for details.
`object_name`	Name of the database object to retrieve metadata.
`object_schema`	Owning schema of the database object.
`branch_name`	Put file to a specific branch.
`commit_details`	Commit Details as a JSON document:`{"message": "Commit message", "author": {"name": "Committing user name", "email": "Email of committing user" } }`
`append`	Append metadata DDL to existing file.

Usage Note

For customized control on the object DDL, you can use `DBMS_METADATA.GET_DDL` along with `DBMS_CLOUD_REPO.PUT_FILE`. In order to get metadata definition of the object, the current user must be privileged to retrieve the object metadata. See DBMS_METADATA for the security requirements of the package.

Example

```
BEGIN
  DBMS_CLOUD_REPO.EXPORT_OBJECT(
     repo        => l_repo,
     object_type => 'PACKAGE',
     object_name => 'MYPACK',
     file_path   => 'mypack.sql'
  );
```

```
END;
/
```

EXPORT_SCHEMA Procedure

This procedure exports metadata of all objects in a schema to the Cloud Code Repository branch identified by the `repo` handle argument.

Syntax

```
PROCEDURE DBMS_CLOUD_REPO.EXPORT_SCHEMA(
    repo            IN CLOB,
    file_path       IN VARCHAR2,
    schema_name     IN VARCHAR2,
    filter_list     IN CLOB         DEFAULT NULL,
    branch_name     IN VARCHAR2     DEFAULT NULL,
    commit_details  IN CLOB         DEFAULT NULL
);
```

Parameters

Parameter	Description
`repo`	Specifies the repository handle. This parameter is mandatory and supported for all cloud providers.
`file_path`	Specifies the name of the schema file to upload to the repo. This parameter is mandatory and supported for all cloud providers.
`schema_name`	Specifies the name of the schema for which a DDL script is to be uploaded to the Cloud Code Repository branch. This parameter is mandatory and supported for all cloud providers.
`filter_list`	Specifies the CLOB of the JSON array that defines the filter conditions to include or exclude the objects whose metadata needs to be exported. This parameter is supported for all cloud providers. JSON parameters for `filter_list` are: • `match_type`: Specifies the type of filter to be applied to the object types or object names. Valid `match_type` values are: – `in/not_in` – `like/not_like` – `equal/not_equal` • `type`: Specifies the type of object by which to filter. • `name`: Specifies the name of the object by which to filter.
`branch_name`	Specifies the repository branch name. This parameter is supported for all cloud providers. If you do not supply a `branch_name` value, the `branch_name` is set to the default repository branch.

Parameter	Description
commit_details	Commit Details as a JSON document {"message": "Commit message", "author": {"name": "Committing user name", "email": "Email of committing user" } } This parameter is supported for all cloud providers. If you do not supply a commit_details value, the commit_details is set to the default commit message that includes the information about the current database session user and database name performing the commit.

Example

```
BEGIN
  DBMS_CLOUD_REPO.EXPORT_SCHEMA(
    repo          => l_repo,
    schema_name   => 'USER1',
    file_path     => 'myschema_ddl.sql'
    filter_list   =>
        to_clob('[
             { "match_type":"equal",
                "type":"table"
             },
             { "match_type":"not_equal",
                "type":"view"
             },
             { "match_type":"in",
                "type":"table",
                "name": " ''EMPLOYEE_SALARY'',''EMPLOYEE_ADDRESS'' "
             },
             { "match_type":"equal",
                "type":"sequence",
                "name": "EMPLOYEE_RECORD_SEQ"
             },
             { "match_type":"like",
                "type":"table",
                "name": "%OFFICE%"
             }
        ]'
     );
  );
END;
/
```

Usage Note

To run the DBMS_CLOUD_REPO.EXPORT_SCHEMA procedure, you must be logged in as the ADMIN user or have EXECUTE privilege on DBMS_CLOUD_REPO.

GET_FILE Procedure and Function

The function downloads the contents of a file from the Cloud Code repository. The procedure allows you to download the contents of a file from the Cloud Code repository and save the file in a directory.

Syntax

```
FUNCTION DBMS_CLOUD_REPO.GET_FILE(
      repo              IN  CLOB,
      file_path         IN  VARCHAR2,
      branch_name       IN  VARCHAR2  DEFAULT NULL,
      tag_name          IN  VARCHAR2  DEFAULT NULL,
      commit_name       IN  VARCHAR2  DEFAULT NULL
) RETURN CLOB;

PROCEDURE DBMS_CLOUD_REPO.GET_FILE(
      repo              IN  CLOB,
      file_path         IN  VARCHAR2,
      directory_name    IN  VARCHAR2,
      target_file_name  IN  VARCHAR2  DEFAULT NULL,
      branch_name       IN  VARCHAR2  DEFAULT NULL,
      tag_name          IN  VARCHAR2  DEFAULT NULL,
      commit_name       IN  VARCHAR2  DEFAULT NULL
);
```

Parameters

Parameter	Description
repo	Specifies the repository handle.
file_path	File path in the repository.
directory_name	Directory object name to save the file contents.
target_file_name	Target file name to save contents in directory.
branch_name	Get file from a specific branch.
tag_name	Get file from a specific Tag.
commit_name	Get file from a specific commit.

Example

```
BEGIN
  DBMS_CLOUD_REPO.GET_FILE(
     repo             =>  l_repo,
     file_path        => 'test3.sql',
     directory_name   => 'DATA_PUMP_DIR',
     target_file_name => 'test2.sql'
  );
END;
/
```

INIT_AWS_REPO Function

This function initializes an AWS repository handle and returns an opaque type.

Syntax

```
FUNCTION DBMS_CLOUD_REPO.INIT_AWS_REPO(
      credential_name IN  VARCHAR2,
      repo_name       IN  VARCHAR2,
      region          IN  VARCHAR2
)  RETURN repo;
```

Parameters

Parameter	Description
credential_name	Credential object specifying AWS CodeCommit accesskey/secretkey.
repo_name	Specifies the repository name.
region	Specifies the AWS region for the CodeCommit repository.

Example

```
BEGIN
  :repo := DBMS_CLOUD_REPO.INIT_AWS_REPO(
                credential_name => 'AWS_CRED',
                repo_name       => 'my_repo',
                region          => 'us-east-1'
           );
END;
/
```

INIT_AZURE_REPO Function

This function initializes an Azure repository handle and returns an opaque type. This function is only supported for Azure cloud provider.

Syntax

```
FUNCTION DBMS_CLOUD_REPO.INIT_AZURE_REPO(
      credential_name IN  VARCHAR2,
      repo_name       IN  VARCHAR2,
      organization    IN  VARCHAR2,
      project         IN  VARCHAR2
)  RETURN repo;
```

Parameters

Parameter	Description
credential_name	Credential object specifying Azure, with a Username and Personal Access Token (PAT).
repo_name	Specifies the repository name.

Parameter	Description
organization	Specifies the Azure DevOps Organization.
project	Azure Team Project name.

Example

```
BEGIN
  :repo := DBMS_CLOUD_REPO.INIT_AZURE_REPO(
                credential_name => 'AZURE_CRED',
                repo_name       => 'my_repo',
                organization    => 'myorg',
                project         => 'myproj',
           );
END;
/
```

INIT_GITHUB_REPO Function

This function initializes a GitHub repository handle and returns an opaque type.

Syntax

```
FUNCTION DBMS_CLOUD_REPO.INIT_GITHUB_REPO(
      credential_name IN  VARCHAR2  DEFAULT NULL,
      repo_name       IN  VARCHAR2,
      owner           IN  VARCHAR2)
RETURN repo;
```

Parameters

Parameter	Description
credential_name	Credential object specifying GitHub. User Email and Personal Access Token (PAT).
repo_name	Specifies the repository name.
owner	Specifies the repository owner.

Example

```
BEGIN
  :repo := DBMS_CLOUD_REPO.INIT_GITHUB_REPO(
                credential_name => 'GITHUB_CRED',
                repo_name       => 'my_repo',
                owner           => 'foo'
           );
END;
/
```

INIT_REPO Function

This function initializes a Cloud Code Repository handle and returns an opaque JSON object. This function is a generic interface to accept a JSON document, and avoids having to change code, you only need to change a JSON document, when moving a code repository from one Cloud Code repository to another Cloud Code repository.

Syntax

```
FUNCTION DBMS_CLOUD_REPO.INIT_REPO(
     params       IN   CLOB)
RETURN CLOB;
```

Parameters

JSON Parameter	Description

`provider`	Cloud code repository provider from the following:
	`DBMS_CLOUD_REPO.GITHUB_REPO ('GITHUB')`
	`DBMS_CLOUD_REPO.AWS_REPO ('AWS')`
	`DBMS_CLOUD_REPO.AZURE_REPO ('AZURE')`

`repo_name`	Specifies the repository name. `DBMS_CLOUD_REPO.PARAM_REPO_NAME`
`owner`	GitHub Repository Owner. `DBMS_CLOUD_REPO.PARAM_OWNER` This parameter is only applicable for GitHub cloud provider.
`region`	AWS Repository Region `DBMS_CLOUD_REPO_PARAM_REGION` This parameter is only applicable for AWS cloud provider.
`organization`	Azure Organization `DBMS_CLOUD_REPO_PARAM_ORGANIZATION` This parameter is only applicable for Azure cloud provider.
`project`	Azure Team Project `DBMS_CLOUD_REPO_PARAM_PROJECT` This parameter is only applicable for Azure cloud provider

Example

```
BEGIN
  :repo := DBMS_CLOUD_REPO.INIT_REPO(
           params => JSON_OBJECT('credential_name' value 'mycred',
                                 'repo_name'       value 'myrepo',
                                 'repo_owner'      value 'foo')
         );
END;
/
```

INSTALL_FILE Procedure

This procedure installs SQL statements from a file in the Cloud Code repository identified by the `repo` handle argument.

Syntax

```
PROCEDURE DBMS_CLOUD_REPO.INSTALL_FILE(
      repo              IN  CLOB,
```

```
      file_path          IN  VARCHAR2,
      branch_name        IN  VARCHAR2  DEFAULT NULL,
      tag_name           IN  VARCHAR2  DEFAULT NULL,
      commit_name        IN  VARCHAR2  DEFAULT NULL,
      stop_on_error      IN  BOOLEAN   DEFAULT TRUE
);
```

Parameters

Parameter	Description
repo	Specifies the repository handle.
file_path	File path in the repository.
branch_name	Branch to install file from a specific branch.
tag_name	Tag to install file from a specific Tag.
commit_name	Commit ID to install file from a specific commit.
stop_on_error	Stop executing the SQL statements on first error.

Usage Notes

- You can install SQL statements containing nested SQL from a Cloud Code repository file using the following:
  - `@`: includes a SQL file with a relative path to the ROOT of the repository.
  - `@@`: includes a SQL file with a path relative to the current file.
- The scripts are intended as schema install scripts and not as generic SQL scripts:
  - Scripts cannot contain SQL*Plus client specific commands.
  - Scripts cannot contain bind variables or parameterized scripts.
  - SQL statements must be terminated with a slash on a new line (/).
  - Scripts can contain DDL, DML PLSQL statements, but direct `SELECT` statements are not supported. Using SELECT within a PL/SQL block is supported.

  Any SQL statement that can be run using `EXECUTE IMMEDIATE` will work if it does not contain bind variables or defines.

Example

```
BEGIN
     DBMS_CLOUD_REPO.INSTALL_FILE(
         repo          => l_repo,
         file_path     => 'test3.sql',
         stop_on_error => FALSE
     );
END;
/
```

INSTALL_SQL Procedure

This procedure installs SQL statements from a buffer given as input.

Syntax

```
PROCEDURE DBMS_CLOUD_REPO.INSTALL_SQL(
      content           IN  CLOB,
      stop_on_error     IN  BOOLEAN  DEFAULT TRUE
);
```

Parameters

Parameter	Descriptions
content	Is the CLOB containing the SQL statements to run.
stop_on_error	Stop executing the SQL statements on first error.

Usage Notes

- The scripts are intended as schema install scripts and not as generic SQL scripts:
  - Scripts cannot contain SQL*Plus client specific commands.
  - Scripts cannot contain bind variables or parameterized scripts.
  - SQL statements must be terminated with a slash on a new line (/).
  - Scripts can contain DDL, DML PLSQL statements, but direct `SELECT` statements are not supported. Using SELECT within a PL/SQL block is supported.

  Any SQL statement that can be run using `EXECUTE IMMEDIATE` will work if it does not contain bind variables or defines.

Example

```
BEGIN
  DBMS_CLOUD_REPO.INSTALL_SQL(
      content   => 'create table t1 (x varchar2(30))' || CHR(10) || '/',
      stop_on_error => FALSE
  );
END;
/
```

LIST_BRANCHES Function

This function lists branches in the Cloud Code Repository branch identified by the `repo` handle argument.

Syntax

```
FUNCTION DBMS_CLOUD_REPO.LIST_BRANCHES(
    repo           IN   CLOB
) RETURN list_branch_ret_tab PIPELINED PARALLEL_ENABLE;
```

Parameters

Parameter	Description
repo	Specifies the repository handle. This parameter is mandatory and supported for all cloud providers.

Example

```
SELECT * FROM DBMS_CLOUD_REPO.LIST_BRANCHES (repo => l_repo);
```

Usage Notes

- This is a pipelined table function with return type as `list_branch_ret_tab`.
- `DBMS_CLOUD_REPO.LIST_BRANCHES` returns the column: `name`, which indicates the name of the Cloud Code Repository branch.

LIST_COMMITS Function

This function lists commits in the Cloud Code Repository branch identified by the `repo` handle argument.

Syntax

```
FUNCTION DBMS_CLOUD_REPO.LIST_COMMITS(
      repo               IN  CLOB,
      branch_name        IN  VARCHAR2  DEFAULT NULL,
      file_path          IN  VARCHAR2 DEFAULT NULL,
      commit_id          IN  VARCHAR2  DEFAULT NULL
) RETURN list_commit_ret_tab PIPELINED PARALLEL_ENABLE;
```

Parameters

Parameter	Description
repo	Specifies the repository handle. This parameter is mandatory and supported for all cloud providers.
branch_name	List commits from a specific branch. This parameter is supported for all cloud providers. If you do not supply a `branch_name` value, the `branch_name` is set to `main`.
file_path	List files under the specified subfolder path in the repository. This parameter is only supported for Git and Azure cloud providers. If you do not supply a `file_path` value, the `file_path` is set to a NULL value.
commit_id	List files starting from the specified `sha/id` This parameter is supported for all cloud providers. If you do not supply a `commit_id` value, the `commit_id` is set to a NULL value.

Example

```
SELECT name FROM DBMS_CLOUD_REPO.LIST_COMMITS(repo => l_repo);
```

Example

```
SELECT name FROM DBMS_CLOUD_REPO.LIST_COMMITS (
    repo        => l_repo,
    commit_id   => '66dd2b23b74cd0afabd11af66c6aa9c550540ba6',
    file_path   => 'sub_dir/test11.sql'
);
```

Usage Notes

- This is a pipelined table function with a return type as `list_commit_ret_tab`.
- `DBMS_CLOUD_REPO.LIST_COMMITS` returns the column: `commit_id`.

LIST_FILES Function

This function downloads a file from Cloud Code repository. Optionally, file content can be accessed from either a specific branch, tag or commit name. By default, the file is accessed from the default repository branch. The results include the file names and additional metadata about the files.

Syntax

```
FUNCTION DBMS_CLOUD_REPO.LIST_FILES(
      repo              IN  CLOB,
      path              IN  VARCHAR2  DEFAULT NULL,
      branch_name       IN  VARCHAR2  DEFAULT NULL,
      tag_name          IN  VARCHAR2  DEFAULT NULL,
      commit_id         IN  VARCHAR2  DEFAULT NULL
) RETURN list_file_ret_tab PIPELINED PARALLEL_ENABLE;
```

Parameters

Parameter	Description
`repo`	Specifies the repository handle.
`path`	List files under the specified subfolder path in the repository.
`branch_name`	List files from a specific branch.
`tag_name`	List files from a specific Tag.
`commit_name`	List files from a specific commit.

Usage Notes

- This is a pipelined table function with return type as `list_file_ret_tab`.
- `DBMS_CLOUD_REPO.LIST_FILES` returns the columns: `id`, `name`, `url`, and `bytes`.

Example

```
SELECT name FROM DBMS_CLOUD_REPO.LIST_FILES(repo => l_repo);

NAME
------------------------
test3.sql
```

LIST_REPOSITORIES Function

This function lists all Cloud Code Repositories identified by the `repo` handle argument. If the `repo` handle has a repository name specified, the function does not restrict the listing to the specified repository name; it lists all repositories accessible to the user.

Syntax

```
FUNCTION DBMS_CLOUD_REPO.LIST_REPOSITORIES(
      repo                 IN   CLOB
) RETURN list_repo_ret_tab PIPELINED PARALLEL_ENABLE;
```

Parameters

Parameter	Description
`repo`	Specifies the repository handle. This parameter is supported by all cloud providers.

Usage Notes

- This is a pipelined table function with return type as `list_repo_ret_tab`.
- `DBMS_CLOUD_REPO.LIST_REPOSITORIES` returns the columns: `id`, `name`, `owner`, `description`, `private`, `url`, `bytes`, `created`, and `last_modified`.

Example

```
SELECT name description FROM DBMS_CLOUD_REPO.LIST_REPOSITORIES(:repo);

NAME                 DESCRIPTION
-------------------- ---------------
TestRepo1            My test repo
```

MERGE_BRANCH Procedure

This procedure merges a repository branch into another specified branch in the Cloud Code Repository identified by the `repo` handle argument. The `MERGE_BRANCH` procedure is currently not supported in Azure.

Syntax

```
PROCEDURE DBMS_CLOUD_REPO.MERGE_BRANCH (
    repo               IN   CLOB,
    branch_name        IN   VARCHAR2,
```

```
    parent_branch_name IN   VARCHAR2 DEFAULT NULL,
    commit_details     IN   CLOB      DEFAULT NULL
 );
```

Parameters

Parameter	Description
`repo`	Specifies the repository handle. This parameter is mandatory and supported for GITHUB and AWS cloud providers.
`branch_name`	Specifies the Git branch name to merge. This parameter is mandatory and supported for all cloud providers.
`target_branch_name`	Specifies the target branch name to merge into. This parameter is mandatory and supported for all cloud providers.
`commit_details`	Commit Details as a JSON document `{"message": "Commit message", "author": {"name": "Committing user name", "email": "Email of committing user" } }` If you do not supply a `commit_details`value, the `commit_details` is set to the default commit message that includes the information about current database session user and database name performing the commit.

Example

```
BEGIN
  DBMS_CLOUD_REPO.MERGE_BRANCH (
    repo               => l_repo,
    branch_name        => 'test_branch',
    target_branch_name => 'main'
    );
END;
/
```

Usage Note

To run the `DBMS_CLOUD_REPO.MERGE_BRANCH` procedure, you must be logged in as the ADMIN user or have `EXECUTE` privilege on `DBMS_CLOUD_REPO`.

PUT_FILE Procedure

This procedure uploads a file to the Cloud Code repository identified by the `repo` handle argument. The procedure is overloaded to support either uploading a file from a directory object or uploading the contents from a BLOB to the repository file.

Syntax

```
PROCEDURE DBMS_CLOUD_REPO.PUT_FILE(
      repo              IN  CLOB,
      file_path         IN  VARCHAR2,
      contents          IN  BLOB,
      branch_name       IN  VARCHAR2  DEFAULT NULL,
      commit_details    IN  CLOB      DEFAULT NULL
```

```
);

PROCEDURE DBMS_CLOUD_REPO.PUT_FILE(
      repo              IN  CLOB,
      file_path         IN  VARCHAR2,
      directory_name    IN  VARCHAR2,
      source_file_name  IN  VARCHAR2  DEFAULT NULL,
      branch_name       IN  VARCHAR2  DEFAULT NULL,
      commit_details    IN  CLOB      DEFAULT NULL
);
```

Parameters

Parameter	Description
repo	Specifies the repository handle.
file_path	File path to upload file in the repository.
contents	BLOB containing the file contents.
directory_name	Directory object name containing the file name.
source_file_name	Source file name to upload to repository.
branch_name	Put file to a specific branch.
commit_details	Commit Details as a JSON document: `{"message": "Commit message", "author": {"name": "Committing user name", "email": "Email of committing user" } }`

Example

```
BEGIN
  DBMS_CLOUD_REPO.PUT_FILE(
     repo   => l_repo,
  );
END;
/
```

UPDATE_REPOSITORY Procedure

This procedure updates a Cloud Code repository identified by the `repo` handle argument. UPDATE_REPOSITORY supports updating the name, description, or private visibility status, as supported by the Cloud Code repository.

Syntax

```
PROCEDURE DBMS_CLOUD_REPO.UPDATE_REPOSITORY(
      repo                IN OUT  CLOB,
      new_name            IN      VARCHAR2 DEFAULT NULL,
      description         IN      CLOB     DEFAULT NULL,
      private             IN      BOOLEAN  DEFAULT NULL
);
```

Parameters

Parameter	Description
repo	Specifies the repository handle. This parameter is supported for all cloud providers.
new_name	New name for repository. This parameter is supported for all cloud providers.
description	A short text description for the repository. This parameter is supported for GITHUB and AWS cloud providers.
private	Repository is private and only accessible with valid credentials. This parameter is supported for the GITHUB cloud provider.

Example

```
BEGIN
  DBMS_CLOUD_REPO.UPDATE_REPOSITORY(
     repo         => l_repo,
     new_name     => 'repo2'
  );
END;
/
```

DBMS_AUTO_PARTITION Package

The `DBMS_AUTO_PARTITION` package provides administrative routines for managing automatic partitioning of schemas and tables.

- CONFIGURE Procedure
  This procedure configures settings for automatic partitioning in Autonomous Database.
- VALIDATE_CANDIDATE_TABLE Function
  This function checks if the given table is a valid candidate for automatic partitioning in Autonomous Database.
- RECOMMEND_PARTITION_METHOD Function
  This function returns a recommendation ID that can be used with `APPLY_RECOMMENDATION` procedure to apply the recommendation, or can be used with `DBA_AUTO_PARTITION_RECOMMENDATIONS` view to retrieve details of the recommendations for automatic partitioning in Autonomous Database.
- APPLY_RECOMMENDATION Procedure
  This procedure applies the given recommendation in an Autonomous Database.
- REPORT_ACTIVITY Function
  This function returns a report of the automatic partitioning operations executed during a specific period in an Autonomous Database.
- REPORT_LAST_ACTIVITY Function
  This function returns a report of the most recent automatic partitioning operation executed in an Autonomous Database.

CONFIGURE Procedure

This procedure configures settings for automatic partitioning in Autonomous Database.

Syntax

```
DBMS_AUTO_PARTITION.CONFIGURE (
     PARAMETER_NAME     IN VARCHAR2,
     PARAMETER_VALUE    IN VARCHAR2,
     ALLOW              IN BOOLEAN    DEFAULT TRUE);
```

Parameters

Parameter	Description
PARAMETER_NAME	Name of the automatic partitioning configuration parameter to update. It can have one of the following values: • AUTO_PARTITION_MODE • AUTO_PARTITION_SCHEMA • AUTO_PARTITION_TABLE • AUTO_PARTITION_REPORT_RETENTION AUTO_PARTITION_MODE sets the mode of automatic partitioning operation, and has one of the following values: • IMPLEMENT: In this mode, automatic partitioning generates a report and modifies the existing table using the recommended partition method. • REPORT ONLY: In this mode, automatic partitioning generates a report but existing tables are not modified. This is the default value. • OFF: In this mode, automatic partitioning is prevented from generating, considering, or applying recommendations. It does not disable existing automatic partitioned tables. AUTO_PARTITION_SCHEMA sets schemas to include or exclude from using automatic partitioning. Its behavior is controlled by the allow parameter. The automatic partitioning process manages two schema lists. 1. Inclusion list is the list of schemas, case-sensitive, that can use automatic partitioning. 2. Exclusion list is the list of schemas, case-sensitive, that cannot use automatic partitioning. Initially, both lists are empty, and all schemas in the database can use automatic partitioning. If the inclusion list contains one or more schemas, then only the schemas listed in the inclusion list can use automatic partitioning. If the inclusion list is empty and the exclusion list contains one or more schemas, then all schemas use automatic partitioning except the schemas listed in the exclusion list. If both lists contain one or more schemas, then all schemas use automatic partitioning except the schemas listed in the exclusion list. AUTO_PARTITION_TABLE sets tables to include or exclude from using auto partitioning. The parameter value is <schema_name>.<table_name>. The automatic partitioning process manages two table lists. 1. Inclusion list is the list of tables, case-sensitive, that can use automatic partitioning. 2. Exclusion list is the list of tables, case-sensitive, that cannot use automatic partitioning. Initially, both lists are empty, and all tables in the database can use automatic partitioning. If the inclusion list contains one or more tables, then only the tables listed in the inclusion list can use automatic partitioning. If the inclusion list is empty and the exclusion list contains one or more tables, then all tables use automatic partitioning except the tables listed in the exclusion list. If both lists contain one or more tables, then all tables use automatic partitioning except the tables listed in the exclusion list. If a table is not on either list, the schema inclusion and exclusion lists decide if a table is a candidate table for automatic partitioning. If there is a conflict between the schema level lists and the table level lists, the table level lists take precedence.

Parameter	Description
	To remove all tables from inclusion and exclusion lists run: `DBMS_AUTO_PARTITION.CONFIGURE('AUTO_PARTITION_TABLE', NULL);` `AUTO_PARTITION_REPORT_RETENTION` sets the number of days for which automatic partitioning logs are retained in the database before they are deleted. An automatic partitioning report cannot be generated for a period beyond the value specified for this value. Default value is 90 days.
`PARAMETER_VALUE`	Value for the configuration setting specified in `parameter_name`. When set to `NULL`, the configuration setting is assigned its default value.
`ALLOW`	Applicable only for the `AUTO_PARTITION_SCHEMA` or `AUTO_PARTITION_TABLE` configuration settings with one of the following values: • `TRUE` adds specified schema or table to the inclusion list. • `FALSE` removes specified schema or table from the exclusion list. • `NULL` removes specified schema or table from the list to which currently assigned. Refer to the description of the `AUTO_PARTITION_SCHEMA` and `AUTO_PARTITION_TABLE` configuration settings for more information about inclusion lists and exclusion lists.

Usage Notes

- You can check the current setting for automatic partitioning configuration using the following SQL:

  ```
  SELECT * FROM DBA_AUTO_PARTITION_CONFIG;
  ```

- Unlike automatic indexing, automatic partitioning does not run periodically as a background task. Automatic partitioning only runs when you invoke it using the `DBMS_AUTO_PARTITION.RECOMMEND_PARTITION_METHOD` function.

VALIDATE_CANDIDATE_TABLE Function

This function checks if the given table is a valid candidate for automatic partitioning in Autonomous Database.

Valid Candidate

To be a valid candidate, the following tests must pass:

- Table passes inclusion and exclusion tests specified by `AUTO_PARTITION_SCHEMA` and `AUTO_PARTITION_TABLE` configuration parameters.
- Table exists and has up-to-date statistics.
- Table is at least 64 GB.
- Table has 5 or more queries in the SQL tuning set that scanned the table.
- Table does not contain a `LONG` data type column.
- Table is not manually partitioned.

- Table is not an external table, an internal/external hybrid table, a temporary table, an index-organized table, or a clustered table.
- Table does not have a domain index or bitmap join index.
- Table is not an advance queuing, materialized view, or flashback archive storage table.
- Table does not have nested tables, or certain other object features.

Returns:

- `VALID` if the table is a valid candidate for autonomous partitioning
- `INVALID: <reason>` if the table is not a valid candidate for autonomous partitioning, and <reason> is a string describing why the table is not a valid candidate.

Syntax

```
DBMS_AUTO_PARTITION.VALIDATE_CANDIDATE_TABLE
    ( SQLSET_OWNER  IN VARCHAR2   DEFAULT 'SYS',
      SQLSET_NAME   IN VARCHAR2   DEFAULT 'SYS_AUTO_STS',
      TABLE_OWNER   IN VARCHAR2,
      TABLE_NAME    IN VARCHAR2)
RETURN VARCHAR2;
```

Parameters

Parameter	Description
`SQLSET_OWNER`, `SQLSET_NAME`	Name of SQL tuning set representing the workload to be evaluated.
`TABLE_OWNER`, `TABLE_NAME`	Name of a table to validate as a candidate for automatic partitioning.

Usage Notes

- As an example, you can check the validity of a sample table, `LINEORDER` in schema `TEST`, with the following SQL:

```
SELECT DBMS_AUTO_PARTITION.VALIDATE_CANDIDATE_TABLE
    ( TABLE_OWNER => 'TEST',
      TABLE_NAME  => 'LINEORDER')
FROM DUAL;
```

RECOMMEND_PARTITION_METHOD Function

This function returns a recommendation ID that can be used with `APPLY_RECOMMENDATION` procedure to apply the recommendation, or can be used with `DBA_AUTO_PARTITION_RECOMMENDATIONS` view to retrieve details of the recommendations for automatic partitioning in Autonomous Database.

Syntax

```
DBMS_AUTO_PARTITION.RECOMMEND_PARTITION_METHOD
    ( SQLSET_OWNER     IN VARCHAR2   DEFAULT 'SYS',
      SQLSET_NAME      IN VARCHAR2   DEFAULT 'SYS_AUTO_STS',
      TABLE_OWNER      IN VARCHAR2   DEFAULT NULL,
```

```
    TABLE_NAME       IN VARCHAR2   DEFAULT NULL,
    TIME_LIMIT       IN INTERVAL DAY TO SECOND DEFAULT INTERVAL '1' DAY,
    REPORT_TYPE      IN VARCHAR2   DEFAULT 'TEXT',
    REPORT_SECTION   IN VARCHAR2   DEFAULT 'SUMMARY',
    REPORT_LEVEL     IN VARCHAR2   DEFAULT 'TYPICAL')
 RETURN RAW;
```

Parameters

Parameter	Description
SQLSET_OWNER, SQLSET_NAME	Name of SQL tuning set representing the workload to be evaluated.
TABLE_OWNER, TABLE_NAME	Name of a table to validate as a candidate for automatic partitioning.
TIME_LIMIT	When the function chooses the tables for which to generate recommendations, TABLE_OWNER and TABLE_NAME are NULL), parameter limits how long the function runs before it stops looking for new candidate tables to partition. Once started processing a table, process will not terminate. It is expected that the function may run longer than this parameter. If this parameter is NULL there is no time limit. The default is 1 day.
REPORT_TYPE	Used to generate report for recommended partition method. See REPORT_ACTIVITY Function for details.
REPORT_SECTION	Used to generate persistent report for recommended partition method. See REPORT_ACTIVITY Function for details.
REPORT_LEVEL	Used to generate report for recommended partition method. See REPORT_ACTIVITY Function for details.

Usage Notes

- The AUTO_PARTITION_MODE controls the actions taken by this function:
  - IMPLEMENT: In this mode, automatic partitioning generates a report and modifies the existing table using the recommended partition method.
  - REPORT ONLY: In this mode, automatic partitioning generates a report generated but existing tables are not modified. This is the default value.
  - OFF: In this mode, automatic partitioning prevented from producing, considering, or applying new recommendations. It does not disable existing automatic partitioned tables.
- Unlike automatic indexing, automatic partitioning does not run periodically as a background task. Automatic partitioning only runs when you invoke it using the DBMS_AUTO_PARTITION.RECOMMEND_PARTITION_METHOD function.

Return Values

This function returns a recommendation ID that can be used as follows: DBMS_AUTO_PARTITION.APPLY_RECOMMENDATION to apply the recommendation,

- Use with DBMS_AUTO_PARTITION.APPLY_RECOMMENDATION to apply the recommendation.

- Use with DBA_AUTO_PARTITION_RECOMMENDATIONS view to retrieve details of the recommendations. For example:

```
SELECT PARTITION_METHOD, PARTITION_KEY
     FROM DBA_AUTO_PARTITION_RECOMMENDATIONS
     WHERE RECOMMENDATION_ID = :RECOMMENDATION_ID;
```

APPLY_RECOMMENDATION Procedure

This procedure applies the given recommendation in an Autonomous Database.

Syntax

```
DBMS_AUTO_PARTITION.APPLY_RECOMMENDATION
    ( RECOMMENDATION_ID  IN RAW,
      TABLE_OWNER        IN VARCHAR2   DEFAULT NULL,
      TABLE_NAME         IN VARCHAR2   DEFAULT NULL);
```

Parameters

Parameter	Description
RECOMMENDATION_ID	Recommendation ID returned from RECOMMEND_PARTITION_METHOD function or queried from DBA_AUTO_PARTITION_RECOMMENDATIONS view.
TABLE_OWNER, TABLE_NAME	When a single recommendation ID has recommendations for multiple tables, this optional parameter allows you to control which tables are partitioned. • If parameters are NULL, partition all tables recommended in the given recommendation ID. • If a table name is given, partition only the named table. • If either TABLE_OWNER or TABLE_NAME is NOT NULL, they must both be NOT NULL.

Usage Note:

Regardless of AUTO_PARTITION_MODE, this procedure raises an ORA-20000: recommendation_id was not found if either there are no accepted recommendations associated with the RECOMMENDATION_ID, or all accepted recommendations associated with the RECOMMENDATION_ID have already been applied. The first case applies if RECOMMENDATION_ID was generated with AUTO_PARTITION_MODE = OFF. The second case applies if RECOMMENDATION_ID was generated with AUTO_PARTITION_MODE = IMPLEMENT.

REPORT_ACTIVITY Function

This function returns a report of the automatic partitioning operations executed during a specific period in an Autonomous Database.

Syntax

```
DBMS_AUTO_PARTITION.REPORT_ACTIVITY
    ( ACTIVITY_START     IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
      ACTIVITY_END       IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
      TYPE               IN VARCHAR2                 DEFAULT 'TEXT',
```

```
      SECTION            IN VARCHAR2                DEFAULT 'ALL',
      LEVEL              IN VARCHAR2                DEFAULT 'TYPICAL')
RETURN CLOB;
```

Parameters

Parameter	Description
ACTIVITY_START	Starting time automatic partitioning operations use for the report. If no value is specified, or NULL is specified, the report is generated for the last automatic partitioning operation that was executed.
ACTIVITY_END	Ending time automatic partitioning operations use for the report. If no value is specified, or NULL is specified, then the report is generated for the last automatic partitioning operation that was executed.
TYPE	Format of the report that has one of the following values: • TEXT (default) • HTML • XML
SECTION	Sections to include in the report that has one of the following values: • SUMMARY - Include only the workload summary in the report • ALL = Include all the sections in the report. (default)
level	Level of information to include in the report that has one of the following values: • TYPICAL - Include typical automatic partitioning information in the report (default). • CHANGED - Include only SQL with changed performance in the report. • IMPROVED - Include only SQL with improved performance in the report. • REGRESSED - Include only SQL with regressed performance in the report. • UNCHANGED - Include only SQL with unchanged performance in the report. • ALL - Include all automatic partitioning information in the report.

Usage Notes

Returns: A performance analysis report for workload executed on database after recommendation is applied. This report is not stored persistently with the recommendation.

REPORT_LAST_ACTIVITY Function

This function returns a report of the most recent automatic partitioning operation executed in an Autonomous Database.

Syntax

```
DBMS_AUTO_PARTITION.REPORT_LAST_ACTIVITY
    ( TYPE    IN VARCHAR2 DEFAULT 'TEXT',
      SECTION IN VARCHAR2 DEFAULT 'ALL',
      LEVEL   IN VARCHAR2 DEFAULT 'TYPICAL')
RETURN CLOB;
```

Parameters

Parameter	Description
TYPE	The output format of the report, see REPORT_ACTIVITY Function for information.
SECTION	The sections included in the report, see REPORT_ACTIVITY Function for information.
LEVEL	The level of information included in the report, see REPORT_ACTIVITY Function for information.

Usage Notes

Returns: A performance analysis report for workload executed on database after latest recommendation is applied. This report is not stored persistently with the recommendation.

DBMS_DATA_ACCESS Package

The DBMS_DATA_ACCESS package provides routines to generate and manage Table Hyperlinks for data sets.

- DBMS_DATA_ACCESS Overview
  Describes the use of the DBMS_DATA_ACCESS package.
- DBMS_DATA_ACCESS Security Model
  Security on this package can be controlled by granting EXECUTE on this package to selected users or roles.
- Summary of DBMS_DATA_ACCESS Subprograms
  This section covers the DBMS_DATA_ACCESS subprograms provided with Autonomous Database.

DBMS_DATA_ACCESS Overview

Describes the use of the DBMS_DATA_ACCESS package.

DBMS_DATA_ACCESS supports these operations:

- Generation of a Table Hyperlink
- Manual invalidation of a Table Hyperlink
- Listing of active Table Hyperlinks

DBMS_DATA_ACCESS Security Model

Security on this package can be controlled by granting EXECUTE on this package to selected users or roles.

When a user has been granted EXECUTE on DBMS_DATA_ACCESS they are able to create, list or invalidate the Table Hyperlinks that are created by the user. In addition, by default the ADMIN user has the following privileges:

- The ADMIN user with PDB_DBA role has EXECUTE privilege on DBMS_DATA_ACCESS.
- The ADMIN user with the PDB_DBA role is able to list or invalidate any Table Hyperlink in an Autonomous Database instance.

Summary of DBMS_DATA_ACCESS Subprograms

This section covers the `DBMS_DATA_ACCESS` subprograms provided with Autonomous Database.

Subprogram	Description
GET_PREAUTHENTICATED_URL Procedure	This procedure generates a Table Hyperlink.
EXTEND_URL Procedure	This procedure extends the life of a Table Hyperlink.
INVALIDATE_URL Procedure	This procedure invalidates a Table Hyperlink.
LIST_ACTIVE_URLS Function	This function lists all the currently active Table Hyperlinks.

- GET_PREAUTHENTICATED_URL Procedure
- EXTEND_URL Procedure
  This procedure extends the life of a Table Hyperlink.
- INVALIDATE_URL Procedure
  This procedure invalidates a Table Hyperlink.
- LIST_ACTIVE_URLS Function
  This function lists all the currently active Table Hyperlinks and Table Hyperlink Groups.

GET_PREAUTHENTICATED_URL Procedure

This procedure generates a Table Hyperlink.

There are two forms, one to generate the Table Hyperlink for a specific object (table or view). The overloaded form, using the `sql_statement` parameter, generates a Table Hyperlink for a SQL statement.

Syntax

```
DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL(
    schema_name           IN VARCHAR2,
    schema_object_name    IN VARCHAR2,
    application_user_id   IN VARCHAR2,
    expiration_minutes    IN NUMBER,
    expiration_count      IN NUMBER,
    service_name          IN VARCHAR2,
    column_lists          IN CLOB,
    inherit_acl           IN BOOLEAN  DEFAULT FALSE,
    result                OUT CLOB);

DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL(
    sql_statement         IN CLOB,
    application_user_id   IN VARCHAR2,
    default_bind_values   IN CLOB,
    expiration_minutes    IN NUMBER,
    expiration_count      IN NUMBER,
    service_name          IN VARCHAR2,
    column_lists          IN CLOB,
    inherit_acl           IN BOOLEAN  DEFAULT FALSE,
    result                OUT CLOB);
```

Parameters

Parameter	Description
schema_name	Specifies the owner of the object.
schema_object_name	Specifies the schema object (table or view).
sql_statement	Specifies the SELECT statement query text. Bind variable support is available for NUMBER and VARCHAR2 column types.
application_user_id	Specifies an application user ID value. When the Table Hyperlink is accessed, the value of application_user_id specified during Table Hyperlink generation is available through: sys_context('DATA_ACCESS_CONTEXT$', 'USER_IDENTITY') You can define VPD Policies that make use of this value in the Application Context to restrict the rows visible to the application user.
default_bind_values	Specifes the default value(s) of one or more bind variables (for a sql_statement specified with bind variables). This allows a table hyperlink consumer to access the table hyperlink data with default bind values, without providing the bind values as query parameters.
expiration_minutes	Duration in minutes of validity of Table Hyperlink. The maximum allowed expiration time is 90 days (129600 minutes). If the value is set to greater than 129600, the value used is 129600 minutes (90 days). If expiration_minutes is specified as a non-null value, expiration_count must not be set to a non-null value. Both cannot be non-null at the same time. Default value: when expiration_minutes is not provided or when expiration_minutes is provided as NULL, the value is set to 90 days (129600 minutes).
expiration_count	Number of accesses allowed on the Table Hyperlink. There is no default value. If expiration_count is not specified and expiration_minutes is not specified, expiration_minutes is set to 90 days (129600 minutes). If expiration_count is specified as a non-null value, expiration_minutes must not be set to a non-null value. Both cannot be non-null at the same time.
service_name	The database service to use for data retrieval when using the Table Hyperlink. Specify the service-level guarantee and resources used to service this Table Hyperlink. For example, access to an object or SQL statement can be mapped to services HIGH or MEDIUM, whereas access to another object or SQL statement can be mapped to the LOW service. Supported values are HIGH, MEDIUM, LOW. The default value is LOW.

Parameter	Description

column_lists	A JSON value that specifies options by column. The supported options specified in the column_lists parameter are one or more of:
	• order_by_columns: specifies the columns that support sorting.
	• filter_columns: specifies the columns that support filtering
	• default_color_columns: specifies to only use the default coloring for the specified columns.
	• group_by_columns: specifies that group by is allowed for the specified columns (viewing the data by grouping the specified column is allowed).
	The column_lists parameter is JSON that contains a list of JSON arrays of columns defining Table Hyperlink functionality. Use this parameter to specify the columns for one or more of the options: order_by_columns, filter_columns, default_color_columns, or group_by_columns.
	The form is:
	`"column_lists" : { "order_by_columns": [order_by_columns_list], "filter_columns": [filter_columns_list], "default_color_columns": [default_color_columns_list], "group_by_columns": [group_by_columns_list] },`
	For example:
	`"column_lists" : { "order_by_columns": ["NAME", "DEPARTMENT"], "filter_columns": ["ID", "NAME", "DEPARTMENT"], "default_color_columns": ["DEPARTMENT"], "group_by_columns": ["DEPARTMENT"] },`
	Default values:
	If column_lists is not specified for order_by_columns and filter_columns options, sorting and filtering is enabled for all columns.
	If column_lists is not specified for group_by_columns, the group by option is not enabled for any column. By default, columns defined to enable as group_by_columns are also be enabled as filter_columns.

inherit_acl	Set the value of this parameter to TRUE to inherit ACLs. When inherit is true, an incoming Table Hyperlink consumer's IP address is validated with the ACL lists on the producer database before allowing access to data. When the parameter is not supplied or the parameter value is set to FALSE, ACL checks are not applied. If the producer database does not have ACLs configured, the inherit_acl value is ignored and data access is allowed without any ACL checks. The default value is FALSE.
result	JSON that indicates the outcome of the operation.

Usage Notes

- There is a limit of 128 active Table Hyperlinks on an Autonomous Database instance.
- When using a Table Hyperlink from a browser, the following options are supported:
  - View the returned data in table format with no coloring (default), by appending the `?view=table` query parameter to the Table Hyperlink.
  - View the returned data in table format and select the column or columns you want colored with preset colors based on column values. To do this, append the `?view=table&colored_column_names=`*`column_name_1`*`,`*`column_name_2`*`,...`*`column_name_n`* query parameter to the Table Hyperlink, where *`column_name_1`* through *`column_name_n`* are the names of the columns you want colored.
  - View the returned data in table format and select a specific column data type you want colored with preset colors, by appending the `?view=table&colored_column_types=`*`data_type`* query parameter. The supported *`data_type`* parameter values are `VARCHAR` and `NONE`.
  - The `sql_statement` parameter value must be a `SELECT` statement. The `SELECT` statement supports bind variables.

    If bind variables are included in the select statement and values are not set in the `default_bind_values` parameter, bind variable values must be appended to the generated Table Hyperlink as a query parameter when accessing the data.

    When you include the `default_bind_values` parameter, when you access the data you can omit the bind variable values when default values are specified in the `default_bind_values` parameter. You can override a default bind variable value specified with `default_bind_values` by explicitly supplying the bind variable value as a query parameter.
- When you generate a Table Hyperlink on an Autonomous Database instance with a private endpoint, the result includes a name `private_preauth_url` with the value of the form: "`https://`*`private-endpoint`*`/adb/p/`*`parurl-token`*`/data`".

  When you generate a Table Hyperlink on an Autonomous Database instance with a private endpoint and the private endpoint is configured with **Allow public access** enabled, the result includes both the `preauth_url` for the public endpoint and `private_preauth_url`.

  See Configure Private Endpoints and Use a Private Endpoint with Public Access Allowed for more information.

Example - Table Hyperlink Generated for a Specific Object

The following example generates a Table Hyperlink for `STUDENTS_VIEW`:

```
DECLARE
   status CLOB;
   BEGIN
   DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL(
      schema_name => 'USER1',
      schema_object_name => 'STUDENTS_VIEW',
      expiration_minutes => 120,
      service_name => 'HIGH',
      result => status);
   dbms_output.put_line(status);
```

```
END;
/
```

Example - Table Hyperlink Generated for a SQL Statement

The following example generates a Table Hyperlink for a `SELECT` SQL statement:

```
DECLARE
   status CLOB;
   par_url_app_string CLOB;
   BEGIN
       par_url_app_string := 1919292929;
       DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL(
            sql_statement => 'SELECT student_id, student_name FROM
STUDENTS_VIEW ORDER BY student_id',
            application_user_id => par_url_app_string,
            expiration_count => 25,
            result => status);
END;
/
```

Example - Table Hyperlink Generated for a SQL Statement with a Bind Variable

The following example uses a bind variable in the `SELECT` statement to generate the Table Hyperlink:

```
set serveroutput on
DECLARE
  status clob;
BEGIN
  DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL(
    sql_statement => 'select * from TREE_DATA WHERE COUNTY = :countyNAME',
    expiration_minutes => 3000,
    result => status);
    dbms_output.put_line('status : '||status);
END;
/
```

To use the generated Table Hyperlink, the bind variable value must be passed. The following example uses the generated Table Hyperlink to access tree data for the first county:

```
https://dataaccess.adb.us-chicago-1.oraclecloudapps.com/adb/p/gTlbq...example/
data?countyNAME=First
```

Use Table Hyperlink to Access Data with Bind Variables

The following example uses a bind variable in the `SELECT` statement and includes the `default_bind_values` parameter to generate the Table Hyperlink:

```
set serveroutput on
DECLARE
  status clob;
BEGIN
  DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL(
```

```
    sql_statement = 'SELECT * FROM TREE_DATA WHERE COUNTY = :countyNAME',
    default_bind_values => '{"countyNAME" : "First"}',
    expiration_minutes => 3000,
    result => status);
    dbms_output.put_line('status : '||status);
END;
/
```

In this case, case the default bind variable value is used and you do not need to provide the value as a query parameter. For example:

```
curl https://dataaccess.adb.us-chicago-1.oraclecloudapps.com/adb/p/
K6X...example/data


curl https://dataaccess.adb.us-chicago-1.oraclecloudapps.com/adb/p/
K6X...example/data

{"items":[
{"COUNTY":"First","SPECIES":"Pine","HEIGHT":16},
{"COUNTY":"First","SPECIES":"Spruce","HEIGHT":6},
{"COUNTY":"First","SPECIES":"Hawthorn","HEIGHT":19},
{"COUNTY":"First","SPECIES":"Cherry","HEIGHT":20},
{"COUNTY":"First","SPECIES":"Chestnut","HEIGHT":51}],
"hasMore":false,
"limit":100,
"offset":0,
"count":6,
"links":
[
{"rel":"self",
"href":"https://dataaccess.adb.us-ashburn-1.oraclecloudapps.com/adb/p/
gTlbq...example/data"}
]}
```

You can override the default bind variable value by explicitly specifying the value as a query parameter. For example:

```
curl https://dataaccess.adb.us-chicago-1.oraclecloudapps.com/adb/p/
K6X...example/data?countyNAME=MAIN
```

See GET_PREAUTHENTICATED_URL Procedure for more information.

Example - Table Hyperlink Generated for a Specific Object with Group By Columns

The following example generates a Table Hyperlink for a specific table with Group By columns specified:

```
DECLARE
   status CLOB;
   BEGIN
      DBMS_DATA_ACCESS.GET_PREAUTHENTICATED_URL(
          schema_name => 'ADMIN',
          schema_object_name    => 'TREE_DATA',
```

```
          expiration_minutes    => 360,
          service_name          => 'HIGH',
          column_lists          => {"group_by_columns": ["COUNTY",
"SPECIES"]}',
          result                => status);

       dbms_output.put_line(status);
    END;
/
```

EXTEND_URL Procedure

This procedure extends the life of a Table Hyperlink.

Syntax:

```
DBMS_DATA_ACCESS.EXTEND_URL(
    id                              IN VARCHAR2,
    extend_expiration_minutes_by    IN NUMBER,
    extend_expiration_count_by      IN NUMBER,
    result                          OUT CLOB);
```

Parameters

Parameter	Description
`id`	Specifies the ID of the Table Hyperlink to extend.
`extend_expiration_minutes_by`	Number of minutes by which to extend expiration time of the Table Hyperlink. The expiration time is set to the current expiration time plus the value of `extend_expiration_minutes_by`. The value for `extend_expiration_minutes_by` plus the current expiration time must not exceed 129600 (which corresponds to 90 days). If `extend_expiration_minutes_by` is null, `extend_expiration_count_by` must not be null. Both cannot be null at the same time. Default value is NULL.
`extend_expiration_count_by`	The number of accesses on the Table Hyperlink is extended by this count. The expiration count is set to the current expiration count plus the value of `extend_expiration_count_by`. If `extend_expiration_count_by` is null, `extend_expiration_minutes_by` must not be null. Both cannot be null at the same time. Default value is null.
`result`	JSON that indicates the outcome of the operation.

Example - Extend Expiration Minutes of Table Hyperlink

```
set serveroutput on
declare
  status clob;
  js_status json_object_t;
  js_arr    json_array_t;
  url_id varchar2(4000);
```

```
begin
  -- Initially sets the expiration time to 60 minutes
  dbms_data_access.get_preauthenticated_url(
    schema_name        => 'SCOTT',     -- Schema name
    schema_object_name => 'EMPLOYEE',  -- Schema object name
    expiration_minutes => 60,          -- Expiration minutes
    service_name       => 'HIGH',
    result             => status);
   js_status := json_object_t.parse(status);
  url_id := js_status.get_string('id');
  dbms_output.put_line('The url id of url: ' || url_id);
  dbms_output.put_line('Initial Expiration Time: ' ||
                       js_status.get_string('expiration_ts'));
  -- Extend the expiration minutes by 1 day, the url would now expire
  -- 24 hours later than the previous expiration time
  dbms_data_access.extend_url(
    id                           => url_id,
    extend_expiration_minutes_by => 1440,
    result                       => status);
   -- List urls created
  status := dbms_data_access.list_active_urls;
  js_arr := json_array_t.parse(status);
  for indx in 0.. js_arr.get_size - 1
  loop
    js_status := TREAT (js_arr.get (indx) AS json_object_t);
    if js_status.get_string('id') = url_id then
      dbms_output.put_line('New Expiration Time : ' ||
                           js_status.get_string('expiration_time'));
      exit;
    end if;
  end loop;
end;
/
```

Example - Extend Expiration Count of Table Hyperlink

```
set serveroutput on
declare  status clob;
  js_status json_object_t;
  js_arr    json_array_t;
  url_id varchar2(4000);
begin
  -- Initially sets the expiration count to 100
  dbms_data_access.get_preauthenticated_url(
    schema_name        => 'SCOTT',     -- Schema name
    schema_object_name => 'EMPLOYEE',  -- Schema object name
    expiration_count   => 100,         -- Expiration count
    service_name       => 'HIGH',
    result             => status);
  js_status := json_object_t.parse(status);
  url_id := js_status.get_string('id');
  dbms_output.put_line('The url id of url: ' || url_id);
  dbms_output.put_line('Initial Expiration Count: ' ||
                       js_status.get_string('expiration_count'));
  -- Extends access count by 100 so url would expire after 200 accesses
```

```
  dbms_data_access.extend_url(
    id                        => url_id,
    extend_expiration_count_by => 100,
    result                    => status);
  -- List urls created
  status := dbms_data_access.list_active_urls;
  js_arr := json_array_t.parse(status);
  for indx in 0.. js_arr.get_size - 1
  loop
    js_status := TREAT (js_arr.get (indx) AS json_object_t);
     if js_status.get_string('id') = url_id then
      dbms_output.put_line('New Expiration Count : ' ||
                           js_status.get_string('expiration_count'));
      exit;
    end if;
  end loop;
end;
/
```

INVALIDATE_URL Procedure

This procedure invalidates a Table Hyperlink.

Syntax

```
DBMS_DATA_ACCESS.INVALIDATE_URL(
    id                  IN VARCHAR2,
    kill_sessions       IN BOOLEAN DEFAULT FALSE,
    result              OUT CLOB);
```

Parameters

Parameter	Description
`id`	Specifies the owner of the object.
`kill_sessions`	By default, existing sessions that may be in the middle of accessing data using a Table Hyperlink are not killed. When TRUE, this parameter specifies that such existing sessions should be killed, so that the invalidation does not leave any ongoing access to the data set. Valid values: `TRUE` | `FALSE`.
`result`	Provides JSON to indicate whether invalidation is a success or a failure (`CLOB`).

LIST_ACTIVE_URLS Function

This function lists all the currently active Table Hyperlinks and Table Hyperlink Groups.

Syntax

```
DBMS_DATA_ACCESS.LIST_ACTIVE_URLS RETURN CLOB;
```

Parameters

Parameter	Description
RETURN	The return value is a JSON array.

Example

```
DECLARE
   result CLOB;
   BEGIN
       result := DBMS_DATA_ACCESS.LIST_ACTIVE_URLS;
       DBMS_OUTPUT.PUT_LINE(result);
   END;
```

```
[{"id":"pT36lYHFGA4s3UXSNBCRO13v3D4_example1",
"created_by":"SCOTT",
"service_name":"HIGH",
"expiration_time":"2025-07-28T16:38:02.723Z",
"expiration_count":10,
"access_count":0,
"created":"2025-04-29T16:38:02.977Z",
"inherit_acl":true,
"is_group_url":false,
"group_ids":[null],
"sql_statement":"select * FROM TREE_DATA WHERE COUNTY = :county"}]
```

Usage Notes

- The behavior of `DBMS_DATA_ACCESS.LIST_ACTIVE_URLS` is dependent on the invoker. If the invoker is ADMIN or any user with `PDB_DBA` role, the function lists all active Table Hyperlinks, regardless of the user who generated the Table Hyperlink. If the invoker is not the ADMIN user and not a user with `PDB_DBA` role, the list includes only the active Table Hyperlinks generated by the invoker.
- When you generate and list a Table Hyperlink on an Autonomous Database instance with a private endpoint, the result includes a name `private_preauth_url` with the value of the form: "`https://`*`private-endpoint`*`/adb/p/`*`parurl-token`*`/data`".

  When you generate and list a Table Hyperlink on an Autonomous Database instance with a private endpoint and the private endpoint is configured with **Allow public access** enabled, the result includes both the `preauth_url` for the public endpoint and `private_preauth_url`.

  See Configure Private Endpoints and Use a Private Endpoint with Public Access Allowed for more information.
- When a Table Hyperlink is a group member the `DBMS_DATA_ACCESS.LIST_ACTIVE_URLS` response entry shows "`group_ids`" with a non-null value that includes one or more IDs. The IDs show Table Hyperlink Group IDs that the Table Hyperlink (group member) is a member of.

DBMS_DCAT Package

The `DBMS_DCAT` package provides functions and procedures to help Autonomous Database users leverage the data discovery and centralized metadata management system of OCI Data Catalog.

Data Catalog harvests metadata from a data lake's object storage assets. The harvesting process creates logical entities, which can be thought of as tables with columns and associated data types. `DBMS_DCAT` procedures and functions connect Autonomous Database to Data Catalog and then synchronize the assets with the database, creating protected schemas and external tables. You can then query object store using those external tables, easily joining external data with data stored in Autonomous Database. This dramatically simplifies the management process; there is a single, centrally managed metadata store that is shared across multiple OCI services (including Autonomous Databases). There are also Autonomous Database dictionary views that allow you to inspect the contents of Data Catalog using SQL, and show you how these Data Catalog entities map to your Autonomous Database schemas and tables.

- Data Catalog Users and Roles
  The `DBMS_DCAT` package supports synced users/schemas, `dcat_admin` users and local users. Users must have the `dcat_sync` role to be able to use this package.
- Required Credentials and IAM Policies
  This topic describes the Oracle Cloud Infrastructure Identity and Access Management (IAM) user credentials and policies required to give Autonomous Database users permission to manage a data catalog and to read from object storage.
- Summary of Connection Management Subprograms
  This table lists the `DBMS_DCAT` package procedures used to create, query and drop Data Catalog connections.
- Summary of Synchronization Subprograms
  Running a synchronization, creating and dropping a synchronization job, and dropping synchronized schemas can be performed with the procedures listed in this table.
- Summary of Data Catalog Views
  Data Catalog integration with Autonomous Database provides numerous tables and views.

Data Catalog Users and Roles

The `DBMS_DCAT` package supports synced users/schemas, `dcat_admin` users and local users. Users must have the `dcat_sync` role to be able to use this package.

Data Catalog Users

- Synced users/schemas
  The synced external tables are organized into database schemas corresponding to Data Asset/Bucket combinations, or according to custom properties set by the user. The synced schemas are automatically created/dropped during Data Catalog synchronization. They are created as no authentication users without the CREATE SESSION privilege. The synced schemas are also created using the protected clause, so that they cannot be altered by local users (not even the PDB admin) and can only be modified through the synchronization.
- User `dcat_admin`

User `dcat_admin` is a local database user that can run a sync and grant READ privilege on synced tables to other users or roles. The user is created as a no authentication user without the CREATE SESSION privilege.

- Local users
  Database users querying the external tables must be explicitly granted READ privileges on the synced external tables by users `dcat_admin` or ADMIN. By default, after the sync is completed, only users `dcat_admin` and ADMIN have access to the synced external tables.

Data Catalog Roles

- `dcat_sync`
  The `dcat_sync` role has all the required privileges for using the `DBMS_DCAT` package. Users must have this role to be able to use the API for navigating the Data Catalog and running the sync.

Required Credentials and IAM Policies

This topic describes the Oracle Cloud Infrastructure Identity and Access Management (IAM) user credentials and policies required to give Autonomous Database users permission to manage a data catalog and to read from object storage.

OCI Data Catalog Credential and Policy Requirements:

- A credential object with permission to manage a Data Catalog instance is required. Credential objects containing OCI native authentication or resource principals credentials are supported. Credential objects based on authentication token user principals are not supported.
  For information on managing credentials, see DBMS_CLOUD for Access Management.

  For OCI native authentication examples, see Example: Creating an OCI Native Authentication Credential Object and Autonomous Database Now Supports Accessing the Object Storage with OCI Native Authentication.

  For examples using resource principal, see Example: Using Autonomous Database Resource Principal and Accessing Oracle Cloud Infrastructure Resources from Your Autonomous Database using Resource Principal.
- The manage Data Catalog privilege is required in order for Autonomous Database to add custom properties to the Data Catalog namespace. These privileges allow you to override schema names, table names, column names and more.

  For further information on Data Catalog permissions, see Permissions Required for Each API Operation.
- The read object storage privilege on buckets is required so that Autonomous Database can query data files.

  For further Oracle Object Storage Policy Examples, see Policy Examples.

AWS Glue Data Catalog Credential and Policy Requirements

The following user credentials and policies are required to give Autonomous Database users permission to access Amazon Web Services (AWS) Glue Data Catalogs and to read from the S3 object storage:

- A credential object with permission to access an AWS Glue Data Catalog is required. For information on managing credentials, see DBMS_CLOUD for Access Management.
  For accessing an AWS Glue Data Catalog the following privileges are required: glue:GetDatabases , glue:GetTables , and glue:GetTable.

In addition, privilege s3:GetBucketLocation is needed during synchronization for generating resolvable https urls pointing to the underlying S3 objects.

- A credential object with permission to access the files stored in S3 is required so that Autonomous Database can query data files.
- AWS credentials are supported. AWS Amazon Resource Names (ARN) credentials are not supported.

Example: Creating an OCI Native Authentication Credential Object

In this example, we create an OCI native authentication credential that can be used when creating a data catalog or an object store credential object. For more details, see DBMS_DCAT SET_DATA_CATALOG_CREDENTIAL Procedure and DBMS_DCAT SET_OBJECT_STORE_CREDENTIAL Procedure respectively.

In OCI native authentication, the `DBMS_CLOUD.CREATE_CREDENTIAL` procedure includes these parameters: `credential_name`, `user_ocid`, `tenancy_ocid`, `private_key`, and `fingerprint`. See DBMS_CLOUD CREATE_CREDENTIAL Procedure for a complete description of this procedure.

The `credential_name` is the name of the credential object. The `user_ocid` and `tenancy_ocid` parameters correspond to the user's and tenancy's OCIDs respectively.

The `private_key` parameter specifies the generated private key in PEM format. Private keys created with a passphrase are not supported. Therefore, we need to make sure we generate a key with no passphrase. See How to Generate an API Signing Key for more details on how to create a private key with no passphrase. Also, the private key that we provide for this parameter must only contain the key itself without any header or footer (e.g. '-----BEGIN RSA PRIVATE KEY-----', '-----END RSA PRIVATE KEY-----').

The `fingerprint` parameter specifies the fingerprint that is obtained either after uploading the public key to the console, or using the OpenSSL commands. See How to Upload the Public Key and How to Get the Key's Fingerprint for further details on obtaining the fingerprint.

Once all the necessary information is gathered and the private key is generated, we're ready to run the following `CREATE_CREDENTIAL` procedure:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL (
    credential_name => 'OCI_NATIVE_CRED',
    user_ocid             =>
'ocid1.user.oc1..aaaaaaaatfn77fe3fxux3o5lego7glqjejrzjsqsrs64f4jsjrhbsk5qzndq'
,
    tenancy_ocid          =>
'ocid1.tenancy.oc1..aaaaaaaapwkfqz3upqklvmelbm3j77nn3y7uqmlsod75rea5zmtmbl574v
e6a',
    private_key           =>
'MIIEogIBAAKCAQEA...t9SH7Zx7a5iV7QZJS5WeFLMUEv+YbYAjnXK+dOnPQtkhOblQwCEY3Hsblj
7Xz7o=',
    fingerprint           =>
'4f:0c:d6:b7:f2:43:3c:08:df:62:e3:b2:27:2e:3c:7a');
END;
/
```

After creating the credential object, it displays in the `dba_credentials` table:

```
SELECT owner, credential_name
FROM dba_credentials
WHERE credential_name LIKE '%NATIVE%';

OWNER CREDENTIAL_NAME
----- ---------------
ADMIN OCI_NATIVE_CRED
```

Example: Using Autonomous Database Resource Principal

In this example, a dynamic group is created that includes appropriate resource members, the dynamic group is given permission to manage a Data Catalog, and then the dynamic group is given permission to read from object storage.

1. Create a dynamic group named `adb-grp-1`. Add a matching rule to `adb-grp-1` that includes the Autonomous Database instance with OCID `ocid1.autonomousdatabase.oc1.iad.abuwcljr...fjkfe` as a resource member.

   Dynamic group matching rule:

   ```
   resource.id = 'ocid1.autonomousdatabase.oc1.iad.abuwcljr...fjkfe'
   ```

2. Define a policy granting the `adb-grp-1` dynamic group full access to the Data Catalog instances, in the `mycompartment` compartment.

   ```
   allow dynamic-group adb-grp-1 to manage data-catalog-family in compartment
   mycompartment
   ```

3. Define a policy that allows the `adb-grp-1` dynamic group to read any bucket in the compartment named `mycompartment`.

   ```
   allow dynamic-group adb-grp-1 to read objects in compartment mycompartment
   ```

Example: Using User Principals

In this example, `user1` is a member of the group `adb-admins`. All members of this group are given permission to manage all data catalogs in `mycompartment`, and to read from object-store in `mycompartment`.

1. Allow users that are members of `adb-admins` to manage all data catalogs within `mycompartment`.

   ```
   allow group adb-admins to manage data-catalog-family in compartment
   mycompartment
   ```

2. Allow users that are members of `adb-admins` to read any object in any bucket within `mycompartment`.

   ```
   allow group adb-admins to read objects in compartment mycompartment
   ```

Summary of Connection Management Subprograms

This table lists the `DBMS_DCAT` package procedures used to create, query and drop Data Catalog connections.

Subprogram	Description
SET_DATA_CATALOG_CONN Procedure	Create a connection to the given data catalog
SET_DATA_CATALOG_CREDENTIAL Procedure	Set the data catalog access credential used by a specific connection to the data catalog
SET_OBJECT_STORE_CREDENTIAL Procedure	Set the credential used by the given unique connection identifier for accessing the Object Store
UNSET_DATA_CATALOG_CONN Procedure	Remove an existing Data Catalog connection

- SET_DATA_CATALOG_CREDENTIAL Procedure
  This procedure sets the Data Catalog access credential used by a specific connection to the Data Catalog.
- SET_OBJECT_STORE_CREDENTIAL Procedure
  This procedure sets the credential that is used by the given unique connection identifier for accessing the Object Store. Changing the Object Store access credential alters all existing synced tables to use the new credential.
- SET_DATA_CATALOG_CONN Procedure
  This procedure creates a connection to the given Data Catalog. The connection is required to synchronize metadata with Data Catalog. An Autonomous Database instance can connect to multiple Data Catalog instances and supports connecting to OCI Data Catalogs and AWS Glue Data Catalogs.
- UNSET_DATA_CATALOG_CONN Procedure
  This procedure removes an existing Data Catalog connection.

SET_DATA_CATALOG_CREDENTIAL Procedure

This procedure sets the Data Catalog access credential used by a specific connection to the Data Catalog.

Syntax

```
PROCEDURE DBMS_DCAT.SET_DATA_CATALOG_CREDENTIAL(
    credential_name VARCHAR2(128) DEFAULT NULL,
    dcat_con_id     VARCHAR2 DEFAULT NULL
  );
```

Parameters

Parameter	Description
`credential_name`	(Optional) The credential used for accessing the Data Catalog.
`dcat_con_id`	The unique Data Catalog connection identifier. This credential is used for the connection identified by `dcat_con_id`. The default is Null.

Usage

This credential must have Manage Data Catalog permissions; see Data Catalog Policies. The default is the resource principal; see Access Cloud Resources by Configuring Policies and Roles.

SET_OBJECT_STORE_CREDENTIAL Procedure

This procedure sets the credential that is used by the given unique connection identifier for accessing the Object Store. Changing the Object Store access credential alters all existing synced tables to use the new credential.

Syntax

```
PROCEDURE DBMS_DCAT.SET_OBJECT_STORE_CREDENTIAL(
    credential_name  VARCHAR2(128),
    dcat_con_id      IN VARCHAR2 DEFAULT NULL
  );
```

Parameters

Parameter	Description
`credential_name`	The credential used by the external tables for accessing the Object Store.
`dcat_con_id`	The unique Data Catalog connection identifier. The default is NULL.

SET_DATA_CATALOG_CONN Procedure

This procedure creates a connection to the given Data Catalog. The connection is required to synchronize metadata with Data Catalog. An Autonomous Database instance can connect to multiple Data Catalog instances and supports connecting to OCI Data Catalogs and AWS Glue Data Catalogs.

Syntax

```
PROCEDURE DBMS_DCAT.SET_DATA_CATALOG_CONN (
   region       VARCHAR2 DEFAULT NULL,
   endpoint     VARCHAR2 DEFAULT NULL,
   catalog_id   VARCHAR2 DEFAULT NULL,
   dcat_con_id  VARCHAR2 DEFAULT NULL,
   catalog_type VARCHAR2 DEFAULT NULL
 );
```

Parameters

Parameter	Description
`region`	The Data Catalog region. If the `endpoint` is specified, `region` is optional. If both `endpoint` and `region` are given, then the `endpoint` takes precedence. Default is NULL.
`endpoint`	The Data Catalog endpoint. If the `region` is specified, `endpoint` is optional. If both `endpoint` and `region` are given, then the `endpoint` takes precedence. Default is NULL.

Parameter	Description
catalog_id	The unique Oracle Cloud Identifier (OCID) for the Data Catalog instance. When connecting to AWS Glue Data Catalogs, catalog_id is optional.
dcat_con_id	A unique Data Catalog connection identifier. This identifier is required when connecting to multiple Data Catalogs and is optional when connecting to only one. It is used to refer to the Data Catalog connection in subsequent calls or when querying views. If no identifier is specified this procedure generates a NULL connection identifier. The following restrictions apply for dcat_con_id: • It must be unique within the Autonomous Database instance. • It must start with a letter. • It may contain alphanumeric characters, underscores (_), dollar signs ($), and pound signs (#). • It must be at least 16 characters long.
catalog_type	The type of data catalog to connect. Allowed values: • OCI_DCAT - OCI Data Catalog • AWS_GLUE - AWS Glue Data Catalog • NULL - The catalog type is automatically detected from the provided region or endpoint.

Usage

You only need to call this procedure once to set the connection. As part of the connection process, Autonomous Database adds custom properties to Data Catalog. These custom properties are accessible to Data Catalog users and allow you to override default names (for schemas, tables and columns) and column data types.

Before creating a connection, credentials must be created and set. For a description of the connection process, see Typical Workflow with Data Catalog for OCI Data Catalogs and User Workflow for Querying with AWS Glue Data Catalog for AWS Glue Data Catalogs.

Example: Connecting with a known OCID

In this example, Autonomous Database is connecting to Data Catalog in the uk-london-1 region. The catalog_id parameter uses the Oracle Cloud Identifier (ocid) for the Data Catalog instance. The type of Data Catalog is automatically determined: AWS Glue Data Catalog or OCI Data Catalog.

```
BEGIN
  DBMS_DCAT.SET_DATA_CATALOG_CONN(
    region=>'uk-london-1',
    catalog_id=>'ocid1.datacatalog.oc1.uk-london-1...');
END;
/
```

Example: Connecting to an AWS Glue Data Catalog

A connection is the association between an Autonomous Database instance and an AWS Glue Data Catalog. After a successful connection, the Autonomous Database instance is able to synchronize with AWS Glue. Each AWS account has one AWS Glue Data Catalog per region and each catalog can be accessed using the corresponding service endpoint for each region. An Autonomous Database instance can be associated with an AWS Glue Data Catalog by invoking the API DBMS_DCAT.SET_DATA_CATALOG_CONN and specify the endpoint for the region where the catalog resides.

See AWS Glue endpoints and quotas.

In this example, Autonomous Database is connecting to an AWS Glue Data Catalog in the `uk-london-1` region. Because this is an AWS Glue Data Catalog connection, the `catalog_id` parameter is not needed.

```
BEGIN
  DBMS_DCAT.SET_DATA_CATALOG_CONN(
    region=>'uk-london-1',
    catalog_type=>'AWS_GLUE'
END;
/
```

UNSET_DATA_CATALOG_CONN Procedure

This procedure removes an existing Data Catalog connection.

> **Note:**
>
> Invoking this procedure drops all of the protected schemas and external tables that were created as part of previous synchronizations. It does not impact the metadata in Data Catalog.

Syntax

```
PROCEDURE DBMS_DCAT.UNSET_DATA_CATALOG_CONN (
    dcat_con_id IN VARCHAR2 DEFAULT NULL
);
```

Parameters

Parameter	Description
`dcat_con_id`	The unique Data Catalog connection identifier. Default is Null.

Summary of Synchronization Subprograms

Running a synchronization, creating and dropping a synchronization job, and dropping synchronized schemas can be performed with the procedures listed in this table.

> **Note:**
>
> On April 4, 2022, the `sync_option` and `grant_read` parameters were added to the `DBMS_DCAT.RUN_SYNC` procedure. To ensure correct performance of scheduled sync jobs created prior to that date, you need to drop and recreate the scheduled sync jobs. See DBMS_DCAT.DROP_SYNC_JOB Procedure and DBMS_DCAT.CREATE_SYNC_JOB Procedure.

Subprogram	Description
CREATE_SYNC_JOB Procedure	Create a scheduler job to invoke RUN_SYNC periodically
DROP_SYNC_JOB Procedure	Drop an existing sync job for the given unique connection identifier
DROP_SYNCED_SCHEMAS Procedure	Drop all previously synchronized schemas for the given unique connection identifier
RUN_SYNC Procedure	Run a synchronization operation

- RUN_SYNC Procedure
  This procedure runs a synchronization operation and is the entry point to the synchronization. As an input, it takes lists of selected data catalog assets, folders and entities and materializes them by creating, dropping, and altering external tables.
- CREATE_SYNC_JOB Procedure
  This procedure creates a scheduler job to invoke RUN_SYNC periodically.
- DROP_SYNC_JOB Procedure
  This procedure drops an existing sync job for the given unique connection identifier.
- DROP_SYNCED_SCHEMAS Procedure
  This procedure drops all previously synchronized schemas for the given unique connection identifier.

RUN_SYNC Procedure

This procedure runs a synchronization operation and is the entry point to the synchronization. As an input, it takes lists of selected data catalog assets, folders and entities and materializes them by creating, dropping, and altering external tables.

The sync_option parameter specifies which operation the RUN_SYNC procedure performs: SYNC, DELETE or REPLACE. The operation is performed over entities within the scope of the synced_objects parameter.

Every call to the RUN_SYNC procedure returns a unique operation_id that can be used to query the USER_LOAD_OPERATIONS view to obtain information about the status of the sync and the corresponding log_table. The DBMS_DCAT$SYNC_LOG view can be queried for easy access to the log_table for the last sync operation executed by the current user. For further details, see DBMS_DCAT$SYNC_LOG View, and Monitoring and Troubleshooting Loads.

> **Note:**
>
> On April 4, 2022, the sync_option and grant_read parameters were added to the RUN_SYNC procedure. To ensure correct performance of scheduled sync jobs created prior to that date, you need to drop and recreate the scheduled sync jobs. See DBMS_DCAT.DROP_SYNC_JOB Procedure and DBMS_DCAT.CREATE_SYNC_JOB Procedure.

Synchronizing Partitioned Logical Entities or Glue Tables

The RUN_SYNC procedure creates a partitioned external table for each logical entity or Glue table when all three of the following apply:

1. The OCI data catalog logical entity or Glue table has one or more partitioned attributes.

2. For OCI data catalogs, the logical entity is derived from a prefix-based filename pattern. Partitioned logical entities derived from regex-based patterns are not supported.
3. For OCI data catalogs, the logical entity is based on partitioned data that follows the hive-style or non-hive folder format. Logical entities based on partitioned data that follow the non-hive style format using object names are not supported.

   - Example 1. Logical entities based on harvested objects that follow the Hive style partitioning format with prefix-based filename patterns.

     Consider the following objects:

     ```
     Bucket: MYBUCKET
     cluster1/db1.db/sales/country=USA/year=2020/month=01/sales1.csv
     cluster1/db1.db/sales/country=USA/year=2020/month=01/sales2.csv
     cluster1/db1.db/sales/country=USA/year=2020/month=02/sales1.csv
     ```

     Harvesting the bucket using a filename pattern with a starting folder prefix of `cluster1/db1.db` generates a logical entity named `SALES` with three partition attributes: `country`, `year`, and `month`. The type for partitioned attributes is `Partition` while the type for non-partitioned attributes is `Primitive`.

   - Example 2. Logical entities based on harvested objects that follow the non-Hive style partitioning format with prefix-based filename patterns.
     Consider the following objects:

     ```
     Bucket: MYBUCKET
     cluster2/db2.db/sales/USA/2020/01/sales1.csv
     cluster2/db2.db/sales/USA/2020/01/sales2.csv
     cluster2/db2.db/sales/USA/2020/02/sales1.csv
     ```

     Harvesting the bucket using a filename pattern with a starting folder prefix of `cluster2/db2.db` generates a logical entity named `SALES` with three partition attributes: `name0`, name1, and `name2`. The only difference between the generated logical entity compared to Example 1, is that the names of partitioned attributes are auto generated, while in Example 1 they are extracted from the URL (`country`, `year`, and `month` respectively).

For a complete end to end example of synchronizing partitioned logical entities, see Example: A Partitioned Data Scenario.

Incremental Synchronization of Partitioned Logical Entities/Glue Tables

Every call to the `RUN_SYNC` procedure specifies a set of OCI data catalog logical entities or AWS Glue tables to be synced with the database. When a logical entity or Glue table is present in two `RUN_SYNC` calls, the second call preserves and possibly alters existing external tables. The following table shows which logical entity or Glue table changes are supported when the logical entity or Glue table is partitioned:

Logical Entity or Glue Table Change	Action
Addition, removal, or update of a partition	All partitions of the external partitioned table are updated, regardless of whether a change has been detected by the data catalog.
Addition of a partitioned attribute	Adding a partitioned column to an external partitioned table is not supported. An exception is raised.

Logical Entity or Glue Table Change	Action
Deletion of a partition attribute	Dropping a partitioned column from an external partitioned table is not supported. An exception is raised.
Renaming of a partitioned attribute	Renaming a partitioned column in an external partitioned table is not supported. An exception is raised.

Syntax

```
PROCEDURE DBMS_DCAT.RUN_SYNC (
    synced_objects   IN  CLOB,
    sync_option      IN VARCHAR2 DEFAULT 'SYNC',
    error_semantics  IN VARCHAR2 DEFAULT 'SKIP_ERRORS',
    log_level        IN VARCHAR2 DEFAULT 'INFO',
    grant_read       IN VARCHAR2 DEFAULT NULL,
    dcat_con_id      IN VARCHAR2 DEFAULT NULL
 );
PROCEDURE DBMS_DCAT.RUN_SYNC (
    synced_objects   IN  CLOB,
    sync_option      IN VARCHAR2 DEFAULT 'SYNC',
    error_semantics  IN VARCHAR2 DEFAULT 'SKIP_ERRORS',
    log_level        IN VARCHAR2 DEFAULT 'INFO',
    grant_read       IN VARCHAR2 DEFAULT NULL,
    operation_id     OUT NOCOPY NUMBER,
    dcat_con_id      IN VARCHAR2 DEFAULT NULL
 );
```

Parameters

Parameter	Description
`synced_objects`	This parameter is a JSON document that specifies the data catalog objects to synchronize. For OCI Data Catalogs, the JSON document specifies a set of entities in multiple granularity: data assets, folders (Object Store buckets) or logical entities. It contains an `asset_list` that is either an array of asset objects or an array containing a single "*" string that stands for 'sync all (object store) data assets in the catalog'. For AWS Glue Data Catalogs, the JSON document specifies a list of tables in multiple granularity: databases, tables. The document specifies a list of databases. Users can restrict the set of tables to be synced by specifying individual tables within a database.

Parameter	Description
sync_option	(Optional) There are three options: • SYNC (Default) - This option ensures that what is in the data catalog, over the synced_objects scope, is represented in the Autonomous Database. If a logical entity or Glue table was deleted from the data catalog, since the last sync operation, then it is deleted in the Autonomous Database. The following operations are performed over the synced_objects scope: – Adds tables for new data catalog entities – Removes tables for deleted data catalog entities – Updates properties (such as name, columns and data types) for existing tables • DELETE - Deletes tables within the synced_objects scope. • REPLACE - Replaces all currently synced objects with the objects within the synced_objects scope.
error_semantics	(Optional) This parameter specifies the error behavior. If set to SKIP_ERRORS, the sync attempts to continue despite errors encountered for individual entities. If set to STOP_ON_ERROR, the procedure fails on the first encountered error. The default is SKIP_ERRORS.
log_level	(Optional) This parameter specifies the following values in increasing level of logging detail: (OFF, FATAL, ERROR, WARN, INFO, DEBUG, TRACE, ALL). The default is INFO.
grant_read	(Optional) This parameter is a list of users/roles that are automatically granted READ privileges on all external tables processed by this invocation of RUN_SYNC. All users/roles in the grant_read list are given READ privileges on all new or already existing external tables that correspond to entities specified by the synced_objects parameter. The RUN_SYNC procedure preserves already granted privileges on synced external tables.
operation_id	(Optional) This parameter is used to find the corresponding entry in USER_LOAD_OPERATIONS for the sync and determine the name of the log table. Note: A version of RUN_SYNC that does not return an operation_id is available so users can query USER_LOAD_OPERATIONS for the latest sync.
dcat_con_id	This parameter is the unique data catalog connection identifier that was specified when the connection to the data catalog was created. See DBMS_DCAT SET_DATA_CATALOG_CONN Procedure. This parameter identifies which connection is used for synchronization and becomes a part of the derived schema name. See Synchronization Mapping for a description of how the schema name is derived. The parameter default is NULL.

Example: Synchronize All OCI Data Catalog Entities

In the following example, all Data Catalog entities are synchronized.

```
EXEC DBMS_DCAT.RUN_SYNC(synced_objects=>'{"asset_list":["*"]}');
```

Example: synced_objects Parameter for Synchronizing All OCI Data Catalog Data Assets

The following is an example synced_objects parameter for synchronizing all (Object Storage) data assets in the Data Catalog.

```
{"asset_list" : ["*"]}
```

Example: `synced_objects` Parameter for Synchronizing Specific OCI Data Catalog Data Assets

The following is an example `synced_objects` parameter for synchronizing two data assets.

```
{"asset_list": [
        {
            "asset_id":"0b320de9-8411-4448-91fb-9e2e7f78fd5f"
        },
        {
            "asset_id":"0b320de9-8411-4448-91fb-9e2e7f74523"
        }
    ]}
```

Example: `synced_objects` Parameter for Synchronizing Specific OCI Data Catalog Entities within a Data Asset

The following shows an example `synced_objects` parameter for synchronizing two entities within the data asset.

```
{"asset_list": [
        {
            "asset_id":"0b320de9-8411-4448-91fb-9e2e7f78fd5f",
            "folder_list":[
                "f1",
                "f2"
           ]
        }
    ]}
```

Example: `synced_objects` Parameter for Synchronizing Specific OCI Data Catalog Folders and Entities within a Data Asset

The following shows an example `synced_objects` parameter for synchronizing two folders and two entities within the data asset.

```
{"asset_list":[
        {
            "asset_id":"0b320de9-8411-4448-91fb-9e2e7f78fd5f",
            "entity_list": [
                "entity1",
                "entity2"
            ],
             "folder_list": [
                "f1",
                "f2"
            ]
        }
    ]}
```

Example: `synced_objects` Parameter for Synchronizing All AWS Glue Data Catalog Databases

The following shows an example `synced_objects` parameter for synchronizing all databases in the AWS Glue Data Catalog.

```
{"database_list":["*"]}
```

Example: `synced_objects` Parameter for Synchronizing Two AWS Glue Data Catalog Databases

The following shows an example `synced_objects` parameter for synchronizing two AWS Glue Data Catalog databases.

```
{"database_list":[
    {"database":"tpcdscsv"},
    {"database":"tpcdsparquet"} ]}
```

Example: `synced_objects` Parameter for Synchronizing Three AWS Glue Data Catalog Databases

The following shows an example `synced_objects` parameter for synchronizing three tables from an AWS Glue Data Catalog database.

```
{"database_list":[
 {"database":"tpcdsparquet",
     "table_list": [ "tpcdsparquet_customer",
                     "tpcdsparquet_item",
                     "tpcdsparquet_web_sales" ] } ]}
```

CREATE_SYNC_JOB Procedure

This procedure creates a scheduler job to invoke `RUN_SYNC` periodically.

It takes as input the set of objects to be synced, the error semantics, the log level, and a repeat interval. See DBMS_DCAT RUN_SYNC Procedure for further details on how synchronization works.

There can only be a single sync job. The `CREATE_SYNC_JOB` procedure fails if another job is already specified, unless the force parameter is set to `TRUE`. If force is set to `TRUE` the previous job is dropped.

If a scheduler job attempts to run while another sync is in progress, the scheduler job fails.

> **Note:**
>
> On April 4, 2022, the `sync_option` and `grant_read` parameters were added to the `RUN_SYNC` procedure. To ensure correct performance of scheduled sync jobs created prior to that date, you need to drop and recreate the scheduled sync jobs. See DBMS_DCAT.DROP_SYNC_JOB Procedure and DBMS_DCAT.CREATE_SYNC_JOB Procedure.

Syntax

```
PROCEDURE DBMS_DCAT.CREATE_SYNC_JOB (
    synced_objects   IN CLOB,
    error_semantics  IN VARCHAR2 DEFAULT 'SKIP_ERRORS',
    log_level        IN VARCHAR2 DEFAULT 'INFO',
    repeat_interval  IN VARCHAR2,
    force            IN VARCHAR2 DEFAULT 'FALSE',
    grant_read       IN VARCHAR2 DEFAULT NULL,
    sync_option      IN VARCHAR2 DEFAULT 'SYNC',
    dcat_con_id      IN VARCHAR2 DEFAULT NULL
);
```

Parameters

Parameter	Description
synced_objects	A JSON object specifying the objects to be synced, as described in the RUN_SYNC procedure.
error_semantics	(Optional) Error behavior, as specified for RUN_SYNC. Default is SKIP_ERRORS.
log_level	(Optional) Logging level, as specified for RUN_SYNC. Default is INFO.
repeat_interval	Repeat interval for the job, with the same semantics as the repeat interval parameter of the DBMS_SCHEDULER.CREATE_JOB procedure. For details on the repeat_interval, see Overview of Creating Jobs.
force	(Optional) If TRUE, existing sync jobs are deleted first. If FALSE, the CREATE_SYNC_JOB procedure fails if a sync job already exists. Default is FALSE.
grant_read	(Optional) List of users/roles to be granted READ on the synced external tables, as described for procedure RUN_SYNC. See DBMS_DCAT.RUN_SYNC Procedure.
sync_option	(Optional) Behavior with respect to entities that have already been synced through a previous RUN_SYNC operation, as described for procedure RUN_SYNC. See DBMS_DCAT.RUN_SYNC Procedure.
dcat_con_id	This parameter is the unique Data Catalog connection identifier that was specified when the connection to Data Catalog was created. See DBMS_DCAT SET_DATA_CATALOG_CONN Procedure. This parameter identifies which connection is used for synchronization and becomes a part of the derived schema name. See Synchronization Mapping for a description of how the schema name is derived. The parameter default is NULL.

DROP_SYNC_JOB Procedure

This procedure drops an existing sync job for the given unique connection identifier.

Syntax

```
PROCEDURE DBMS_DCAT.DROP_SYNC_JOB (
    dcat_con_id IN VARCHAR2 DEFAULT NULL
  );
```

Parameters

Parameter	Description
dcat_con_id	The unique Data Catalog connection identifier. The default is NULL.

DROP_SYNCED_SCHEMAS Procedure

This procedure drops all previously synchronized schemas for the given unique connection identifier.

Syntax

```
PROCEDURE DBMS_DCAT.DROP_SYNCED_SCHEMAS (
    dcat_con_id IN VARCHAR2 DEFAULT NULL
  );
```

Parameters

Parameter	Description
dcat_con_id	The unique Data Catalog connection identifier. The default is NULL.

Summary of Data Catalog Views

Data Catalog integration with Autonomous Database provides numerous tables and views.

These tables and views help you understand:

- Available Data Catalog assets. Get information about any type of Data Catalog asset - including databases, object stores, and more.
- Information about the Data Catalog Object Storage assets and entities that have been synchronized with Autonomous Database. This includes details about how Data Catalog items (assets, folders and entities) map to Autonomous Database objects (i.e. schemas and external tables).
- Metadata sync executions. Review details about sync jobs, including any issues that may have occurred during synchronization.

This table lists the tables and views provided by the DBMS_DCAT package.

View	Description
ALL_CLOUD_CATALOG_DATABASES View	Display information about OCI Data Catalog data assets and AWS Glue Data Catalog databases
ALL_CLOUD_CATALOG_TABLES View	Used to display information about data entities for OCI Data Catalogs and tables for AWS Glue Data Catalogs
ALL_DCAT_ASSETS View	List data catalog assets that this database is authorized to access
ALL_DCAT_ATTRIBUTES View	List data catalog attributes this database is authorized to access
ALL_DCAT_CONNECTIONS View	A view that contains information about the data catalog(s) connected to this instance
ALL_DCAT_ENTITIES View	Lists logical entities this database is authorized to access

View	Description
ALL_DCAT_FOLDERS View	List metadata for the Object Storage buckets containing the data files for the Logical Entities
ALL_DCAT_GLOBAL_ACCESSIBLE_CATALOGS View	List all accessible catalogs across all regions, along with the level of access privileges for each catalog
ALL_DCAT_LOCAL_ACCESSIBLE_CATALOGS View	List all accessible catalogs in the current region, along with the level of access privileges for each catalog
ALL_GLUE_DATABASES View	Lists the AWS Glue Data Catalog databases that the data catalog credential is authorized to access
ALL_GLUE_TABLES View	Shows all AWS Glue Data Catalog tables that the data catalog credential is authorized to access
DCAT_ATTRIBUTES View	List the mapping of logical entity attributes to external table columns
DCAT_ENTITIES View	Describes the mapping of logical entities to external tables
DBMS_DCAT$SYNC_LOG View	Provides easy access to the log table for the last sync operation executed by the current user

- ALL_CLOUD_CATALOG_DATABASES View
  Use view `ALL_CLOUD_CATALOG_DATABASES` to display information about OCI Data Catalog data assets and AWS Glue Data Catalog databases.
- ALL_CLOUD_CATALOG_TABLES View
  View `ALL_CLOUD_CATALOG_TABLES` is used to display information about data entities for OCI Data Catalogs and tables for AWS Glue Data Catalogs.
- ALL_DCAT_ASSETS View
  The Data Catalog assets that this database is authorized to access.
- ALL_DCAT_ATTRIBUTES View
  The Data Catalog attributes this database is authorized to access.
- ALL_DCAT_CONNECTIONS View
  A view that contains information about the data catalog(s) connected to this instance.
- ALL_DCAT_ENTITIES View
  The Data Catalog logical entities this database is authorized to access.
- ALL_DCAT_FOLDERS View
  Metadata for the Object Storage buckets containing the data files for the Logical Entities.
- ALL_DCAT_GLOBAL_ACCESSIBLE_CATALOGS View
  This view lists all accessible catalogs across all regions, along with the level of access privileges for each catalog.
- ALL_DCAT_LOCAL_ACCESSIBLE_CATALOGS View
  This view lists all accessible catalogs in the current region, along with the level of access privileges for each catalog.
- ALL_GLUE_DATABASES View
  The AWS Glue Data Catalog databases that the data catalog credential is authorized to access.
- ALL_GLUE_TABLES View
  This view shows all AWS Glue Data Catalog tables that the data catalog credential is authorized to access.
- DCAT_ATTRIBUTES View
  Lists the mapping of logical entity attributes to external table columns.
- DCAT_ENTITIES View
  Describes the mapping of logical entities to external tables.

- DBMS_DCAT$SYNC_LOG View
  The DBMS_DCAT$SYNC_LOG view provides easy access to the log table for the last sync operation executed by the current user.

ALL_CLOUD_CATALOG_DATABASES View

Use view ALL_CLOUD_CATALOG_DATABASES to display information about OCI Data Catalog data assets and AWS Glue Data Catalog databases.

Column	Description
DCAT_CON_ID	CON1
CATALOG_ID	Data catalog unique identifier. **OCI Data Catalog example**: ocid1.datacatalog.oc1.ap-mumbai-1.….y35a **AWS Glue Data Catalog example**: NULL 579294766787
NAME	Name of the data asset (OCI)/ database (AWS Glue). **OCI Data Catalog example**: OBJECT_STORE_AT_ASHBURN **AWS Glue Data Catalog example**: OBJECT_STORE_AT_N_CALIFORNIA
DESCRIPTION	Description of the data asset (OCI)/ database (AWS Glue). **OCI Data Catalog example**: Data stored in S3 (N. California) **AWS Glue Data Catalog example**: Data stored in S3 (N. California)
TIME_CREATED	The date and time the data asset (OCI) / databases (AWS Glue) were created in the data catalog. **OCI Data Catalog example**: 26-SEP-22 10.56.01.395000 PM +00:00 **AWS Glue Data Catalog example**: 2022-06-15T09:45:35+01:00

Column	Description

DETAILS	JSON document with metadata about each data entity (OCI) / database (AWS Glue).
	OCI Data Catalog example:
	
	`{`
	`  "catalog-id": "ocid1.datacatalog.oc1.ap-mumbai-1.amaaa...",`
	`  "description": null,`
	`  "display-name": "OBJECT_STORE_AT_ASHBURN",`
	`  "external-key": "https://swiftobjectstorage.us-ashburn-1....",`
	`  "key": "bc95181c-3ac3-4959-9e5f-4e460d3fb82a",`
	`  "lifecycle-state": "ACTIVE",`
	`  "time-created": "2022-09-26T22:56:01.395000+00:00",`
	`  "type-key": "3ea65bc5-f60d-477a-a591-f063665339f9",`
	`  "uri": "/dcat/20190325/dataAssets/bc95181c-3ac3-4959-9e5f-4e460d3fb82a"`
	`}`
	
	AWS Glue Data Catalog example:
	
	`{`
	`    "Name": "dbmsdcatpoc",`
	`    "Parameters": {`
	`        "somekey": "somevalue"`
	`    },`
	`    "CreateTime": "2022-06-15T09:45:35+01:00",`
	`    "CreateTableDefaultPermissions": [`
	`        {`
	`            "Principal": {`
	`                "DataLakePrincipalIdentifier": "IAM_ALLOWED_PRINCIPALS"`
	`            },`
	`            "Permissions": [`
	`                "ALL"`
	`            ]`
	`        }`
	`    ],`
	`    "CatalogId": "579294766787"`
	`}`

ALL_CLOUD_CATALOG_TABLES View

View ALL_CLOUD_CATALOG_TABLES is used to display information about data entities for OCI Data Catalogs and tables for AWS Glue Data Catalogs.

Column	Description
DCAT_CON_ID	Unique identifier of the data catalog. The connection id. **OCI Data Catalog example**: CON1 **AWS Glue Data Catalog example**: CON1

Column	Description
CATALOG_ID	Data catalog unique identifier. **OCI Data Catalog example**: ocid1.datacatalog.oc1.ap-mumbai-1.….y35a **AWS Glue Data Catalog example**: NULL 579294766787
DATABASE_NAME	Name of the data asset (OCI)/ database (AWS Glue). **OCI Data Catalog example**: OBJECT_STORE_AT_ASHBURN **AWS Glue Data Catalog example**: OBJECT_STORE_AT_N_CALIFORNIA
NAME	Name of the data entity (OCI) / table (AWS Glue). **OCI Data Catalog example**: BIKES_TRIPS **AWS Glue Data Catalog example**: BIKES_TRIPS
DESCRIPTION	Description of the data entity (OCI) / table (AWS Glue). **OCI Data Catalog example**: Table storing bike trips **AWS Glue Data Catalog example**: Table storing bike trips
TIME_CREATED	The date and time the data entity (OCI) / table (AWS Glue) was created in the data catalog. **OCI Data Catalog example**: 26-SEP-22 10.56.01.395000 PM +00:00 **AWS Glue Data Catalog example**: 2022-06-15T09:45:35+01:00
TIME_UPDATED	Last time a change was made to the data entity (OCI) / table (AWS Glue). **OCI Data Catalog example**: 26-SEP-22 10.56.01.395000 PM +00:00 **AWS Glue Data Catalog example**: 2022-06-15T09:45:35+01:00

Column	Description

DETAILS	JSON document with metadata about each each data entity (OCI) / table (AWS Glue)
	OCI Data Catalog example:
	
	`{`
	`  "business-name": null,`
	`  "data-asset-key": "bc95181c-3ac3-4959-9e5f-...",`
	`  "description": null,`
	`  "display-name": "bikes_trips",`
	`  "external-key": "LE: https://swiftobjectstorage.us-ashburn-1.oraclecloud.com/v1/..._trips",`
	`  "folder-key": "9c4b542d-d6eb-4b83-bf59-...",`
	`  "folder-name": "hive",`
	`  "is-logical": true,`
	`  "is-partition": false,`
	`  "key": "fde30a69-a07c-478a-ab62-...",`
	`  "lifecycle-state": "ACTIVE",`
	`  "object-storage-url": "https://objectstorage.us-ashburn-1.oraclecloud.com/n/...",`
	`  "path": "OBJECT_STORE_AT_ASHBURN/hive/hive",`
	`  "pattern-key": "db21b3f1-1508-4045-aa80-...",`
	`  "properties": {`
	`  "default": {`
	`  "CONTENT-LENGTH": "4310321",`
	`  "LAST-MODIFIED": "Fri, 9 Oct 2020 20:16:52 UTC",`
	`  "archivedPECount": "0",`
	`  "dataEntityExpression": "{logicalEntity:[^/]+}.db/{logicalEntity:[^/]+}/.*",`
	`  "harvestedFile": "bikes.db/trips/p_start_month=2019-09/000000_0",`
	`  "patternName": "bikes_trips"`
	`  },`
	`  "harvestProps": {`
	`  "characterset": "UTF8",`
	`  "compression": "none",`
	`  "type": "PARQUET"`
	`  }`
	`  },`
	`  "realized-expression": "bikes.db/trips/.*",`
	`  "time-created": "2022-09-26T22:56:35.063000+00:00",`
	`  "time-updated": "2022-09-26T22:56:35.063000+00:00",`
	`  "type-key": "6753c3af-7f88-44b9-be52-1d57bef462fb",`
	`  "updated-by-id": "ocid1.user.oc1..r5l3tov7a",`
	`  "uri": "/dcat/20190325/dataAssets/bc95181c-3ac3-4959-9e5f-..."`
	`}`
	
	AWS Glue Data Catalog example:
	
	`{`
	`  "Name": "bikes_trips",`

Column	Description


```
    "DatabaseName": "dbmsdcatpoc",
    "Owner": "owner",
    "CreateTime": "2022-06-23T13:24:20+01:00",
    "UpdateTime": "2022-06-23T13:24:20+01:00",
    "LastAccessTime": "2022-06-23T13:24:20+01:00",
    "Retention": 0,
    "StorageDescriptor": {
        "Columns": [
            {
                "Name": "trip_duration",
                "Type": "int"
            },
            {
                "Name": "start_month",
                "Type": "string"
            }, ...
        ],
        "Location": "s3://dbmsdcatpoc/hive/bikes.db/trips/",
        "InputFormat":
"org.apache.hadoop.hive.ql.io.parquet.MapredParquetInputForma
t",
        "OutputFormat":
"org.apache.hadoop.hive.ql.io.parquet.MapredParquetOutputForm
at",
        "Compressed": false,
        "NumberOfBuckets": -1,
        "SerdeInfo":
 {            "SerializationLibrary":
"org.apache.hadoop.hive.ql.io.parquet.serde.ParquetHiveSerDe"
,
            "Parameters": {
                "serialization.format": "1"
            }
        },
        "BucketColumns": [],
        "SortColumns": [],
        "Parameters": {
            "CrawlerSchemaDeserializerVersion": "1.0",
            "CrawlerSchemaSerializerVersion": "1.0",
            "UPDATED_BY_CRAWLER": "crawler-bikes",
            "averageRecordSize": "86",
            "classification": "parquet",
            "compressionType": "none",
            "objectCount": "12",
            "recordCount": "404947",
            "sizeKey": "35312159",
            "typeOfData": "file"
        },
        "StoredAsSubDirectories": false
    },
    "PartitionKeys": [
```

Column	Description

	{
	"Name": "p_start_month",
	"Type": "string"
	}
	],
	"TableType": "EXTERNAL_TABLE",
	"Parameters": {
	"CrawlerSchemaDeserializerVersion": "1.0",
	"CrawlerSchemaSerializerVersion": "1.0",
	"UPDATED_BY_CRAWLER": "crawler-bikes",
	"averageRecordSize": "86",
	"classification": "parquet",
	"compressionType": "none",
	"objectCount": "12",
	"recordCount": "404947",
	"sizeKey": "35312159",
	"typeOfData": "file"
	},
	"CreatedBy": "arn:aws:sts::579294766787:assumed-role/
	AWSGlueServiceRole-dbmsdcat/AWS-Crawler",
	"IsRegisteredWithLakeFormation": false,
	"CatalogId": "579294766787",
	"VersionId": "0"
	}

Example

ALL_DCAT_ASSETS View

The Data Catalog assets that this database is authorized to access.

Column	Datatype	Description
DCAT_CON_ID	VARCHAR2 (4000)	Connection identifier that is unique within the instance
KEY	VARCHAR2(4000)	Asset key
DISPLAY_NAME	VARCHAR2(4000)	Asset display name
DESCRIPTION	VARCHAR2(4000)	Asset description
CATALOG_ID	VARCHAR2(4000)	OCID for the Data Catalog containing the asset
EXTERNAL_KEY	VARCHAR2(4000)	Base Object Storage URI for the asset
URI	VARCHAR2(4000)	Asset URI for the Data Catalog API
TIME_CREATED	TIMESTAMP(6) WITH TIMEZONE	The date and time the data asset was created
TYPE_KEY	VARCHAR2(4000)	The key of the data asset type (currently, only Object Storage data assets are supported). Type keys can be found via the '/types' Data Catalog endpoint.

Column	Datatype	Description
LIFECYCLE_STATE	VARCHAR2(4000)	The current state of the data asset. For more information on possible life cycle states, see the Data Catalog DataAsset Reference for a list of possible states for `lifecycleState`.

ALL_DCAT_ATTRIBUTES View

The Data Catalog attributes this database is authorized to access.

Column	Datatype	Description
DCAT_CON_ID	VARCHAR2 (4000)	Connection identifier that is unique within the instance
KEY	NUMBER	Attribute key
DISPLAY_NAME	VARCHAR2(4000)	Attribute display name
BUSINESS_NAME	VARCHAR2(4000)	Attribute business name
DESCRIPTION	VARCHAR2(4000)	Attribute description
DATA_ASSET_KEY	VARCHAR2(4000)	Data asset key
FOLDER_KEY	VARCHAR2(4000)	Folder key
ENTITY_KEY	VARCHAR2(4000)	Entity key
EXTERNAL_KEY	VARCHAR2(4000)	Unique external key for the attribute
LENGTH	NUMBER	Maximum allowed length of the attribute value
PRECISION	NUMBER	Precision of the attribute value (usually applies to float data type)
SCALE	NUMBER	Scale of the attribute value (usually applies to float data type)
IS_NULLABLE	NUMBER	Identifies if this attribute can be assigned null values
URI	VARCHAR2(4000)	URI to the attribute instance in the Data Catalog API
LIFECYCLE_STATE	VARCHAR2(4000)	The current state of the attribute. For more information on possible life cycle states, see the Data Catalog Attribute Reference for a list of possible states for `lifecycleState`.
TIME_CREATED	TIMESTAMP(6) WITH TIME ZONE	The date and time the attribute was created
EXTERNAL_DATA_TYPE	VARCHAR2(4000)	Data type of the attribute as defined in the external system
MIN_COLLECTION_COUNT	NUMBER	Minimum number of elements, if the type of the attribute is a collection type
MAX_COLLECTION_COUNT	NUMBER	Maximum number of elements, if the type of the attribute is a collection type
DATATYPE_ENTITY_KEY	VARCHAR2(4000)	Entity key that represents the datatype of this attribute, applicable if this attribute is a complex type
EXTERNAL_DATATYPE_ENTITY_KEY	VARCHAR2(4000)	External entity key that represents the datatype of this attribute, applicable if this attribute is a complex type

Column	Datatype	Description
PARENT_ATTRIBUTE_KEY	VARCHAR2(4000)	Attribute key that represents the parent attribute of this attribute, applicable if the parent attribute is of complex datatype
EXTERNAL_PARENT_ATTRIBUTE_KEY	VARCHAR2(4000)	External attribute key that represents the parent attribute of this attribute, applicable if the parent attribute is of complex type
PATH	VARCHAR2(4000)	Full path of the attribute

ALL_DCAT_CONNECTIONS View

A view that contains information about the data catalog(s) connected to this instance.

Column	Datatype	Description
DCAT_CON_ID	VARCHAR2(4000)	Connection identifier that is unique within the instance
COMPARTMENT_ID	VARCHAR2(4000)	OCID for the compartment where the Data Catalog instance resides
INSTANCE_ID	VARCHAR2(4000)	OCID for the Data Catalog instance
REGION	VARCHAR2(4000)	Region for the Data Catalog instance
ENDPOINT	VARCHAR2(4000)	Endpoint for the Data Catalog instance
CREATED	TIMESTAMP	When the Data Catalog instance was created
NAME	VARCHAR2(4000)	Name of the Data Catalog instance
LAST_UPDATED	TIMESTAMP	Timestamp of the last update of the connection to the Data Catalog instance
LATEST_OPERATION_ID	NUMBER	The id of the last synchronization operation
DATA_CATALOG_CREDENTIAL	VARCHAR2(128)	Credential used for accessing the Data Catalog
OBJECT_STORE_CREDENTIAL	VARCHAR2(128)	Credential used by the external table driver for accessing the Object Store

ALL_DCAT_ENTITIES View

The Data Catalog logical entities this database is authorized to access.

Column	Datatype	Description
DCAT_CON_ID	VARCHAR2(4000)	Connection identifier that is unique within the instance
CATALOG_ID	VARCHAR2(4000)	OCID for the Data Catalog containing the asset
KEY	VARCHAR2(4000)	Entity key
DISPLAY_NAME	VARCHAR2(4000)	Entity display name
BUSINESS_NAME	VARCHAR2(4000)	Entity business name
DESCRIPTION	VARCHAR2(4000)	Logical entity description
DATA_ASSET_KEY	VARCHAR2(4000)	Asset key
FOLDER_KEY	VARCHAR2(4000)	Folder unique key
FOLDER_NAME	VARCHAR2(4000)	Folder name (bucket)

Column	Datatype	Description
EXTERNAL_KEY	VARCHAR2(4000)	External key for the logical entity
PATTERN_KEY	VARCHAR2(4000)	Key of the associated pattern for the logical entity
REALIZED_EXPRESSION	VARCHAR2(4000)	The regular expression used to obtain the files for this logical entity
PATH	VARCHAR2(4000)	Full path for the logical entity
TIME_CREATED	TIMESTAMP(6) WITH TIME ZONE	Date and time the entity was created
TIME_UPDATED	TIMESTAMP(6) WITH TIME ZONE	Last time a change was made to the data entity
UPDATED_BY_ID	VARCHAR2(4000)	OCID of the user who updated this object in the Data Catalog
URI	VARCHAR2(4000)	URI of the entity instance in the API
LIFECYCLE_STATE	VARCHAR2(4000)	The current state of the entity. For more information on possible life cycle states, see the Data Catalog Entity Reference for a list of possible states for `lifecycleState`.

ALL_DCAT_FOLDERS View

Metadata for the Object Storage buckets containing the data files for the Logical Entities.

Column	Datatype	Description
DCAT_CON_ID	VARCHAR2(4000)	Connection identifier that is unique within the instance
CATALOG_ID	VARCHAR2(4000)	OCID for the Data Catalog containing the asset
KEY	VARCHAR2(4000)	Folder key
DISPLAY_NAME	VARCHAR2(4000)	Folder display name
BUSINESS_NAME	VARCHAR2(4000)	Folder business name
DESCRIPTION	VARCHAR2(4000)	Folder description
DATA_ASSET_KEY	VARCHAR2(4000)	Key for the data asset containing the folder
PARENT_FOLDER_KEY	VARCHAR2(4000)	Key for the parent folder (currently, this is the data asset key)
PATH	VARCHAR2(4000)	Full path for the folder
EXTERNAL_KEY	VARCHAR2(4000)	Object Storage URI for the bucket
TIME_EXTERNAL	TIMESTAMP(6) WITH TIMEZONE	The last modified timestamp of this folder
TIME_CREATED	TIMESTAMP(6) WITH TIMEZONE	The date/time the folder was created
URI	VARCHAR2(4000)	URI to the folder instance in the Data Catalog API.
LIFECYCLE_STATE	VARCHAR2(4000)	The current state of the folder. For more information on possible life cycle states, see the Data Catalog Folder Reference for a list of possible states for `lifecycleState`.

ALL_DCAT_GLOBAL_ACCESSIBLE_CATALOGS View

This view lists all accessible catalogs across all regions, along with the level of access privileges for each catalog.

Column	Datatype	Description
CATALOG_ID	VARCHAR2(4000)	Catalog OCID
CATALOG_NAME	VARCHAR2(4000)	Name of the catalog
CATALOG_REGION	VARCHAR2(4000)	Name of the catalog region
CATALOG_SCORE	NUMBER	The catalog score is a numeric value calculated from the privileges configured for the Data Catalog access credential. A higher catalog score means greater privileges, which may equate to a higher likelihood that this catalog is intended for use with this Autonomous Database instance.

ALL_DCAT_LOCAL_ACCESSIBLE_CATALOGS View

This view lists all accessible catalogs in the current region, along with the level of access privileges for each catalog.

Column	Datatype	Description
CATALOG_ID	VARCHAR2(4000)	Catalog OCID
CATALOG_NAME	VARCHAR2(4000)	Name of the catalog
CATALOG_SCORE	NUMBER	The catalog score is a numeric value calculated from the privileges configured for the Data Catalog access credential. A higher catalog score means greater privileges, which may equate to a higher likelihood that this catalog is intended for use with this Autonomous Database instance.

ALL_GLUE_DATABASES View

The AWS Glue Data Catalog databases that the data catalog credential is authorized to access.

Column	Data Type	Description
DCAT_CON_ID	VARCHAR2(4000)	Unique identifier of data catalog connection id.
CATALOG_ID	VARCHAR2(255)	Data Catalog unique identifier.
NAME	VARCHAR2(255)	Name of the database.
DESCRIPTION	VARCHAR2(2048)	Description of the database.
LOCATION_URI	VARCHAR2(1024)	The location of the database.
CREATE_TIME	TIMESTAMP	The time that the database was created in the data catalog.

Column	Data Type	Description
PARAMETERS	CLOB	JSON document with key-value pairs that define parameters and properties of the database.
TARGET_DATABASE	VARCHAR2(4000)	JSON document that describes a target database for resource linking in AWS.

ALL_GLUE_TABLES View

This view shows all AWS Glue Data Catalog tables that the data catalog credential is authorized to access.

Column	Data Type	Description
DCAT_CON_ID	VARCHAR2(4000)	Unique identifier of data catalog connection id.
CATALOG_ID	VARCHAR2(255)	Catalog identifier
DATABASE_NAME	VARCHAR2(255)	Database name
NAME	VARCHAR2(255)	Table name
TABLE_TYPE	VARCHAR2(255)	Table type
CLASSIFICATION	VARCHAR2(255)	
DESCRIPTION	VARCHAR2(2048)	Table description
OWNER	VARCHAR2(255)	Table owner
CREATED_BY	VARCHAR2(255)	Table creator
CREATE_TIME	TIMESTAMP	The time the table was created in the data catalog.
LAST_ANALYZED_TIME	TIMESTAMP	The last time column statistics were computed for this table.
LAST_ACCESS_TIME	TIMESTAMP	The last time the table was accessed.
UPDATE_TIME	TIMESTAMP	The last time the table was updated.
IS_REGISTERED_WITH_LAKE_FORMATION	NUMBER	Indicates whether the table is registered with AWS lake formation.
PARAMETERS	CLOB	JSON document with key-value pairs that define properties of the table.
PARTITION_KEYS	CLOB	JSON document with a list of columns by which the table is partitioned.
RETENTION	NUMBER	The retention time for this table.
STORAGE_DESCRIPTION	CLOB	JSON document with information about the physical storage of a table.
TARGET_TABLE	VARCHAR2(4000)	JSON document describing a target table used for resource linking in AWS.
VERSION_ID	VARCHAR2(255)	The version identifier for the table.
VIEW_EXPANDED_TEXT	CLOB	Introduced by AWS Glue for compatibility with Hive. Not used by AWS Glue.
VIEW_ORIGINAL_TEXT	CLOB	Introduced by AWS Glue for compatibility with Hive. Not used by AWS Glue.

DCAT_ATTRIBUTES View

Lists the mapping of logical entity attributes to external table columns.

Column	Datatype	Description
DCAT_CON_ID	VARCHAR2 (4000)	Connection identifier that is unique within the instance
ASSET_KEY	VARCHAR2(4000)	Data Catalog asset key
ENTITY_KEY	VARCHAR2(4000)	Data Catalog entity key
ATTRIBUTE_KEY	VARCHAR2(4000)	Data Catalog attribute key
ORACLE_COLUMN_NAME	VARCHAR2(128)	Mapped column name

DCAT_ENTITIES View

Describes the mapping of logical entities to external tables.

Column	Datatype	Description
DCAT_CON_ID	VARCHAR2(4000)	Connection identifier that is unique within the instance
ASSET_KEY	VARCHAR2(4000)	Data Catalog asset key
ENTITY_KEY	VARCHAR2(4000)	Data Catalog entity key
FOLDER_KEY	VARCHAR2(4000)	Data Catalog folder key
ORACLE_TABLE_NAME	VARCHAR2(128)	Mapped table name
ORACLE_SCHEMA_NAME	VARCHAR2(128)	Mapped schema name
ENTITY_ORACLE_DB_SCHEMA	VARCHAR2(4000)	The entity's oracle-db-schema custom property used to derive the schema
ASSET_ORACLE_DB_SCHEMA	VARCHAR2(4000)	The data asset's oracle-db-schema custom property used to derive the schema
FOLDER_ORACLE_DB_SCHEMA	VARCHAR2(4000)	The folder's oracle-db-schema custom property used to derive the schema

DBMS_DCAT$SYNC_LOG View

The DBMS_DCAT$SYNC_LOG view provides easy access to the log table for the last sync operation executed by the current user.

Every call to the RUN_SYNC procedure is logged to a new log table, pointed to by the LOGFILE_TABLE field of USER_LOAD_OPERATIONS. The log tables are automatically dropped after 2 days, and users can clear all sync logs using the DELETE_ALL_OPERATIONS Procedure where type is DCAT_SYNC.

The DBMS_DCAT$SYNC_LOG view automatically identifies the latest log table. The schema for the DBMS_DCAT$SYNC_LOG view is described below and the access permissions are identical to those of the individual log tables. By default READ is granted to the dbms_dcat role and to the ADMIN user.

The log tables have the following format:

Column	Datatype	Description
LOG_TIMESTAMP	TIMESTAMP	Timestamp for the log entry.
LOG_LEVEL	VARCHAR2(32)	The entry log level can have one of the following values: OFF, FATAL, ERROR, WARN, INFO, DEBUG, TRACE, ALL.
LOG_DETAILS	VARCHAR2(32767)	The log message.

DBMS_INMEMORY_ADMIN Package

The DBMS_INMEMORY_ADMIN package enables you to manage the In-Memory area and In-Memory FastStart.

- Summary of DBMS_INMEMORY_ADMIN Subprograms
  This table summarizes the subprograms included in the DBMS_INMEMORY_ADMIN package.

Summary of DBMS_INMEMORY_ADMIN Subprograms

This table summarizes the subprograms included in the DBMS_INMEMORY_ADMIN package.

Subprogram	Description
AIM_GET_PARAMETER Procedure	This procedure obtains the current values for parameters that control Automatic In-Memory.
AIM_SET_PARAMETER Procedure	The procedure customizes the execution environment of Automatic In-Memory.
FASTSTART_DISABLE Procedure	This procedure disables the In-Memory FastStart (IM FastStart) feature.
FASTSTART_ENABLE Procedure	This procedure enables IM FastStart and assigns a tablespace.
FASTSTART_MIGRATE_STORAGE Procedure	This procedure moves all IM FastStart data and metadata from the existing tablespace to the specified new tablespace.
GET_FASTSTART_TABLESPACE Function	This function returns the name of the tablespace that is currently designated for IM FastStart.
IME_CAPTURE_EXPRESSIONS Procedure	This procedure captures the 20 most frequently accessed (“hottest”) expressions in the database in the specified time interval.
IME_CLOSE_CAPTURE_WINDOW Procedure	This procedure signals the end of the current expression capture window.
IME_DROP_ALL_EXPRESSIONS Procedure	This procedure drops all SYS_IME expression virtual columns in the database.
IME_GET_CAPTURE_STATE Procedure	This procedure returns the current capture state of the expression capture window and the timestamp of the most recent modification.
IME_OPEN_CAPTURE_WINDOW Procedure	This procedure signals the beginning of an expression capture window.
IME_POPULATE_EXPRESSIONS Procedure	This procedure forces the population of expressions captured in the latest invocation of DBMS_INMEMORY_ADMIN.IME_CAPTURE_EXPRESSIONS.
POPULATE_WAIT Function	Initiates population of all INMEMORY objects that have a priority greater than or equal to the specified priority, and sets a timeout interval within which population must occur.

Subprogram	Description
SET_SGA_PERCENTAGE Procedure	This procedure modifies the maximum In-Memory size for your database.

- SET_SGA_PERCENTAGE Procedure
  Use `DBMS_INMEMORY_ADMIN.SET_SGA_PERCENTAGE` to modify the maximum In-Memory size for your database.

SET_SGA_PERCENTAGE Procedure

Use `DBMS_INMEMORY_ADMIN.SET_SGA_PERCENTAGE` to modify the maximum In-Memory size for your database.

Syntax

```
DBMS_INMEMORY_ADMIN.SET_SGA_PERCENTAGE (
     sga_percentage IN NUMBER);
```

Parameter	Description
`sga_percentage`	This specifies a value for the maximum In-Memory size for your Autonomous Database. The supplied `sga_percentage` value must be in the range of 0 and 70. This parameter is mandatory.

Usage Note

To run `DBMS_INMEMORY_ADMIN.SET_SGA_PERCENTAGE` you must be logged in as the ADMIN user or have the `EXECUTE` privilege on `DBMS_INMEMORY_ADMIN`.

Example

```
BEGIN
     DBMS_INMEMORY_ADMIN.SET_SGA_PERCENTAGE(60);
END;
/
```

DBMS_MAX_STRING_SIZE Package

The `DBMS_MAX_STRING_SIZE` package provides an interface for checking and changing the value of the `DBMS_MAX_STRING_SIZE` initialization parameter.

- CHECK_MAX_STRING_SIZE Function
  This function checks if the `MAX_STRING_SIZE` parameter can be updated to a given value and returns a list of violations that would prevent the parameter from being updated.
- MODIFY_MAX_STRING_SIZE Procedure
  This procedure updates the value of the `MAX_STRING_SIZE` parameter to a given value.

CHECK_MAX_STRING_SIZE Function

This function checks if the `MAX_STRING_SIZE` parameter can be updated to a given value and returns a list of violations that would prevent the parameter from being updated.

Syntax

```
DBMS_MAX_STRING_SIZE.CHECK_MAX_STRING_SIZE(
       new_value    IN VARCHAR2)
RETURN DBMS_MAX_STRING_SIZE_TBL;
```

Parameters

Parameter	Description
`new_value`	Specifies the new `MAX_STRING_SIZE` parameter value to be set. The only valid value is: `'STANDARD'`.

Usage Notes

If the return list is empty, then there are no violations and the `MAX_STRING_SIZE` update can be performed.

Example

```
SELECT * FROM TABLE(DBMS_MAX_STRING_SIZE.CHECK_MAX_STRING_SIZE('STANDARD'));

TYPE   OBJECT_OWNER OBJECT_NAME COLUMN_NAME
REASON
------ ------------ ----------- --------------
-------------------------------------------
COLUMN ADMIN        SALES       CUST_NOTES     Physical column exceeds
STANDARD length limit
1 rows selected.
```

MODIFY_MAX_STRING_SIZE Procedure

This procedure updates the value of the `MAX_STRING_SIZE` parameter to a given value.

Syntax

```
DBMS_MAX_STRING_SIZE.MODIFY_MAX_STRING_SIZE(
         new_value    IN VARCHAR2);
```

Where: *user_account* is the user account name (schema name).

Parameters

Parameter	Description
`new_value`	Specifies the new `MAX_STRING_SIZE` parameter value to be set. The only valid value is: `'STANDARD'`.

Usage Notes

- Using `DBMS_MAX_STRING_SIZE.MODIFY_MAX_STRING_SIZE` is a one-way change that cannot be reverted. After a database is switched back to the `STANDARD` style of supporting a maximum length of 4000 bytes for the `VARCHAR2`, `NVARCHAR2`, and `RAW` data types, you cannot re-enable `EXTENDED` data types.
- The ADMIN user is granted `EXECUTE` privilege `WITH GRANT OPTION` clause on `DBMS_MAX_STRING_SIZE`. Oracle recommends that you do not `GRANT EXECUTE` on this package to other users.
- The error `ORA-20000` is raised if any object exists that would prevent `MAX_STRING_SIZE` from being updated.
- The ADMIN user is granted `EXECUTE` privilege `WITH GRANT OPTION` clause on `DBMS_MAX_STRING_SIZE`. Oracle recommends that you do not `GRANT EXECUTE` on this package to other users.

Example

```
SELECT NAME, VALUE FROM V$PARAMETER WHERE NAME = 'max_string_size';

NAME                 VALUE
max_string_size      EXTENDED


   BEGIN
     DBMS_MAX_STRING_SIZE.MODIFY_MAX_STRING_SIZE('STANDARD');
   END;
/


PL/SQL procedure successfully completed.

SELECT NAME, VALUE FROM V$PARAMETER WHERE NAME = 'max_string_size';

NAME                 VALUE
max_string_size      STANDARD
```

DBMS_PIPE Package

The `DBMS_PIPE` package lets two or more sessions in the same instance communicate.

Oracle Autonomous Database supports core `DBMS_PIPE` functionality as available in Oracle Database 19c, plus extensions.

See DBMS_PIPE for details about the core `DBMS_PIPE` functionality provided in Oracle Database.

- DBMS_PIPE Overview for Singleton Pipes
  Pipe functionality has several potential applications: external service interface, debugging, independent transactions, and alerts.
- Summary of DBMS_PIPE Subprograms for Singleton Pipes
  This table lists the `DBMS_PIPE` subprograms and briefly describes them.

- DBMS_PIPE Overview for Persistent Messaging Pipes
  Pipe functionality has several potential applications: external service interface, debugging, independent transactions, and alerts.
- Summary of DBMS_PIPE Subprograms for Persistent Messaging
  This table lists the `DBMS_PIPE` subprograms and briefly describes them.

DBMS_PIPE Overview for Singleton Pipes

Pipe functionality has several potential applications: external service interface, debugging, independent transactions, and alerts.

On Autonomous Database the DBMS_PIPE package has extended functionality to support singleton pipes.

Singleton pipe features in `DBMS_PIPE` provide the following:

- Ability to cache and retrieve a custom message, of up to 32,767 bytes, in Oracle database memory. The message size maximum of 32,767 bytes is applicable to all Pipes, including Singleton Pipes. Previous versions of `DBMS_PIPE` had a smaller maximum message size.
- Share the cached message across multiple database sessions with concurrent reads.
- Cache Invalidation methods:
  - Explicit cache invalidation controlled by user.
  - Cache invalidation after a user specified parameter (`shelflife`) time interval (in seconds).
- Declarative and easy to use PL/SQL APIs for caching.
- Supports both Read-Only and Read-Write databases.

A Singleton Pipe can be any one of the supported `DBMS_PIPE` types:

- **Implicit Pipe:** Automatically created when a message is sent with an unknown pipe name using the `DBMS_PIPE.SEND_MESSAGE` function.
- **Explicit Pipe:** Created using the `DBMS_PIPE.CREATE_PIPE` function with a user specified pipe name.
- **Public Pipe:** Accessible by any user with `EXECUTE` permission on `DBMS_PIPE` package
- **Private Pipe:** Accessible by sessions with the same user as the pipe creator.

Summary of DBMS_PIPE Subprograms for Singleton Pipes

This table lists the `DBMS_PIPE` subprograms and briefly describes them.

Table 6-1 DBMS_PIPE Package Subprograms

Subprogram	Description
CREATE_PIPE Function	Creates a pipe (necessary for private pipes)
NEXT_ITEM_TYPE Function	Returns datatype of next item in buffer
PACK_MESSAGE Procedures	Builds message in local buffer
PURGE Procedure	Purges contents of named pipe
RECEIVE_MESSAGE Function	Copies message from named pipe into local buffer
RESET_BUFFER Procedure	Purges contents of local buffer

Table 6-1 (Cont.) DBMS_PIPE Package Subprograms

Subprogram	Description
REMOVE_PIPE Function	Removes the named pipe
SEND_MESSAGE Function	Sends message on named pipe: This implicitly creates a public pipe if the named pipe does not exist
UNIQUE_SESSION_NAME Function	Returns unique session name
UNPACK_MESSAGE Procedures	Accesses next item in buffer

- CREATE_PIPE Function
  This function explicitly creates a public or private pipe. If the `private` flag is `TRUE`, then the pipe creator is assigned as the owner of the private pipe.
- RECEIVE_MESSAGE Function
  This function copies the message into the local message buffer.
- SEND_MESSAGE Function
  This function sends a message on the named pipe.

CREATE_PIPE Function

This function explicitly creates a public or private pipe. If the `private` flag is `TRUE`, then the pipe creator is assigned as the owner of the private pipe.

Explicitly-created pipes can only be removed by calling `REMOVE_PIPE`, or by shutting down the instance.

In order to create a Singleton Pipe, set the `singleton` parameter to `TRUE`. The following arguments are applicable to Singleton Pipes:

- `singleton`: Indicates that the pipe should be created as a Singleton Pipe (default value: `FALSE`).
- `shelflife`: Optionally specify a shelflife expiration (in seconds) of cached message in the Singleton Pipe. It can be used for implicit invalidation of message in Singleton Pipe.

  The message `shelflife` in Singleton Pipe can also be specified when you send a message message (see SEND_MESSAGE Function).

Syntax

```
DBMS_PIPE.CREATE_PIPE (
   pipename     IN VARCHAR2,
   maxpipesize  IN INTEGER DEFAULT 66536,
   private      IN BOOLEAN DEFAULT TRUE,
   singleton    IN BOOLEAN DEFAULT FALSE,
   shelflife    IN INTEGER DEFAULT 0)
RETURN INTEGER;
```

Parameters

Table 6-2 CREATE_PIPE Function Parameters

Parameter	Description
pipename	Name of the pipe you are creating. You must use this name when you call SEND_MESSAGE and RECEIVE_MESSAGE. This name must be unique across the instance. Caution: Do not use pipe names beginning with ORA$. These are reserved for use by procedures provided by Oracle. Pipename should not be longer than 128 bytes, and is case insensitive. At this time, the name cannot contain Globalization Support characters.
maxpipesize	The maximum size allowed for the pipe, in bytes. The total size of all of the messages on the pipe cannot exceed this amount. The message is blocked if it exceeds this maximum. The default maxpipesize is 66536 bytes. The maxpipesize for a pipe becomes a part of the characteristics of the pipe and persists for the life of the pipe. Callers of SEND_MESSAGE with larger values cause the maxpipesize to be increased. Callers with a smaller value use the existing, larger value. The default maxpipesize of 65536 is applicable for all Pipes.
private	Uses the default, TRUE, to create a private pipe. Public pipes can be implicitly created when you call SEND_MESSAGE.
singleton	Use TRUE to create a Singleton Pipe. Default value: FALSE
shelflife	Expiration time in seconds of a message cached in Singleton Pipe. After the specified shelflife time is exceeded, the message is no longer accessible from the Pipe. The parameter shelflife is only applicable to a Singleton Pipe. Default value is 0, implying that the message never expires.

Return Values

Table 6-3 CREATE_PIPE Function Return Values

Return	Description
0	Successful. If the pipe already exists and the user attempting to create it is authorized to use it, then Oracle returns 0, indicating success, and any data already in the pipe remains.
6	Failed to convert existing pipe to singleton pipe. • Implicit pipe with more than one existing message cannot be converted to a Singleton Pipe. • For an explicit pipe that is not Singleton, DBMS_PIPE.SEND_MESSAGE cannot send a message with singleton argument set to TRUE.
7	A non-zero value is given for the shelflife parameter and the pipe is not a singleton pipe.

Table 6-3 (Cont.) CREATE_PIPE Function Return Values

Return	Description
ORA-23322	Failure due to naming conflict. If a pipe with the same name exists and was created by a different user, then Oracle signals error ORA-23322, indicating the naming conflict.

Exceptions

Table 6-4 CREATE_PIPE Function Exception

Exception	Description
Null pipe name	Permission error: Pipe with the same name already exists, and you are not allowed to use it.

Example

Create a Singleton Pipe with shelflife of 1 hour.

```
DECLARE
  l_status INTEGER;
BEGIN
  l_status := DBMS_PIPE.create_pipe(pipename  => 'MY_PIPE1',
                                    private   => TRUE,
                                    singleton => TRUE,
                                    shelflife => 3600);
END;
/
```

RECEIVE_MESSAGE Function

This function copies the message into the local message buffer.

Syntax

```
DBMS_PIPE.RECEIVE_MESSAGE (
   pipename      IN VARCHAR2,
   timeout       IN INTEGER      DEFAULT maxwait,
   cache_func    IN VARCHAR2     DEFAULT NULL)
RETURN INTEGER;
```

Parameters

Table 6-5 RECEIVE_MESSAGE Function Parameters

Parameter	Description
pipename	Name of the pipe on which you want to receive a message. Names beginning with ORA$ are reserved for use by Oracle.

Table 6-5 (Cont.) RECEIVE_MESSAGE Function Parameters

Parameter	Description
timeout	Time to wait for a message, in seconds. A timeout of 0 lets you read without blocking. The timeout does not include the time spent in execution cache function specified in the cache_func parameter. Default value: is the constant MAXWAIT, which is defined as 86400000 (1000 days).
cache_func	Cache function name to automatically cache a message in a Singleton Pipe. The name of the function should be fully qualified with the owner schema: • *OWNER.FUNCTION_NAME* • *OWNER.PACKAGE.FUNCTION_NAME* Default value: NULL

Return Values

Table 6-6 RECEIVE_MESSAGE Function Return Values

Return	Description
0	Success
1	Timed out. If the pipe was implicitly-created and is empty, then it is removed.
2	Record in the pipe is too large for the buffer.
3	An interrupt occurred.
8	Cache function can only be specified when using a Singleton Pipe.
ORA-23322	User has insufficient privileges to read from the pipe.

Usage Notes

To receive a message from a pipe, first call RECEIVE_MESSAGE. When you receive a message, it is removed from the pipe; hence, a message can only be received once. For implicitly-created pipes, the pipe is removed after the last record is removed from the pipe.

If the pipe that you specify when you call RECEIVE_MESSAGE does not already exist, then Oracle implicitly creates the pipe and waits to receive the message. If the message does not arrive within a designated timeout interval, then the call returns and the pipe is removed.

After receiving the message, you must make one or more calls to UNPACK_MESSAGE to access the individual items in the message. The UNPACK_MESSAGE procedure is overloaded to unpack items of type DATE, NUMBER, VARCHAR2, and there are two additional procedures to unpack RAW and ROWID items. If you do not know the type of data that you are attempting to unpack, then call NEXT_ITEM_TYPE to determine the type of the next item in the buffer.

Cache Function Parameter

Singleton Pipes support cache function to automatically cache a message in the pipe in case of the following two scenarios:

- Singleton Pipe is empty.
- Message in Singleton Pipe is invalid due to `shelflife` time elapsed.

The name of the function should be fully qualified with the owner schema:

- *OWNER.FUNCTION_NAME*
- *OWNER.PACKAGE.FUNCTION_NAME*

To use a cache function the current session user that invokes `DBMS_PIPE.RECEIVE_MESSAGE` must have required privileges to execute the cache function.

Cache Function Syntax

```
CREATE OR REPLACE FUNCTION cache_function_name (
       pipename  IN VARCHAR2
) RETURN INTEGER;
```

Parameter	Datatype	Description
pipename	VARCHAR2	Name of the Singleton Pipe.

Return	Description
0	Success
Non-zero	Failure value returned from DBMS_PIPE.RECEIVE_MESSAGE

Define a cache function to provide encapsulation and abstraction of complexity from the reader sessions of Singleton Pipe. The typical operations within a cache function would be:

- Create a Singleton Pipe, for an Explicit Pipe, using `DBMS_PIPE.CREATE_PIPE`.
- Create the message to cache in the Singleton Pipe.
- Send Message to Singleton Pipe, optionally specifying a `shelflife` for the implicit message.

Exceptions

Table 6-7 RECEIVE_MESSAGE Function Exceptions

Exception	Description
Null pipe name	Permission error. Insufficient privilege to remove the record from the pipe. The pipe is owned by someone else.

Example

```
DECLARE
  l_status INTEGER;
BEGIN
  l_status := DBMS_PIPE.receive_message(pipename   => 'MY_PIPE1',
                                        timeout    => 1,
                                        cache_func =>
'MY_USER.MY_CACHE_FUNC');
```

```
END;
/
```

SEND_MESSAGE Function

This function sends a message on the named pipe.

The message is contained in the local message buffer, which was filled with calls to `PACK_MESSAGE`. You can create a pipe explicitly using `CREATE_PIPE`, otherwise, it is created implicitly.

To create an implicit Singleton pipe, set the `singleton` parameter to `TRUE`. The following arguments are applicable to Singleton Pipes:

- `singleton`: Indicates that the pipe should be created as a Singleton Pipe (default value: `FALSE`).
- `shelflife`: Optionally specify a shelflife expiration of cached message in the Singleton Pipe. It can be used for implicit invalidation of message in Singleton Pipe.
  This argument is applicable for implicit as well as explicit Singleton pipes. A `shelflife` value specified in SEND_MESSAGE Function overwrites the `shelflife` specified for the Explicit Singleton Pipe in CREATE_PIPE Function and will be the default for any new messages cached in the Singleton Pipe.

Syntax

```
DBMS_PIPE.SEND_MESSAGE (
    pipename     IN VARCHAR2,
    timeout      IN INTEGER DEFAULT MAXWAIT,
    maxpipesize  IN INTEGER DEFAULT 65536,
    singleton    IN BOOLEAN DEFAULT FALSE,
    shelflife    IN INTEGER DEFAULT 0)
RETURN INTEGER;
```

Parameters

Table 6-8 SEND_MESSAGE Function Parameters

Parameter	Description
`pipename`	Name of the pipe on which you want to place the message. If you are using an explicit pipe, then this is the name that you specified when you called `CREATE_PIPE`. Caution: Do not use pipe names beginning with '`ORA$`'. These names are reserved for use by procedures provided by Oracle. Pipename should not be longer than 128 bytes, and is case-insensitive. At this time, the name cannot contain Globalization Support characters.
`timeout`	Time to wait while attempting to place a message on a pipe, in seconds. The default value is the constant `MAXWAIT`, which is defined as 86400000 (1000 days).

Table 6-8 (Cont.) SEND_MESSAGE Function Parameters

Parameter	Description
maxpipesize	Maximum size allowed for the pipe, in bytes. The total size of all the messages on the pipe cannot exceed this amount. The message is blocked if it exceeds this maximum. The default is 65536 bytes. The maxpipesize for a pipe becomes a part of the characteristics of the pipe and persists for the life of the pipe. Callers of SEND_MESSAGE with larger values cause the maxpipesize to be increased. Callers with a smaller value simply use the existing, larger value. Specifying maxpipesize as part of the SEND_MESSAGE procedure eliminates the need for a separate call to open the pipe. If you created the pipe explicitly, then you can use the optional maxpipesize parameter to override the creation pipe size specifications. The default maxpipesize of 65536 is applicable for all Pipes.
singleton	Use TRUE to create a Singleton Pipe. Default value: FALSE
shelflife	Expiration time in seconds of a message cached in Singleton Pipe. After the specified shelflife time is exceeded, the message is no longer accessible from the Pipe. The parameter shelflife is only applicable to a Singleton Pipe. Default value is 0, implying that the message never expires.

Return Values

Table 6-9 SEND_MESSAGE Function Return Values

Return	Description
0	Success. If the pipe already exists and the user attempting to create it is authorized to use it, then Oracle returns 0, indicating success, and any data already in the pipe remains. If a user connected as SYSDBS/SYSOPER re-creates a pipe, then Oracle returns status 0, but the ownership of the pipe remains unchanged.
1	Timed out. This procedure can timeout either because it cannot get a lock on the pipe, or because the pipe remains too full to be used. If the pipe was implicitly-created and is empty, then it is removed.
3	An interrupt occurred. If the pipe was implicitly created and is empty, then it is removed.
6	Failed to convert existing pipe to singleton pipe. • Implicit pipe with more than one existing message cannot be converted to a Singleton Pipe. • For an explicit pipe that is not Singleton, DBMS_PIPE.SEND_MESSAGE cannot send a message with singleton argument set to TRUE.
7	A non-zero value is given for the shelflife parameter and the pipe is not a singleton pipe.

Table 6-9 (Cont.) SEND_MESSAGE Function Return Values

Return	Description
ORA-23322	Insufficient privileges. If a pipe with the same name exists and was created by a different user, then Oracle signals error ORA-23322, indicating the naming conflict.

Exceptions

Table 6-10 SEND_MESSAGE Function Exception

Exception	Description
Null pipe name	Permission error. Insufficient privilege to write to the pipe. The pipe is private and owned by someone else.

DBMS_PIPE Overview for Persistent Messaging Pipes

Pipe functionality has several potential applications: external service interface, debugging, independent transactions, and alerts.

On Autonomous Database the DBMS_PIPE package has extended functionality to support persistent messaging pipes.

Persistent Messages in DBMS_PIPE:

- Support the ability to send and retrieve very large messages.
- Support a large number of pipe messages.
- Support sharing of messages within a single database, across multiple databases, and across databases in different regions.
- Support multiple pipes using the same Cloud Object Store location URI.

  Persistent messaging functionality allows two or more database sessions to communicate with messages that are stored in Cloud Object Store. Using this functionality messages in a pipe can be made available to only the current database or they can be made available to multiple databases in the same region or across different regions.

  A Persistent Messaging Pipe can be any one of the supported DBMS_PIPE types:

  - **Implicit Pipe:** Automatically created when a message is sent with an unknown pipe name using the DBMS_PIPE.SEND_MESSAGE function.
  - **Explicit Pipe:** Created using the DBMS_PIPE.CREATE_PIPE function with a user specified pipe name.
  - **Public Pipe:** Accessible by any user with EXECUTE permission on DBMS_PIPE package.
  - **Private Pipe:** Accessible by sessions with the same user as the pipe creator.

> **Note:**
>
> When send and receive messages across different databases using persistent messages, Oracle recommends you call `DBMS_PIPE.CREATE_PIPE` before sending or receiving messages. Creating an explicit pipe with `DBMS_PIPE.CREATE_PIPE` ensures that a pipe is created with the access permissions you want, either public or private (by setting the `PRIVATE` parameter to `FALSE` or using the default value `TRUE`).

DBMS_PIPE Limitation

The `DBMS_PIPE` package does not support sending messages between databases that use different character sets. For example, if you have one Autonomous Database instances that uses AL32UTF8 and another instance that uses WE8MSWIN1252, you cannot send messages with `DBMS_PIPE` between these two databases. In this case, you the system will raise the error `ORA-12704` if you attempt to send messages with `DBMS_PIPE` between these two databases.

See Choose a Character Set for Autonomous Database for more information.

Summary of DBMS_PIPE Subprograms for Persistent Messaging

This table lists the `DBMS_PIPE` subprograms and briefly describes them.

Table 6-11 DBMS_PIPE Package Subprograms

Subprogram	Description
CREATE_PIPE Function	Creates a pipe (necessary for private pipes).
GET_CREDENTIAL_NAME Function	Returns the global `credential_name` variable value.
GET_LOCATION_URI Function	Returns the global `location_uri` variable value that is used as a default location URI for use when a message is stored in Cloud Object Store.
NEXT_ITEM_TYPE Function	Returns datatype of next item in buffer.
PACK_MESSAGE Procedures	Builds message in local buffer.
RECEIVE_MESSAGE Function	Copies message from named pipe into local buffer.
RESET_BUFFER Procedure	Purges contents of local buffer.
REMOVE_PIPE Function	Removes the named pipe.
SEND_MESSAGE Function	Sends message on a named pipe: This implicitly creates a public pipe if the named pipe does not exist.
SET_CREDENTIAL_NAME Procedure	Sets the `credential_name` variable that is used as a default credential for messages that are stored in Cloud Object Store.
SET_LOCATION_URI Procedure	Sets the global `location_uri` variable that is used as a default location URI for messages that are stored in Cloud Object Store.
UNIQUE_SESSION_NAME Function	Returns unique session name.
UNPACK_MESSAGE Procedures	Accesses next item in buffer.

- CREATE_PIPE Function
  This function explicitly creates a public or private pipe. If the `private` flag is `TRUE`, then the pipe creator is assigned as the owner of the private pipe.

- GET_CREDENTIAL_NAME Function
  This function returns the global `credential_name` variable value for use when messages are stored to Cloud Object Store.
- GET_LOCATION_URI Function
  This function returns the global `location_uri` variable value that can be used as a default location URI when pipe messages are stored to Cloud Object Store.
- RECEIVE_MESSAGE Function
  This function copies the message into the local message buffer.
- SEND_MESSAGE Function
  This function sends a message on the named pipe.
- SET_CREDENTIAL_NAME Procedure
  This procedure sets the `credential_name` variable that is used as a default credential when pipe messages are stored in Cloud Object Store.
- SET_LOCATION_URI Procedure
  This procedure sets the global `location_uri` variable.

CREATE_PIPE Function

This function explicitly creates a public or private pipe. If the `private` flag is `TRUE`, then the pipe creator is assigned as the owner of the private pipe.

Explicitly-created pipes can only be removed by calling `REMOVE_PIPE`, or by shutting down the instance.

Syntax

```
DBMS_PIPE.CREATE_PIPE (
   pipename     IN VARCHAR2,
   maxpipesize  IN INTEGER DEFAULT 66536,
   private      IN BOOLEAN DEFAULT TRUE)
RETURN INTEGER;
```

Parameters

Table 6-12 CREATE_PIPE Function Parameters

Parameter	Description
`pipename`	Name of the pipe you are creating. You must use this name when you call `SEND_MESSAGE` and `RECEIVE_MESSAGE`. This name must be unique across the instance. Caution: Do not use pipe names beginning with `ORA$`. These are reserved for use by procedures provided by Oracle. Pipename should not be longer than 128 bytes, and is case insensitive. At this time, the name cannot contain Globalization Support characters.

Table 6-12 (Cont.) CREATE_PIPE Function Parameters

Parameter	Description
`maxpipesize`	The maximum size allowed for the pipe, in bytes. The total size of all of the messages on the pipe cannot exceed this amount. The message is blocked if it exceeds this maximum. The default `maxpipesize` is 66536 bytes. The `maxpipesize` for a pipe becomes a part of the characteristics of the pipe and persists for the life of the pipe. Callers of `SEND_MESSAGE` with larger values cause the `maxpipesize` to be increased. Callers with a smaller value use the existing, larger value. The default `maxpipesize` of 65536 is applicable for all Pipes.
`private`	Uses the default, `TRUE`, to create a private pipe. Public pipes can be implicitly created when you call `SEND_MESSAGE`.

Return Values

Table 6-13 CREATE_PIPE Function Return Values

Return	Description
`0`	Successful. If the pipe already exists and the user attempting to create it is authorized to use it, then Oracle returns 0, indicating success, and any data already in the pipe remains.
`ORA-23322`	Failure due to naming conflict. If a pipe with the same name exists and was created by a different user, then Oracle signals error `ORA-23322`, indicating the naming conflict.

Exceptions

Table 6-14 CREATE_PIPE Function Exception

Exception	Description
`Null pipe name`	Permission error: Pipe with the same name already exists, and you are not allowed to use it.

Example

Create an explicit private named `MY_PIPE1`

```
DECLARE
  l_status INTEGER;
BEGIN
  l_status := DBMS_PIPE.create_pipe(
      pipename  => 'MY_PIPE1',
      private   => TRUE);
END;
/
```

GET_CREDENTIAL_NAME Function

This function returns the global `credential_name` variable value for use when messages are stored to Cloud Object Store.

Syntax

```
DBMS_PIPE.GET_CREDENTIAL_NAME
         RETURN VARCHAR2;
```

Return Values

Return Value	Description
`credential_name`	The name of the credential to access the Cloud Object Storage.

Example

```
DECLARE
  credential_name     VARCHAR2(400)
BEGIN
  credential_name := DBMS_PIPE.GET_CREDENTIAL_NAME;
END;
/
```

GET_LOCATION_URI Function

This function returns the global `location_uri` variable value that can be used as a default location URI when pipe messages are stored to Cloud Object Store.

Syntax

```
DBMS_PIPE.GET_LOCATION_URI
        RETURN VARCHAR2;
```

Return Value

Return Value	Description
`location_uri`	The object URI.

Example

```
DECLARE
  location_uri     VARCHAR2(400)
BEGIN
  location_uri := DBMS_PIPE.GET_LOCATION_URI;
END;
/
```

RECEIVE_MESSAGE Function

This function copies the message into the local message buffer.

Syntax

```
DBMS_PIPE.RECEIVE_MESSAGE (
   pipename          IN VARCHAR2,
   timeout           IN INTEGER  DEFAULT maxwait,
   credential_name   IN VARCHAR2 DEFAULT null,
   location_uri      IN VARCHAR2)
RETURN INTEGER;
```

Parameters

Table 6-15 RECEIVE_MESSAGE Function Parameters

Parameter	Description
`pipename`	Name of the pipe on which you want to receive a message. Names beginning with `ORA$` are reserved for use by Oracle.
`timeout`	Time to wait for a message, in seconds. A timeout of 0 lets you read without blocking. The timeout does not include the time spent running the cache function specified with the `cache_func` parameter. Default value: is the constant `MAXWAIT`, which is defined as 86400000 (1000 days).
`credential_name`	The credential name for the cloud store used to store messages. The `credential_name` is a package argument that is by default initialized as `NULL`. You can set this value before calling `DBMS_PIPE.RECEIVE_MESSAGE`. The passed parameter value takes precedence over the global variable's value. The credential object must have `EXECUTE` and `READ/WRITE` privileges by the user running `DBMS_PIPE.RECEIVE_MESSAGE`. The `credential_name` value can be an OCI resource principal, Azure service principal, Amazon Resource Name (ARN), or a Google service account. See Configure Policies and Roles to Access Resources for more information on resource principal based authentication.
`location_uri`	The location URI for the cloud store used to store messages. The `location_uri` is a global variable that by default is initialized as `NULL`. You can set this value before calling `DBMS_PIPE.RECEIVE_MESSAGE`. The passed parameter value takes precedence over the global variable's value.

Return Values

Table 6-16 RECEIVE_MESSAGE Function Return Values

Return	Description
`0`	Success

Table 6-16 (Cont.) RECEIVE_MESSAGE Function Return Values

Return	Description
1	Timed out. If the pipe was implicitly-created and is empty, then it is removed.
2	Record in the pipe is too large for the buffer.
3	An interrupt occurred.
ORA-23322	User has insufficient privileges to read from the pipe.

Usage Notes

- To receive a message from a pipe, first call `RECEIVE_MESSAGE`. When you receive a message, it is removed from the pipe; hence, a message can only be received once. For implicitly-created pipes, the pipe is removed after the last record is removed from the pipe.
- If the pipe that you specify when you call `RECEIVE_MESSAGE` does not already exist, then Oracle implicitly creates the pipe and waits to receive the message. If the message does not arrive within a designated timeout interval, then the call returns and the pipe is removed.
- After receiving the message, you must make one or more calls to `UNPACK_MESSAGE` to access the individual items in the message. The `UNPACK_MESSAGE` procedure is overloaded to unpack items of type `DATE`, `NUMBER`, `VARCHAR2`, and there are two additional procedures to unpack `RAW` and `ROWID` items. If you do not know the type of data that you are attempting to unpack, then call `NEXT_ITEM_TYPE` to determine the type of the next item in the buffer.
- Persistent messages are guaranteed to either be written or read by exactly one process. This prevents message content inconsistency due to concurrent writes and reads. Using a persistent messaging pipe, `DBMS_PIPE` allows only one operation, sending a message or a receiving message to be active at a given time. However, if an operation is not possible due to an ongoing operation, the process retries periodically until the `timeout` value is reached.
- If you use Oracle Cloud Infrastructure Object Storage to store messages, you can use Oracle Cloud Infrastructure Native URIs or Swift URIs. However, the location URI and the credential must match in type as follows:
  - If you use a native URI format to access Oracle Cloud Infrastructure Object Storage, you must use Native Oracle Cloud Infrastructure Signing Keys authentication in the credential object.
  - If you use Swift URI format to access Oracle Cloud Infrastructure Object Storage, you must use an auth token authentication in the credential object.

Exceptions

Table 6-17 RECEIVE_MESSAGE Function Exceptions

Exception	Description
`Null pipe name`	Permission error. Insufficient privilege to remove the record from the pipe. The pipe is owned by someone else.

SEND_MESSAGE Function

This function sends a message on the named pipe.

The message is contained in the local message buffer, which was filled with calls to `PACK_MESSAGE`. You can create a pipe explicitly using `CREATE_PIPE`, otherwise, it is created implicitly.

Syntax

```
DBMS_PIPE.SEND_MESSAGE (
    pipename          IN VARCHAR2,
    timeout           IN INTEGER DEFAULT MAXWAIT,
    credential_name   IN VARCHAR2 DEFAULT null,
    location_uri      IN VARCHAR2 )
RETURN INTEGER;
```

Parameters

Table 6-18 SEND_MESSAGE Function Parameters

Parameter	Description
`credential_name`	The credential name for the cloud store used to store messages. The `credential_name` is a package argument that is by default initialized as `NULL`. You can set this value before calling `DBMS_PIPE.SEND_MESSAGE`. The passed parameter value takes precedence over the global variable's value. The credential object must have `EXECUTE` and `READ/WRITE` privileges by the user running `DBMS_PIPE.SEND_MESSAGE`. The `credential_name` value can be an OCI resource principal, Azure service principal, Amazon Resource Name (ARN), or a Google service account. See Configure Policies and Roles to Access Resources for more information on resource principal based authentication.
`location_uri`	The location URI for the cloud store used to store messages. The `location_uri` is a global variable that by default is initialized as `NULL`. You can set this value before calling `DBMS_PIPE.SEND_MESSAGE`. The passed parameter value takes precedence over the global variable's value.
`maxpipesize`	Maximum size allowed for the pipe, in bytes. The total size of all the messages on the pipe cannot exceed this amount. The message is blocked if it exceeds this maximum. The default is 65536 bytes. The `maxpipesize` for a pipe becomes a part of the characteristics of the pipe and persists for the life of the pipe. Callers of `SEND_MESSAGE` with larger values cause the `maxpipesize` to be increased. Callers with a smaller value simply use the existing, larger value. Specifying `maxpipesize` as part of the `SEND_MESSAGE` procedure eliminates the need for a separate call to open the pipe. If you created the pipe explicitly, then you can use the optional `maxpipesize` parameter to override the creation pipe size specifications. The default `maxpipesize` of 65536 is applicable for all Pipes.

Table 6-18 (Cont.) SEND_MESSAGE Function Parameters

Parameter	Description
pipename	Name of the pipe on which you want to place the message. If you are using an explicit pipe, then this is the name that you specified when you called CREATE_PIPE. Caution: Do not use pipe names beginning with 'ORA$'. These names are reserved for use by procedures provided by Oracle. Pipename should not be longer than 128 bytes, and is case-insensitive. At this time, the name cannot contain Globalization Support characters.
timeout	Time to wait while attempting to place a message on a pipe, in seconds. The default value is the constant MAXWAIT, which is defined as 86400000 (1000 days).

Return Values

Table 6-19 SEND_MESSAGE Function Return Values

Return	Description
0	Success. If the pipe already exists and the user attempting to create it is authorized to use it, then Oracle returns 0, indicating success, and any data already in the pipe remains. If a user connected as SYSDBS/SYSOPER re-creates a pipe, then Oracle returns status 0, but the ownership of the pipe remains unchanged.
1	Timed out. This procedure can timeout either because it cannot get a lock on the pipe, or because the pipe remains too full to be used. If the pipe was implicitly-created and is empty, then it is removed.
3	An interrupt occurred. If the pipe was implicitly created and is empty, then it is removed.
ORA-23322	Insufficient privileges. If a pipe with the same name exists and was created by a different user, then Oracle signals error ORA-23322, indicating the naming conflict.

Usage Notes

- Persistent messages are guaranteed to either be written or read by exactly one process. This prevents message content inconsistency due to concurrent writes and reads. Using a persistent messaging pipe, DBMS_PIPE allows only one operation, sending a message or a receiving message to be active at a given time. However, if an operation is not possible due to an ongoing operation, the process retries periodically until the timeout value is reached.
- If you use Oracle Cloud Infrastructure Object Storage to store messages, you can use Oracle Cloud Infrastructure Native URIs or Swift URIs. However, the location URI and the credential must match in type as follows:
  - If you use a native URI format to access Oracle Cloud Infrastructure Object Storage, you must use Native Oracle Cloud Infrastructure Signing Keys authentication in the credential object.

- If you use Swift URI format to access Oracle Cloud Infrastructure Object Storage, you must use an auth token authentication in the credential object.

Exceptions

Table 6-20 SEND_MESSAGE Function Exception

Exception	Description
`Null pipe name`	Permission error. Insufficient privilege to write to the pipe. The pipe is private and owned by someone else.

SET_CREDENTIAL_NAME Procedure

This procedure sets the `credential_name` variable that is used as a default credential when pipe messages are stored in Cloud Object Store.

Syntax

```
DBMS_PIPE.SET_CREDENTIAL_NAME (
   credential_name    IN VARCHAR2 );
```

Parameters

Parameter	Description
`credential_name`	The name of the credential to access the Cloud Object Storage. The `credential_name` value can be an OCI resource principal, Azure service principal, Amazon Resource Name (ARN), or a Google service account. See Configure Policies and Roles to Access Resources for more information on resource principal based authentication.

Usage Note

If you use Oracle Cloud Infrastructure Object Storage to store messages, you can use Oracle Cloud Infrastructure Native URIs or Swift URIs. However, the location URI and the credential must match in type as follows:

- If you use a native URI format to access Oracle Cloud Infrastructure Object Storage, you must use Native Oracle Cloud Infrastructure Signing Keys authentication in the credential object.
- If you use Swift URI format to access Oracle Cloud Infrastructure Object Storage, you must use an auth token authentication in the credential object.

Example

```
BEGIN
     DBMS_PIPE.SET_CREDENTIAL_NAME(
       credential_name =>  'my_cred1');
END;
/
```

SET_LOCATION_URI Procedure

This procedure sets the global `location_uri` variable.

Syntax

```
DBMS_PIPE.SET_LOCATION_URI (
   location_uri    IN VARCHAR2 );
```

Parameter

Parameter	Description
`location_uri`	Object or file URI. The format of the URI depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD URI Formats.

Usage Note

If you use Oracle Cloud Infrastructure Object Storage to store messages, you can use Oracle Cloud Infrastructure Native URIs or Swift URIs. However, the location URI and the credential must match in type as follows:

- If you use a native URI format to access Oracle Cloud Infrastructure Object Storage, you must use Native Oracle Cloud Infrastructure Signing Keys authentication in the credential object.
- If you use Swift URI format to access Oracle Cloud Infrastructure Object Storage, you must use an auth token authentication in the credential object.

Example

```
BEGIN
  DBMS_PIPE.GET_LOCATION_URI(
      location_uri  => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname1/');
END;
/
```

DBMS_PROXY_SQL

Describes the procedures to enable and disable query offload for an elastic pool leader.

- DISABLE_READ_ONLY_OFFLOAD Procedure
  This procedure disables query offload for an Autonomous Database elastic pool leader or for an elastic pool member.
- ENABLE_READ_ONLY_OFFLOAD Procedure
  This procedure enables query offload for an Autonomous Database elastic pool leader or for an elastic pool member.

DISABLE_READ_ONLY_OFFLOAD Procedure

This procedure disables query offload for an Autonomous Database elastic pool leader or for an elastic pool member.

Syntax

```
DBMS_PROXY_SQL.DISABLE_READ_ONLY_OFFLOAD;
```

Example

```
EXEC DBMS_PROXY_SQL.DISABLE_READ_ONLY_OFFLOAD;
```

ENABLE_READ_ONLY_OFFLOAD Procedure

This procedure enables query offload for an Autonomous Database elastic pool leader or for an elastic pool member.

Syntax

```
DBMS_PROXY_SQL.ENABLE_READ_ONLY_OFFLOAD(
      module_name       IN CLOB,
      action_name       IN CLOB);
```

Parameters

Parameter	Description
`module_name`	Specifies a list of module names as a comma-separated list. The list specifies the modules where queries are considered for offload (where a session's module name matches a value in the list). The default value for this parameter is `NULL`, which means that a session's module name can be any value and the session is considered for offloading.
`action_name`	Specifies a list of action names as a comma-separated list. The list specifies the action names where queries are considered for offloading (where a session's action name matches a value in the list). The default value for this parameter is `NULL`, which means that a session's action name can be any value and the session is considered for offloading

Usage Notes

- If both `module_name` and `action_name` are specified, a session's module name must match a value in the list of module names and it's action name must match a value in the list of action names for the session to be considered for offload.
- When query offload is enabled for a session you can find the name of the Refreshable Clone to which queries are offloaded. For example:

  ```
  SELECT sys_context('userenv', 'con_name') from dual;
  ```

  If queries are not being offloaded to a Refreshable Clone, this query shows the name of the elastic pool leader (or the name of the elastic pool member).
- The views `DBA_PROXY_SQL_ACTIONS` and `DBA_PROXY_SQL_MODULES` display the list of modules or actions that are configured for query offload. You must query these views from a session that is not enabled for query offload.

  See DBA_PROXY_SQL Views for more information.

- When you offload queries from a list of sessions that you specify by module or action name, the module name and action name set with `DBMS_APPLICATION_INFO` are case sensitive. `DBMS_PROXY_SQL.ENABLE_READ_ONLY_OFFLOAD` adheres to this case sensitive behavior. For example, when a session has a lowercase module name, or mixed case module name, the case must match in the parameter values when you enable query offload with `DBMS_PROXY_SQL.ENABLE_READ_ONLY_OFFLOAD` and you include the `module_name` or `action_name` parameters.

  You can check the module name and action name for the current session using `DBMS_APPLICATION_INFO.READ_MODULE`:

```
set serveroutput on;
declare
  l_mod varchar2(50);
  l_act varchar2(50);
BEGIN
DBMS_APPLICATION_INFO.READ_MODULE(l_mod, l_act);
  DBMS_OUTPUT.PUT_LINE('l_mod: ' || l_mod);
  DBMS_OUTPUT.PUT_LINE('l_act: ' || l_act);
END;
/
```

Examples

```
EXEC DBMS_PROXY_SQL.ENABLE_READ_ONLY_OFFLOAD;
```

```
DECLARE
   mod_values clob := to_clob('["mod1", "mod2"]');
   act_values clob := to_clob('["act1", "act2"]');
BEGIN
   DBMS_PROXY_SQL.ENABLE_READ_ONLY_OFFLOAD(
      module_name => mod_values,
      action_name => act_values);
END;
/
```

DBMS_SHARE Package

The `DBMS_SHARE` subprograms and views are used to share local data with external systems.

- Summary of Share Producer Subprograms
  This table lists the `DBMS_SHARE` package procedures and functions used to produce shares for recipients.
- Summary of Share Consumer Subprograms
  This table lists the `DBMS_SHARE` package procedures and functions used to consume shares.
- Summary of Share Producer Views
  This table lists the share producer views for the `DBMS_SHARE` package.
- Summary of Share Consumer Views
  This table lists the `DBMS_SHARE` package views.

- DBMS_SHARE Constants
  These constants are used by the `DBMS_SHARE` package.

Summary of Share Producer Subprograms

This table lists the `DBMS_SHARE` package procedures and functions used to produce shares for recipients.

Subprogram	Description
ADD_TO_SHARE Procedure	Add a table or view to a share.
ASSERT_SHAREABLE_OBJECT Procedure	Return without error, if the object exists and can be shared.
ASSERT_SHARING_ID Procedure	Run basic validation checks against a sharing id and return one in canonical form.
CAN_CREATE_SHARE Function	This function checks to see if the current schema can create share recipients.
CAN_CREATE_SHARE_RECIPIENT Function	This function checks to see if the current schema can create share recipients.
CLEAR_RECIPIENT_EVENTS Procedure	Clear events from the share recipient event log.
CLEAR_SHARE_EVENTS Procedure	Clear events from the share event log.
CREATE_BEARER_TOKEN_CREDENTIAL Procedure	Create a credential suitable for use with delta share providers.
CREATE_CLOUD_STORAGE_LINK Procedure	Create a named cloud storage URI link.
CREATE_OR_REPLACE_CLOUD_STORAGE_LINK Procedure	Create or replace a named cloud storage URI.
CREATE_OR_REPLACE_SHARE_RECIPIENT Procedure	Create or replace a share recipient.
CREATE_SHARE Procedure	Create a named share object.
CREATE_SHARE_RECIPIENT Procedure	Create a new share recipient.
DROP_CLOUD_STORAGE_LINK Procedure	Drop a cloud storage link.
DROP_RECIPIENT Procedure	Drop a recipient.
DROP_SHARE Procedure	Drop a share and all of its contents.
DROP_SHARE_LINK_VIEW Procedure	Drop a view that was created by the `CREATE_SHARE_LINK_VIEW` procedure.
DROP_SHARE_VERSION Procedure	Drop a single share version.
DROP_SHARE_VERSIONS Procedure	Drop a range of share versions.
DROP_UNUSED_SHARE_VERSIONS Procedure	Drop any share version that is not currently in use.
ENABLE_SCHEMA Procedure	Enable or disable a schema for sharing.
GET_ACTIVATION_LINK Function	Generate the link that gets put into emails to the authorized recipient.

Subprogram	Description
GET_PUBLISHED_IDENTITY Function	Get data about the current user that was set by SET_PUBLISHED_IDENTITY.
GET_RECIPIENT_PROPERTY Function	Return the value of a property for a recipient.
GET_SHARE_PROPERTY Function	Get the property value of an existing share.
GET_SHARE_TABLE_PROPERTY Function	Get the property value of an existing share table.
GRANT_TO_RECIPIENT Procedure	Grant access on a share to a specific recipient.
POPULATE_SHARE_PROFILE Procedure	Generate a delta profile for a recipient.
PUBLISH_SHARE Procedure	Publish a share and return immediately.
PUBLISH_SHARE_WAIT Procedure	Publish a share and wait until the background job is complete.
PURGE_DETACHED_FILES Procedure	Delete or forget parquet files that have become detached from their shares.
REMOVE_FROM_SHARE Procedure	Remove a table or view from a share.
RENAME_RECIPIENT Procedure	Rename a recipient.
RENAME_SHARE Procedure	Rename a share.
RENAME_SHARE_LINK Procedure	Rename a registered share link.
RENAME_SHARE_SCHEMA Procedure	Rename a share schema.
RENAME_SHARE_TABLE Procedure	Rename a share table.
REVOKE_FROM_RECIPIENT Procedure	Revoke access on a share from a specific recipient.
SET_CURRENT_SHARE_VERSION Procedure	Change the current version of a share.
SET_PUBLISHED_IDENTITY Procedure	Set data about the current user that will be supplied to recipients of published ORACLE shares.
SET_RECIPIENT_LOG_LEVEL Procedure	Change the log level for an existing share recipient.
SET_SHARE_LOG_LEVEL Procedure	Change the log level for an existing share.
SET_STORAGE_CREDENTIAL Procedure	Set the access credential name for the given storage.
STOP_JOB Procedure	Stop a running share job.
UNPUBLISH_SHARE Procedure	Unpublish a share.
UPDATE_DEFAULT_RECIPIENT_PROPERTY Procedure	Update the default recipient property values.
UPDATE_DEFAULT_SHARE_PROPERTY Procedure	Update the default share property values.
UPDATE_RECIPIENT_PROPERTY Procedure	Update a property of an existing recipient.

Subprogram	Description
UPDATE_SHARE_JOB_PROPERTY Procedure	Modify properties of a running share job.
UPDATE_SHARE_PROPERTY Procedure	Update a property of an existing share.
UPDATE_SHARE_TABLE_PROPERTY Procedure	Update the property value of an existing share table.
VALIDATE_CREDENTIAL Function	Validate a credential name, converting it to canonical form first if required.
VALIDATE_SHARE_STORAGE Procedure	Check to see if the given storage is suitable for versioned shares.
WAIT_FOR_JOB Procedure	This procedure waits until the specified share job is complete.

- ADD_TO_SHARE Procedure
  Add a table or view to a share. The object becomes visible to any external user who has been granted access to the share.
- ASSERT_SHAREABLE_OBJECT Procedure
  Return without error, if the object exists and can be shared.
- ASSERT_SHARING_ID Procedure
  Run basic validation checks against a sharing id and return one in canonical form. An exception is raised if the id is obviously invalid.
- CAN_CREATE_SHARE Function
  This function checks to see if the current schema can create share recipients. If shares can be created, a 1 is returned and 0 otherwise.
- CAN_CREATE_SHARE_RECIPIENT Function
  This function checks to see if the current schema can create share recipients. If shares can be created a 1 is returned, and 0 otherwise.
- CLEAR_RECIPIENT_EVENTS Procedure
  Clear events from the share recipient event log.
- CLEAR_SHARE_EVENTS Procedure
  Clear events from the share event log.
- CREATE_BEARER_TOKEN_CREDENTIAL Procedure
  Create a credential suitable for use with delta share providers.
- CREATE_CLOUD_STORAGE_LINK Procedure
  Create a named cloud storage URI link. A cloud storage link is a named association between an OCI bucket URI, and a local credential name.
- CREATE_OR_REPLACE_CLOUD_STORAGE_LINK Procedure
  Create or replace a named cloud storage URI. A cloud storage link is a named association between an OCI bucket URI, and a local credential name.
- CREATE_OR_REPLACE_SHARE_RECIPIENT Procedure
  Create or replace a share recipient. You must provide at least an email address or sharing id.
- CREATE_SHARE Procedure
  Create a named share object.
- CREATE_SHARE_RECIPIENT Procedure
  Create a new share recipient.
- DROP_CLOUD_STORAGE_LINK Procedure
  Drop a cloud storage link.

- DROP_RECIPIENT Procedure
  Drop a recipient. All access to the recipient will be revoked.
- DROP_SHARE Procedure
  Drop a share and all of its contents. Future access to the share by consumers will end.
- DROP_SHARE_LINK_VIEW Procedure
  Drop a view that was created by the `CREATE_SHARE_LINK_VIEW` procedure.
- DROP_SHARE_VERSION Procedure
  Drop a single share version. Note that you cannot drop the current version.
- DROP_SHARE_VERSIONS Procedure
  Drop a range of share versions. Note that you cannot drop the current version using this procedure.
- DROP_UNUSED_SHARE_VERSIONS Procedure
  Drop any share version that is not currently in use.
- ENABLE_SCHEMA Procedure
  Enable or disable a schema for sharing. This procedure must be run by the ADMIN user.
- GET_ACTIVATION_LINK Function
  Generate the link that gets put into emails to the authorized recipient. This activation link leads to the download page, where the recipient clicks a button to get the delta profile.
- GET_PUBLISHED_IDENTITY Function
  Get data about the current user that was set by `SET_PUBLISHED_IDENTITY`.
- GET_RECIPIENT_PROPERTY Function
  Return the value of a property for a recipient.
- GET_SHARE_PROPERTY Function
  Get the property value of an existing share.
- GET_SHARE_TABLE_PROPERTY Function
  Get the property value of an existing share table.
- GRANT_TO_RECIPIENT Procedure
  Grant access on a share to a specific recipient. The share and recipient must both belong to the same schema.
- POPULATE_SHARE_PROFILE Procedure
  Generate a delta profile for a recipient. You could print this to the screen or upload it somewhere. For example, to an object bucket using `DBMS_CLOUD.EXPORT_DATA`.
- PUBLISH_SHARE Procedure
  Publish a share and return immediately.
- PUBLISH_SHARE_WAIT Procedure
  Publish a share and wait until the background job is complete. The publication continues even if the call is interrupted.
- PURGE_DETACHED_FILES Procedure
  Delete or forget parquet files that have become detached from their shares.
- REMOVE_FROM_SHARE Procedure
  Remove a table or view from a share.
- RENAME_RECIPIENT Procedure
  Rename a recipient. This procedure only changes the local name of the recipient. The external definition of the recipient, for example the name of the OAUTH user or sharing id, is not changed.

- RENAME_SHARE Procedure
  Rename a share. Care should be take with this procedure since the change effects any existing consumers whose access is based on the previous name. Consumers are not notified directly or updated.
- RENAME_SHARE_LINK Procedure
  Rename a registered share link.
- RENAME_SHARE_SCHEMA Procedure
  Rename a share schema.
- RENAME_SHARE_TABLE Procedure
  Rename a share table.
- REVOKE_FROM_RECIPIENT Procedure
  Revoke access on a share from a specific recipient.
- SET_CURRENT_SHARE_VERSION Procedure
  Change the current version of a share.
- SET_PUBLISHED_IDENTITY Procedure
  Set data about the current user that will be supplied to recipients of published ORACLE shares.
- SET_RECIPIENT_LOG_LEVEL Procedure
  Change the log level for an existing share recipient.
- SET_SHARE_LOG_LEVEL Procedure
  Change the log level for an existing share.
- SET_STORAGE_CREDENTIAL Procedure
  Set the credential name used by the current user when it attempts to access the given storage.
- STOP_JOB Procedure
  Attempt to stop a running share job. The procedure should return quickly, but it may take some time for the associated job to stop.
- UNPUBLISH_SHARE Procedure
  Unpublish a share.
- UPDATE_DEFAULT_RECIPIENT_PROPERTY Procedure
  Update the default recipient property values. This procedure requires the user to have admin privileges.
- UPDATE_DEFAULT_SHARE_PROPERTY Procedure
  Update the default share property values.
- UPDATE_RECIPIENT_PROPERTY Procedure
  Update a property of an existing recipient.
- UPDATE_SHARE_JOB_PROPERTY Procedure
  Modify properties of a running share job. The procedure should return quickly, but it may take some time for the changes to take effect.
- UPDATE_SHARE_PROPERTY Procedure
  Update a property of an existing share.
- UPDATE_SHARE_TABLE_PROPERTY Procedure
  Update the property value of an existing share table.
- VALIDATE_CREDENTIAL Function
  Validate a credential name, converting it to canonical form first if required.
- VALIDATE_SHARE_STORAGE Procedure
  Check to see if the given storage is suitable for versioned shares.

- WAIT_FOR_JOB Procedure
  This procedure waits until the specified share job is complete.

ADD_TO_SHARE Procedure

Add a table or view to a share. The object becomes visible to any external user who has been granted access to the share.

Syntax

```
PROCEDURE ADD_TO_SHARE
(
    share_name           IN VARCHAR2,
    table_name           IN VARCHAR2,
    owner                IN VARCHAR2 := NULL,
    share_table_name     IN VARCHAR2 := NULL,
    share_schema_name    IN VARCHAR2 := NULL,
    object_metadata      IN SYS.JSON_OBJECT_T := NULL,
    replace_existing     IN BOOLEAN := FALSE,
    share_owner          IN VARCHAR2 := NULL,
    auto_commit          IN BOOLEAN := FALSE
);
```

Parameters

Parameter	Description
share_name	The name of an existing share to which the object is granted.
table_name	The name of the entity to share (for example, a table or view name).
owner	The owner of the entity to share. The default is to current schema.
share_table_name	The externally visible name of the table. By default, this is the uppercased table_name.
share_schema_name	The externally visible schema where the table will be placed. By default this is the uppercased table owner. The schema is created automatically if it does not already exist.
object_metadata	Optional metadata to associate with the shared entity.
replace_existing	If TRUE, and this share_table_name already exists, the existing table_name is dropped from the share and replaced with this table_name. If FALSE, and this share_table_name already exists, an exception is raised indicating the share table is already used.
share_owner	The owner of the share.
auto_commit	If TRUE, this procedure call commits changes that are not visible externally until the commit takes place. The default value is FALSE, which means that the user must COMMIT after running this call in order to make the change visible.

ASSERT_SHAREABLE_OBJECT Procedure

Return without error, if the object exists and can be shared.

Syntax

```
PROCEDURE ASSERT_SHAREABLE_OBJECT
(
  object_name          IN VARCHAR2,
  object_owner         IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
`object_name`	The name of the object.
`object_owner`	The owner of the object. Defaults to current schema.

ASSERT_SHARING_ID Procedure

Run basic validation checks against a sharing id and return one in canonical form. An exception is raised if the id is obviously invalid.

Syntax

```
PROCEDURE ASSERT_SHARING_ID
(
    sharing_id              IN OUT NOCOPY VARCHAR2,
    out_sharing_id_type     IN OUT NOCOPY VARCHAR2
);
```

Parameters

Parameter	Description
`sharing_id`	The id to check.
`out_sharing_id_type`	The type of the id, if valid. For example, `TENANCY` or `DATABASE`.

CAN_CREATE_SHARE Function

This function checks to see if the current schema can create share recipients. If shares can be created, a 1 is returned and 0 otherwise.

Syntax

```
FUNCTION CAN_CREATE_SHARE
RETURN NUMBER;
```

Example: Before and after enabling the admin schema

```
SQL> select dbms_share.can_create_share from dual;
 CAN_CREATE_SHARE
----------------
               0
SQL> exec dbms_share.enable_schema('admin')
PL/SQL procedure successfully completed.
SQL> select dbms_share.can_create_share from dual;
CAN_CREATE_SHARE
----------------
               1
```

CAN_CREATE_SHARE_RECIPIENT Function

This function checks to see if the current schema can create share recipients. If shares can be created a 1 is returned, and 0 otherwise.

Syntax

```
FUNCTION CAN_CREATE_SHARE_RECIPIENT
RETURN NUMBER;
```

CLEAR_RECIPIENT_EVENTS Procedure

Clear events from the share recipient event log.

Syntax

```
PROCEDURE CLEAR_RECIPIENT_EVENTS
(
   recipient_name       IN VARCHAR2,
   from_time            IN TIMESTAMP WITH TIME ZONE := NULL,
   to_time              IN TIMESTAMP WITH TIME ZONE := NULL,
   recipient_owner      IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
recipient_name	The local name of the share recipient.
from_time	Earliest time for events that should be cleared or NULL.
to_time	Latest time for events that should be cleared or NULL.
recipient_owner	The schema that owns the recipient.

CLEAR_SHARE_EVENTS Procedure

Clear events from the share event log.

Syntax

```
PROCEDURE CLEAR_SHARE_EVENTS
(
    share_name           IN VARCHAR2,
    from_time            IN TIMESTAMP WITH TIME ZONE := NULL,
    to_time              IN TIMESTAMP WITH TIME ZONE := NULL,
    share_owner          IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
share_name	The name of the share.
from_time	Earliest time for events that should be cleared or NULL.
to_time	Latest time for events that should be cleared or NULL.
share_owner	The schema that owns the share.

CREATE_BEARER_TOKEN_CREDENTIAL Procedure

Create a credential suitable for use with delta share providers.

This is similar to the CREATE_CREDENTIALS call, but it takes explicit values instead of a delta sharing profile. See CREATE_CREDENTIALS Procedure and Function for further information.

Syntax

```
PROCEDURE CREATE_BEARER_TOKEN_CREDENTIAL
(
   credential_name      IN VARCHAR2,
   bearer_token         IN VARCHAR2 := NULL,
   token_endpoint       IN VARCHAR2 := NULL,
   client_id            IN VARCHAR2 := NULL,
   client_secret        IN VARCHAR2 := NULL,
   token_refresh_rate   IN PLS_INTEGER := 3600
);
```

Parameters

Parameter	Description
credential_name	The name of the new credential.
bearer_token	The bearer token, if known.
token_endpoint	The endpoint to call to get a new token.
client_id	The username to send to the token_endpoint.
client_secret	The password to send to the token_endpoint.
token_refresh_rate	An optional refresh time, in seconds.

Example: Credential with a fixed bearer token

In its simplest form, this procedure is equivalent to calling `DBMS_CREDENTIAL` with a user name of `'BEARER_TOKEN'` and the bearer token itself as password.

```
SQL> exec dbms_share.create_bearer_token_credential('MY_FIXED_CREDENTIAL',
'FF42322D27D4C2DEE05392644664351E')
PL/SQL procedure successfully completed.
SQL> select username from user_credentials where credential_name =
'MY_FIXED_CREDENTIAL';
USERNAME
--------------------------------------------------------------------------------
----------------------------------------
BEARER_TOKEN
```

Example: Credential with a renewable bearer token

Create a credential that contains a short lived bearer token obtained from a token endpoint. The bearer token will be refreshed once every hour using a second credential, which is populated from the client id and secret.

```
SQL> BEGIN
  2    dbms_share.create_bearer_token_credential(
  3      credential_name=>'MY_RENEWABLE_CREDENTIAL',
  4      token_endpoint=>'https://myserver/ords/share_provider/oauth/token',
  5      client_id=>'VXGQ_44s6qJ-K4WHUNM2yQ..',
  6      client_secret=>'y9ddppgwEmZl7adDHFQndw..');
  7  END;
  8  /
 PL/SQL procedure successfully completed.
 SQL> select credential_name, username from user_credentials where
credential_name LIKE 'MY_RENEWABLE_CREDENTIAL%';
 CREDENTIAL_NAME                             USERNAME
------------------------------------------
-------------------------------------
MY_RENEWABLE_CREDENTIAL                      BEARER_TOKEN
MY_RENEWABLE_CREDENTIAL$TOKEN_REFRESH_CRED   VXGQ_44s6qJ-K4WHUNM2yQ..
```

CREATE_CLOUD_STORAGE_LINK Procedure

Create a named cloud storage URI link. A cloud storage link is a named association between an OCI bucket URI, and a local credential name.

> **Note:**
>
> Use the SET_STORAGE_CREDENTIAL Procedure to add a credential to the storage link.

Syntax

```
PROCEDURE CREATE_CLOUD_STORAGE_LINK
(
```

```
  storage_link_name     IN VARCHAR2,
  uri                   IN VARCHAR2,
  owner                 IN VARCHAR2 := NULL,
  metadata              IN SYS.JSON_OBJECT_T := NULL,
  auto_commit           IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
`storage_link_name`	The name of the cloud storage link. The name of the link should follow standard Oracle naming conventions.
`uri`	The URI of the storage bucket. The URI should be of the form `https://objectstorage.`*`region`*`.oraclecloud.com/n/`*`namespace-string`*`/b/`*`bucket`*`/o`
`owner`	The owner of the storage link. Leave as NULL for the current user.
`metadata`	(Optional) A JSON metadata document containing additional information.
`auto_commit`	The default is TRUE. If TRUE, this transaction is committed. If FALSE, the user must commit the transaction. Changes are not visible until the commit takes place.

Example

In this example, a cloud storage link named `MY_SHARE_STORAGE` is created on the given URL.

```
SQL> BEGIN
  2    dbms_share.create_cloud_storage_link(
  3      'MY_SHARE_STORAGE',
  4      'https://objectstorage.../n/abcdef/b/my_bucket/o' );
  5  END;
  6  /

PL/SQL procedure successfully completed.

SQL> select storage_link_name from user_lineage_cloud_storage_links;

STORAGE_LINK_NAME
--------------------------------------------------------------------------------
----------
MY_SHARE_STORAGE
```

CREATE_OR_REPLACE_CLOUD_STORAGE_LINK Procedure

Create or replace a named cloud storage URI. A cloud storage link is a named association between an OCI bucket URI, and a local credential name.

> **Note:**
>
> Use the SET_STORAGE_CREDENTIAL Procedure to add a credential to the storage link.

Syntax

```
PROCEDURE CREATE_OR_REPLACE_CLOUD_STORAGE_LINK
(
    storage_link_name    IN VARCHAR2,
    uri                  IN VARCHAR2,
    owner                IN VARCHAR2 := NULL,
    metadata             IN SYS.JSON_OBJECT_T := NULL,
    auto_commit          IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
storage_link_name	The name of the cloud storage link. The name of the link should follow standard Oracle naming conventions. See Database Object Names and Qualifiers
uri	The URI of the storage bucket. The URI should be of the form *https://objectstorage.region.oraclecloud.com/n/namespace-string/b/bucket/o*
owner	The owner of the storage link. Defaults to the current schema.
metadata	An optional JSON metadata document containing additional information.
auto_commit	If TRUE, the changes automatically commit after creating the link. The default is TRUE.

CREATE_OR_REPLACE_SHARE_RECIPIENT Procedure

Create or replace a share recipient. You must provide at least an email address or sharing id.

Syntax

```
PROCEDURE CREATE_OR_REPLACE_SHARE_RECIPIENT
(
    recipient_name       IN VARCHAR2,
    description          IN VARCHAR2 := NULL,
    recipient_owner      IN VARCHAR2 := NULL,
    email                IN VARCHAR2 := NULL,
    sharing_id           IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
recipient_name	The local name of the share recipient. Some names are not allowed (e.g. MY_TENANCY).
description	A description of the recipient.
recipient_owner	The schema that owns the recipient.
email	An email that will be registered for the OAUTH user.
sharing_id	The sharing id of the recipient, from GET_SHARING_ID Function.

CREATE_SHARE Procedure

Create a named share object.

Syntax

```
PROCEDURE CREATE_SHARE
(
  share_name           IN VARCHAR2,
  share_type           IN VARCHAR2 := SHARE_TYPE_VERSIONED,
  storage_link_name    IN VARCHAR2 := NULL,
  storage_link_owner   IN VARCHAR2 := NULL,
  description          IN VARCHAR2 := NULL,
  public_description   IN VARCHAR2 := NULL,
  configuration        IN SYS.JSON_OBJECT_T := NULL,
  force_create         IN BOOLEAN  := FALSE,
  share_owner          IN VARCHAR2 := NULL,
  auto_commit          IN BOOLEAN := FALSE,
  log_level            IN PLS_INTEGER := LOG_LEVEL_BASIC,
  run_storage_tests    IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
share_name	The name of the share. This name is uppercased since delta shares are case insensitive. The name follows standard Oracle conventions, so it must be 128 characters or fewer and must be double quoted if it is not a simple identifier. The only difference is that it will be uppercased even if it is double quoted.
share_type	The type of share. For information on constants used for this parameter, see descriptions for **Share Types** in DBMS_SHARE Constants.
storage_link_name	The name of the cloud storage link where the objects are created. The user must have read/write access to this storage and have the ability to create pre-authenticated requests on the storage. The parameter is required for versioned shares, and optional for current shares.
storage_link_owner	The owner of the cloud storage link where objects are created.
description	A local description for the share.
public_description	A public description for the share.
configuration	A configuration document that defines how objects are created.
force_create	Set force_create to TRUE to redefine the share if it exists.
share_owner	The owner of the share.
auto_commit	If TRUE, this procedure call commits changes that are not visible externally until the commit takes place. The default value is FALSE, which means that the user must COMMIT after running this call in order to make the change visible.
log_level	Event logging level. This controls the amount of detail recorded in the ALL_SHARE_EVENTS View. For information on constants used for this parameter, see descriptions for **Log Level** in DBMS_SHARE Constants.

Parameter	Description
run_storage_tests	If this parameter is TRUE then DBMS_SHARE runs tests to verify that the specified share storage link has the correct privileges. If this parameter is FALSE, then the procedure does not run any checks at creation time. This may lead to errors later, during publication or consumption of the share. Oracle recommends that you specify TRUE for this parameter.

CREATE_SHARE_RECIPIENT Procedure

Create a new share recipient.

Syntax

```
PROCEDURE CREATE_SHARE_RECIPIENT
(
    recipient_name       IN VARCHAR2,
    description          IN VARCHAR2 := NULL,
    recipient_owner      IN VARCHAR2 := NULL,
    email                IN VARCHAR2 := NULL,
    sharing_id           IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
recipient_name	The local name of the share recipient. Some names are not allowed, for example: MY_TENANCY.
description	A description of the recipient.
recipient_owner	The schema that owns the recipient.
email	An email that will be registered for the OAUTH user. You must provide at least one of email or sharing id.
sharing_id	The sharing id of the recipient from GET_SHARING_ID Function. You must provide at least one of email or sharing id.

DROP_CLOUD_STORAGE_LINK Procedure

Drop a cloud storage link.

Syntax

```
PROCEDURE DROP_CLOUD_STORAGE_LINK
(
    storage_link_name    IN VARCHAR2,
    owner                IN VARCHAR2 := NULL,
    auto_commit          IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
storage_link_name	The name of the cloud storage link to drop.
owner	The owner of the cloud storage link. Leave as NULL for the current user.
auto_commit	If TRUE, the code automatically commits after dropping the link. The default is TRUE.

DROP_RECIPIENT Procedure

Drop a recipient. All access to the recipient will be revoked.

Syntax

```
PROCEDURE DROP_RECIPIENT
(
    recipient_name       IN VARCHAR2,
    owner                IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
recipient_name	The name of the share recipient.
owner	The schema that defines the recipient.

DROP_SHARE Procedure

Drop a share and all of its contents. Future access to the share by consumers will end.

Syntax

```
PROCEDURE DROP_SHARE
(
  share_name           IN VARCHAR2,
  share_owner          IN VARCHAR2 := NULL,
  destroy_objects      IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
share_name	The name of the share to drop.
share_owner	The owner of the share to drop.
destroy_objects	If TRUE, delete all objects created on behalf of the share. The default is TRUE.

DROP_SHARE_LINK_VIEW Procedure

Drop a view that was created by the `CREATE_SHARE_LINK_VIEW` procedure.

See CREATE_SHARE_LINK_VIEW Procedure for further information.

Syntax

```
PROCEDURE DROP_SHARE_LINK_VIEW
(
    view_name            IN VARCHAR2,
    view_owner           IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
`view_name`	The name of the new view.
`view_owner`	The view owner. The default is the current schema.

DROP_SHARE_VERSION Procedure

Drop a single share version. Note that you cannot drop the current version.

Syntax

```
PROCEDURE DROP_SHARE_VERSION
(
    share_name           IN VARCHAR2,
    share_version        IN NUMBER,
    destroy_objects      IN BOOLEAN := TRUE,
    force_drop           IN BOOLEAN := FALSE,
    share_owner          IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
`share_name`	The name of the share.
`share_version`	The version to drop. You can not drop the current version.
`destroy_objects`	Destroy any associated object in cloud storage, if applicable.
`force_drop`	Drop the share version even if there is an outstanding PAR file on the version.
`share_owner`	The owner of the share.

DROP_SHARE_VERSIONS Procedure

Drop a range of share versions. Note that you cannot drop the current version using this procedure.

Syntax

```
PROCEDURE DROP_SHARE_VERSIONS
(
    share_name           IN VARCHAR2,
    from_share_version   IN NUMBER,
    to_share_version     IN NUMBER,
    destroy_objects      IN BOOLEAN := TRUE,
    force_drop           IN BOOLEAN := FALSE,
    share_owner          IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
share_name	The name of the share.
from_share_version	The lowest version to drop.
to_share_version	The highest version to drop.
destroy_objects	Destroy any associated object in cloud storage, if applicable.
force_drop	Drop the share version even if there is an outstanding PAR file on the version.
share_owner	The owner of the share.

DROP_UNUSED_SHARE_VERSIONS Procedure

Drop any share version that is not currently in use.

Syntax

```
PROCEDURE DROP_UNUSED_SHARE_VERSIONS
(
    share_name           IN VARCHAR2,
    destroy_objects      IN BOOLEAN := TRUE,
    share_owner          IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
share_name	The name of the share.
destroy_objects	Destroy any associated object in cloud storage, if applicable.
share_owner	The owner of the share.

ENABLE_SCHEMA Procedure

Enable or disable a schema for sharing. This procedure must be run by the ADMIN user.

Users can consume delta shares without being enabled with this procedure, but they cannot create or publish shares. Sharing is disabled by default for all schemas, including ADMIN.

Syntax

```
PROCEDURE ENABLE_SCHEMA
(
    schema_name           IN VARCHAR2,
    enabled               IN BOOLEAN := TRUE,
    privileges            IN PLS_INTEGER := NULL
);
```

Parameters

Parameter	Description
schema_name	The schema to enable or disable.
enable	TRUE to enable, FALSE to disable.
privileges	The privileges argument has bitmap values. If you leave the argument as NULL, it defaults to PRIV_CREATE_SHARE + PRIV_CREATE_RECIPIENT + PRIV_CONSUME_ORACLE_SHARE. Bitmap values are as follows: • PRIV_CREATE_SHARE Allow users to create an publish shares in their own schema. PRIV_CREATE_SHARE CONSTANT PLS_INTEGER := 1; • PRIV_CREATE_RECIPIENT Allow users to create share recipients in their own schema. PRIV_CREATE_RECIPIENT CONSTANT PLS_INTEGER := 2; • PRIV_CONSUME_ORACLE_SHARE Allow users to consume Oracle-to-Oracle live shares. PRIV_CONSUME_ORACLE_SHARE CONSTANT PLS_INTEGER := 4; • PRIV_ORDS_ACL Grant an ACL to the user on the local ORDS endpoint. This is required for the user to generate bearer tokens on locally created shares. PRIV_ORDS_ACL CONSTANT PLS_INTEGER := 8;

GET_ACTIVATION_LINK Function

Generate the link that gets put into emails to the authorized recipient. This activation link leads to the download page, where the recipient clicks a button to get the delta profile.

Syntax

```
FUNCTION GET_ACTIVATION_LINK
(
  recipient_name       IN VARCHAR2,
  recipient_owner      IN VARCHAR2 := NULL,
  expiration           IN PLS_INTEGER := 259200,
  invalidate_previous  IN BOOLEAN := TRUE
)
RETURN VARCHAR2;
```

Parameters

Parameter	Description
recipient_name	The local name of the recipient.
recipient_owner	The schema that owns the recipient.
expiration	Number of seconds before the activation link expires.
invalidate_previous	If TRUE, which is the default, a previously generated activation link is made invalid. If FALSE, a previously generated activation link remains valid.

Example: Print activation link to the screen

```
SQL> exec dbms_share.create_share_recipient('new_recipient',
email=>'anyone@example.com')

PL/SQL procedure successfully completed.

SQL> column PROFILE format A200
SQL> variable sprof varchar2(32767)
SQL> declare
  2    profile sys.json_object_t;
  3  begin
  4    dbms_share.populate_share_profile('NEW_RECIPIENT', profile);
  5    :sprof := profile.to_string;
  6  end;
  7  /
PL/SQL procedure successfully completed.

SQL> exec
dbms_output.put_line(dbms_share.get_activation_link('NEW_RECIPIENT'))
http://.../ords/_adpshr/delta-sharing/download?key=43BA....YXJlX3Rlc3Q=

PL/SQL procedure successfully completed.
```

GET_PUBLISHED_IDENTITY Function

Get data about the current user that was set by `SET_PUBLISHED_IDENTITY`.

Syntax

```
FUNCTION GET_PUBLISHED_IDENTITY
  RETURN CLOB;
```

Example

```
SQL> declare
  2    id_json json_object_t := json_object_t();
  3  begin
  4    id_json.put('name', 'Demo Publisher');
  5    id_json.put('description', 'Documentation Share Provider');
  6    id_json.put('contact', 'null@example.com');
  7    dbms_share.set_published_identity(id_json);
  8  end;
  9  /

PL/SQL procedure successfully completed.

SQL> select json_query(dbms_share.get_published_identity, '$' pretty)
"Published Identity"
  2  from dual;

Published Identity-
--------------------------------------------------------------------------------
-
{
  "name" : "Demo Publisher",
  "description" : "Documentation Share Provider",
  "contact" : "null@example.com"
}
```

GET_RECIPIENT_PROPERTY Function

Return the value of a property for a recipient.

Syntax

```
FUNCTION GET_RECIPIENT_PROPERTY
(
  recipient_name       IN VARCHAR2,
  recipient_property   IN VARCHAR2,
  recipient_owner      IN VARCHAR2 := NULL
)RETURN VARCHAR2;
```

Parameters

Parameter	Description
recipient_name	The name of the recipient.
recipient_property	The property to get. These properties include: • PROP_SHARE_DESC • PROP_SHARE_PUBLIC_DESC • PROP_SHARE_VERSION_ACCESS • PROP_SHARE_JOB_NAME • PROP_SHARE_SPLIT_SIZE • PROP_SHARE_LOG_LEVEL • PROP_SHARE_JOB_DOP • PROP_SHARE_JOB_CLASS • PROP_SHARE_JOB_PRIORITY • PROP_SHARE_VERSION_ACCESS For information on constants used for this parameter, see descriptions for **Share Properties** in DBMS_SHARE Constants.
recipient_owner	The owner of the recipient. Defaults to current user.

GET_SHARE_PROPERTY Function

Get the property value of an existing share.

Syntax

```
FUNCTION GET_SHARE_PROPERTY
(
  share_name       IN VARCHAR2,
  share_property   IN VARCHAR2,
  share_owner      IN VARCHAR2 := NULL
)
RETURN VARCHAR2
```

Parameters

Parameter	Description
share_name	The name of the share.
share_property	The property value to get. For information on constants used for this parameter, see descriptions for **Share Properties** in DBMS_SHARE Constants.
share_owner	The owner of the share. Defaults to current user.

GET_SHARE_TABLE_PROPERTY Function

Get the property value of an existing share table.

Syntax

```
FUNCTION GET_SHARE_TABLE_PROPERTY
(
  share_name           IN VARCHAR2,
```

```
  share_table_name     IN VARCHAR2,
  share_table_property IN VARCHAR2,
  share_schema_name    IN VARCHAR2 := NULL,
  share_owner          IN VARCHAR2 := NULL
)RETURN VARCHAR2;
```

Parameters

Parameter	Description
share_name	The name of the share.
share_table_name	The name of the share table.
share_table_property	The share table property to update. For information on constants used for this parameter, see descriptions for **Share Table Properties** in DBMS_SHARE Constants.
share_schema_name	The name of the share schema. Defaults to uppercase of current user.
share_owner	The owner of the share.

GRANT_TO_RECIPIENT Procedure

Grant access on a share to a specific recipient. The share and recipient must both belong to the same schema.

Syntax

```
PROCEDURE GRANT_TO_RECIPIENT
(
    share_name           IN VARCHAR2,
    recipient_name       IN VARCHAR2,
    owner                IN VARCHAR2 := NULL,
    auto_commit          IN BOOLEAN := FALSE
);
```

Parameters

Parameter	Description
share_name	The name of the share to grant.
recipient_name	The name of the recipient.
owner	The owner of both the share and recipient.
auto_commit	The auto_commit parameter is ignored. This procedure will always commit.

POPULATE_SHARE_PROFILE Procedure

Generate a delta profile for a recipient. You could print this to the screen or upload it somewhere. For example, to an object bucket using DBMS_CLOUD.EXPORT_DATA.

Syntax

```
PROCEDURE POPULATE_SHARE_PROFILE
(
    recipient_name       IN VARCHAR2,
```

```
    share_profile        IN OUT NOCOPY SYS.JSON_OBJECT_T,
    recipient_owner      IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
recipient_name	The local name of the recipient.
share_profile	The share profile, without bearer token.
recipient_owner	The schema that owns the recipient.

Example: Print the profile to the screen.

```
SQL> EXEC DBMS_SHARE.CREATE_SHARE_RECIPIENT('new_recipient',
email=>'anyone@example.com')

PL/SQL procedure successfully completed.
SQL> column PROFILE format A200
SQL> variable sprof varchar2(32767)
SQL> declare
  2    profile sys.json_object_t;
  3  begin
  4    dbms_share.populate_share_profile('NEW_RECIPIENT', profile);
  5    :sprof := profile.to_string;
  6  end;
  7  /

PL/SQL procedure successfully completed.

SQL> select json_query(:sprof, '$' pretty) "PROFILE" from dual;
PROFILE--------------------------------------------------------------------------
-----------------------------------------------
{
  "shareCredentialsVersion" : 1,
  "endpoint" : "https://.../ords/share_test/_delta_sharing/",
  "bearerToken" : "mc7puvhqCpU6xjTOjRdl_w",
  "tokenEndpoint" : "https://.../ords/share_test/oauth/token",
  "clientID" : "VXGQ_44s6qJ-K4WHUNM2yQ..",
  "clientSecret" : "y9ddppgwEmZl7adDHFQndw.."
}
```

PUBLISH_SHARE Procedure

Publish a share and return immediately.

The publication will continue in the background. You can check on the status of the job by querying the `USER_SHARE_JOBS` view. See USER_SHARE_JOBS View for further information.

Syntax

```
PROCEDURE PUBLISH_SHARE
(
    share_name           IN VARCHAR2,
```

```
    share_owner           IN VARCHAR2 := NULL,
    drop_prior_versions   IN BOOLEAN := FALSE,
    share_job_dop         IN NUMBER := NULL,
    share_job_priority    IN NUMBER := NULL,
    job_class             IN VARCHAR2 := 'DEFAULT_JOB_CLASS',
    force_job_class       IN BOOLEAN := FALSE
);
```

Parameters

Parameter	Description
share_name	The name of the share to publish.
share_owner	The share owner, which must be the current user or NULL.
drop_prior_versions	Set to TRUE if you want to drop all prior versions of the share. Note that versions are only dropped if there are no outstanding Pre-Authenticated Requests (PARs).
share_job_dop	Specify the maximum number of dbms_scheduler jobs that will be used to publish the share. Leave as NULL to use the system default number.
share_job_priority	Specify a relative priority of this share publication compared to others. If two shares are being published at the same time, then the one with the highest priority is processed first even if it was started later.
job_class	The scheduler job class, from all_scheduler_job_classes, that are used to publish the share.
force_job_class	Use the specifiedjob_class even if the admin has defined a different default job class.

PUBLISH_SHARE_WAIT Procedure

Publish a share and wait until the background job is complete. The publication continues even if the call is interrupted.

Syntax

```
PROCEDURE PUBLISH_SHARE_WAIT
(
       share_name           IN VARCHAR2,
       share_owner          IN VARCHAR2 := NULL,
       drop_prior_versions  IN BOOLEAN := FALSE,
       share_job_dop        IN NUMBER := NULL,
       share_job_priority   IN NUMBER := NULL,
       job_class            IN VARCHAR2 := 'DEFAULT_JOB_CLASS',
       force_job_class      IN BOOLEAN := FALSE
);
```

Parameters

Parameter	Description
share_name	The name of the share to publish.
share_owner	The share owner, which must be the current user or NULL.
drop_prior_versions	Set to TRUE if you want to drop all prior versions of the share. Note that versions are only dropped if there are no outstanding Pre-Authenticated Requests (PARs).

Parameter	Description
share_job_dop	Specify the maximum number of dbms_scheduler jobs that will be used to publish the share. Leave as NULL to use the system default number.
share_job_priority	Specify a relative priority of this share publication compared to others. If two shares are being published at the same time, then the one with the highest priority is processed first even if it was started later.
job_class	The scheduler job class, from all_scheduler_job_classes, that are used to publish the share.
force_job_class	Use the specifiedjob_class even if the admin has defined a different default job class.

PURGE_DETACHED_FILES Procedure

Delete or forget parquet files that have become detached from their shares.

Syntax

```
PROCEDURE PURGE_DETACHED_FILES
(
    file_pattern          IN VARCHAR2 := '%',
    credential_name       IN VARCHAR2 := NULL,
    purge_mode            IN PLS_INTEGER := PURGE_DROP,
    owner_id              IN NUMBER := SYS_CONTEXT('USERENV',
'CURRENT_USERID')
);
```

A version of PURGE_DETACHED_FILES that returns the ID of the purge job, if any.

```
PROCEDURE PURGE_DETACHED_FILES
(
    purge_job_id          IN OUT NOCOPY NUMBER,
    file_pattern          IN VARCHAR2 := '%',
    credential_name       IN VARCHAR2 := NULL,
    purge_mode            IN PLS_INTEGER := PURGE_DROP,
    owner_id              IN NUMBER := SYS_CONTEXT('USERENV',
'CURRENT_USERID')
);
```

Parameters

Parameter	Description
purge_job_id	The purge job ID.
file_pattern	An optional LIKE pattern for the files that are to be purged.
credential_name	An optional credential to be used to delete the files.

Parameter	Description
purge_mode	Specifies how the files are purged. Purge mode values include: • PURGE_DROP Attempt to drop the files, using the specified credential. If the file cannot be dropped, then the file will continue to be listed in the *_SHARE_DETACHED_FILES views. PURGE_DROP CONSTANT PLS_INTEGER := 1; • PURGE_DROP_FORCE Attempt to drop the files, using the specified credential. The file will be removed from the *_SHARE_DETACHED_FILES views even if the attempt to drop the file fails again. PURGE_DROP_FORCE CONSTANT PLS_INTEGER := 2; • PURGE_FORGET Remove the files *_SHARE_DETACHED_FILES views without attempting to drop them. PURGE_FORGET CONSTANT PLS_INTEGER := 3;
owner_id	The owner ID whose files are to be purged.

REMOVE_FROM_SHARE Procedure

Remove a table or view from a share.

Syntax

```
PROCEDURE REMOVE_FROM_SHARE
(
    share_name           IN VARCHAR2,
    share_table_name     IN VARCHAR2,
    share_schema_name    IN VARCHAR2 := NULL,
    share_owner          IN VARCHAR2 := NULL,
    auto_commit          IN BOOLEAN := FALSE
);
```

Parameters

Parameter	Description
share_name	The name of an existing share to which the object is revoked.
share_table_name	The name of the share table to revoke. This must match the externally visible name, so the input will be uppercased.
share_schema_name	The name of the share schema. Defaults to uppercase of current user.
share_owner	The owner of the share.
auto_commit	If TRUE, this procedure call commits changes that are not visible externally until the commit takes place. The default value is FALSE, which means that the user must COMMIT after running this call in order to make the change visible.

RENAME_RECIPIENT Procedure

Rename a recipient. This procedure only changes the local name of the recipient. The external definition of the recipient, for example the name of the OAUTH user or sharing id, is not changed.

Syntax

```
PROCEDURE RENAME_RECIPIENT
(
    old_recipient_name   IN VARCHAR2,
    new_recipient_name   IN VARCHAR2,
    owner                IN VARCHAR2 := NULL,
    auto_commit          IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
old_recipient_name	The current name of the share recipient.
new_recipient_name	The new name of the share recipient.
owner	The schema that defines the recipient.
auto_commit	If TRUE, the changes are automatically committed. Changes are not visible externally until the commit takes place. The default is TRUE.

RENAME_SHARE Procedure

Rename a share. Care should be take with this procedure since the change effects any existing consumers whose access is based on the previous name. Consumers are not notified directly or updated.

Syntax

```
PROCEDURE RENAME_SHARE
(
    old_share_name       IN VARCHAR2,
    new_share_name       IN VARCHAR2,
    share_owner          IN VARCHAR2 := NULL,
    auto_commit          IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
old_share_name	The current name of the share.
new_share_name	The new name of the share.
share_owner	The owner of the share.
auto_commit	If TRUE, the changes are automatically committed. Changes are not visible externally until the commit takes place. The default is TRUE.

RENAME_SHARE_LINK Procedure

Rename a registered share link.

Syntax

```
PROCEDURE RENAME_SHARE_LINK
(
    old_name             IN VARCHAR2,
    new_name             IN VARCHAR2,
    link_owner           IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
old_name	The current name of the share link.
new_name	The new name of the link.
link_owner	The owner of the share link. Defaults to current schema.

RENAME_SHARE_SCHEMA Procedure

Rename a share schema.

Syntax

```
PROCEDURE RENAME_SHARE_SCHEMA
(
    share_name           IN VARCHAR2,
    old_schema_name      IN VARCHAR2,
    new_schema_name      IN VARCHAR2,
    share_owner          IN VARCHAR2 := NULL,
    auto_commit          IN BOOLEAN := FALSE
);
```

Parameters

Parameter	Description
share_name	The name of the share.
old_schema_name	The old name of the schema.
new_schema_name	The new name of the schema.
share_owner	The owner of the share. Defaults to the current schema.
auto_commit	If TRUE, the changes are automatically committed. Changes are not visible externally until the commit takes place. The default is FALSE.

RENAME_SHARE_TABLE Procedure

Rename a share table.

Syntax

```
PROCEDURE RENAME_SHARE_TABLE
(
    share_name           IN VARCHAR2,
    share_schema_name    IN VARCHAR2,
    old_share_table_name IN VARCHAR2,
    new_share_table_name IN VARCHAR2,
    share_owner          IN VARCHAR2 := NULL,
    auto_commit          IN BOOLEAN := FALSE
);
```

Parameters

Parameter	Description
share_name	The name of the share.
share_schema_name	The name of the share schema.
old_share_table_name	The old name of the share table.
new_share_table_name	The new name of the share table.
share_owner	The owner of the share.
auto_commit	If TRUE, the changes are automatically committed. Changes are not visible externally until the commit takes place. The default is FALSE.

REVOKE_FROM_RECIPIENT Procedure

Revoke access on a share from a specific recipient.

Syntax

```
PROCEDURE REVOKE_FROM_RECIPIENT
(
    share_name           IN VARCHAR2,
    recipient_name       IN VARCHAR2,
    owner                IN VARCHAR2 := NULL,
    auto_commit          IN BOOLEAN := FALSE
);
```

Parameters

Parameter	Description
share_name	The name of the share to revoke.
recipient_name	The name of the recipient.
owner	The owner of the share and recipient.
auto_commit	If TRUE, the changes are automatically committed. Changes are not visible externally until the commit takes place. The default is FALSE.

SET_CURRENT_SHARE_VERSION Procedure

Change the current version of a share.

Syntax

```
PROCEDURE SET_CURRENT_SHARE_VERSION
(
    share_name          IN VARCHAR2,
    share_version       IN NUMBER,
    share_owner         IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
share_name	The name of the share.
share_version	The new version or NULL. The version must exist and must be valid. If the share_version is NULL, then no version will be marked as CURRENT and the share will be "unpublished".
share_owner	The owner of the share. Defaults to the current schema.

SET_PUBLISHED_IDENTITY Procedure

Set data about the current user that will be supplied to recipients of published ORACLE shares.

Syntax

```
PROCEDURE SET_PUBLISHED_IDENTITY
(
    metadata              IN SYS.JSON_OBJECT_T := NULL
);
```

Parameters

Parameter	Description
metadata	If the provider identity has already been set, then only the items that the caller wants to update need to be included in the JSON. Supplying a NULL value for an item causes that item to be removed from the stored identity. However, "name", "description" and "contact" cannot be removed in this way. If the metadata argument is NULL, the existing provider identity is deleted. This can only happen if the current user has no published shares. If the provider identity has not yet been set, the JSON must contain at least: • name (<=128 bytes) • description (<=4000 bytes) • contact (<=4000 bytes) values Additional items may be included at the caller's discretion.

Example: Include additional items in the JSON metadata

In addition to the required "name", "description", and "contact", this example includes additional item "schedule".

```
{
  "name" : "A sample share provider",
  "description" : "Test of share provider metadata",
  "contact" : "provider1@example.com".
  "schedule" : "New data provided on alternate rainy Fridays"
}
```

Example: Update "description" and remove "schedule"

To update "description" and remove "schedule", the following JSON can be used:

```
{
  "description" : "The Share Provider You Can Trust!(tm)",
  "schedule" : null
}
```

SET_RECIPIENT_LOG_LEVEL Procedure

Change the log level for an existing share recipient.

Syntax

```
PROCEDURE SET_RECIPIENT_LOG_LEVEL
(
  recipient_name       IN VARCHAR2,
  log_level            IN PLS_INTEGER,
  recipient_owner      IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
recipient_name	The local name of the share recipient.
log_level	Event logging level. For information on constants used for this parameter, see descriptions for **Log Level** in DBMS_SHARE Constants.
recipient_owner	The schema that owns the recipient.

SET_SHARE_LOG_LEVEL Procedure

Change the log level for an existing share.

Syntax

```
PROCEDURE SET_SHARE_LOG_LEVEL
(
  share_name           IN VARCHAR2,
  log_level            IN PLS_INTEGER,
```

```
  share_owner           IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
share_name	The name of the share.
log_level	Event logging level. For information on constants used for this parameter, see descriptions for **Log Level** in DBMS_SHARE Constants.
share_owner	The owner of the share.

SET_STORAGE_CREDENTIAL Procedure

Set the credential name used by the current user when it attempts to access the given storage.

Syntax

```
PROCEDURE SET_STORAGE_CREDENTIAL
(
    storage_link_name    IN VARCHAR2,
    credential_name      IN VARCHAR2,
    owner                IN VARCHAR2 := NULL,
    check_if_exists      IN BOOLEAN := TRUE,
    auto_commit          IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
storage_link_name	The name of a cloud storage link previously created using CREATE_CLOUD_STORAGE_LINK Procedure.
credential_name	The name of a local credential that gives access to the storage. The credentials used for share storage must be able to read, write, delete, and manage pre-authenticated requests. See Using Pre-Authenticated Requests
owner	The owner of the cloud storage link. Leave as NULL for the current user.
check_if_exists	If check_if_exists is TRUE, then the function will also confirm that the credential exists for the current user.
auto_commit	The default is TRUE. If TRUE, this transaction is committed. If FALSE, the user must commit the transaction. Changes are not visible until the commit takes place.

STOP_JOB Procedure

Attempt to stop a running share job. The procedure should return quickly, but it may take some time for the associated job to stop.

Syntax

```
PROCEDURE STOP_JOB
(
  share_job_id          IN NUMBER,
```

```
  share_job_owner       IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
share_job_id	The ID of the share job.
share_job_owner	The owner of the job. Defaults to the current schema.

UNPUBLISH_SHARE Procedure

Unpublish a share.

Syntax

```
PROCEDURE UNPUBLISH_SHARE
(
    share_name            IN VARCHAR2,
    out_share_job_id      IN OUT NOCOPY NUMBER,
    share_owner           IN VARCHAR2 := NULL,
    drop_all_versions     IN BOOLEAN := FALSE,
    restart_versions      IN BOOLEAN := FALSE
);
```

Parameters

Parameter	Description
share_name	The name of the share.
out_share_job_id	The ID of any share job that needs to be run to process this command.
share_owner	The owner of the share. Defaults to the current schema.
drop_all_versions	If set to TRUE, all existing versions along with all associated storage is dropped. If drop_all_versions is FALSE, then all existing versions continue to exist, but the share is not visible to recipients. A subsequent call to PUBLISH_SHARE or SET_CURRENT_SHARE_VERSION makes them visible again.
restart_versions	Restart version numbering. If FALSE, then the next published version number will be the same as it would have been if no versions had been dropped. If TRUE, then the next published version will be set to one more than the highest existing version. When used in conjunction with drop_all_versions, this has the effect of resetting the share to its original state. Note that this may confuse any existing delta clients, so it should be used with care.

UPDATE_DEFAULT_RECIPIENT_PROPERTY Procedure

Update the default recipient property values. This procedure requires the user to have admin privileges.

Syntax

```
PROCEDURE UPDATE_DEFAULT_RECIPIENT_PROPERTY
(
  recipient_property   IN VARCHAR2,
  new_value_vc         IN VARCHAR2
);
```

Parameters

Parameter	Description
recipient_property	The property to update. For information on constants used for this parameter, see descriptions for **Share Recipient Properties** in DBMS_SHARE Constants.
new_value_vc	The new property value.

UPDATE_DEFAULT_SHARE_PROPERTY Procedure

Update the default share property values.

Syntax

```
PROCEDURE UPDATE_DEFAULT_SHARE_PROPERTY
(
    share_property       IN VARCHAR2,
    new_value            IN VARCHAR2
);
```

Parameters

Parameter	Description
share_property	The property to update. For information on constants used for this parameter, see descriptions for **Share Properties** in DBMS_SHARE Constants. These properties can be read using the ALL_SHARE_DEFAULT_SETTINGS View.
new_value	The new property value.

UPDATE_RECIPIENT_PROPERTY Procedure

Update a property of an existing recipient.

Syntax

```
PROCEDURE UPDATE_RECIPIENT_PROPERTY
(
```

```
    recipient_name       IN VARCHAR2,
    recipient_property   IN VARCHAR2,
    new_value            IN VARCHAR2,
    recipient_owner      IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
recipient_name	The name of the recipient.
recipient_property	The property to update. These properties include: • PROP_RECIPIENT_EMAIL • PROP_RECIPIENT_DESCRIPTION • PROP_RECIPIENT_LOG_LEVEL • PROP_RECIPIENT_SHARING_ID • PROP_RECIPIENT_PROFILE_URL • PROP_RECIPIENT_PAR_TYPE • PROP_RECIPIENT_MIN_PAR_LIFETIME • PROP_RECIPIENT_PAR_LIFETIME • PROP_RECIPIENT_TOKEN_LIFETIME For information on constants used for this parameter, see descriptions for **Share Recipient Recipient Properties** in DBMS_SHARE Constants.
new_value	The new property value.
recipient_owner	The owner of the recipient. Defaults to the current schema.

UPDATE_SHARE_JOB_PROPERTY Procedure

Modify properties of a running share job. The procedure should return quickly, but it may take some time for the changes to take effect.

Syntax

```
PROCEDURE UPDATE_SHARE_JOB_PROPERTY
(
    share_job_id         IN NUMBER,
    share_property       IN VARCHAR2,
    new_value            IN VARCHAR2,
    share_job_owner      IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
share_job_id	The ID of the share job.
share_property	The property to update. For information on constants used for this parameter, see descriptions for **Share Job Properties** in DBMS_SHARE Constants.
new_value	The new property value.
share_job_owner	The owner of the job. Defaults to the current schema.

UPDATE_SHARE_PROPERTY Procedure

Update a property of an existing share.

Syntax

```
PROCEDURE UPDATE_SHARE_PROPERTY
(
    share_name           IN VARCHAR2,
    share_property       IN VARCHAR2,
    new_value            IN VARCHAR2,
    share_owner          IN VARCHAR2 := NULL,
    auto_commit          IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
share_name	The name of the share.
share_property	The property to update. For information on constants used for this parameter, see descriptions for **Share Properties** in DBMS_SHARE Constants.
new_value	The new property value.
share_owner	The owner of the share. Defaults to the current schema.
Auto_commit	If TRUE (the default), the changes are automatically committed. If FALSE, the user must commit the changes. Changes are visible externally after the commit takes place.

UPDATE_SHARE_TABLE_PROPERTY Procedure

Update the property value of an existing share table.

Syntax

```
PROCEDURE UPDATE_SHARE_TABLE_PROPERTY
(
    share_name           IN VARCHAR2,
    share_table_name     IN VARCHAR2,
    share_table_property IN VARCHAR2,
    new_value            IN VARCHAR2,
    share_schema_name    IN VARCHAR2 := NULL,
    share_owner          IN VARCHAR2 := NULL,
    auto_commit          IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
share_name	The name of the share.
share_table_name	The name of the share table.

Parameter	Description
share_table_property	The property to update. These properties include: • PROP_SHARE_TABLE_SPLIT_METHOD • PROP_SHARE_TABLE_SPLIT_COLUMNS • PROP_SHARE_TABLE_ORDER_COLUMNS • PROP_SHARE_TABLE_SHARE_COLUMNS • PROP_SHARE_TABLE_SPLIT_SIZE • PROP_SHARE_TABLE_GATHER_STATS • PROP_SHARE_TABLE_SPLIT_ROWS • PROP_SHARE_TABLE_FLASHBACK
new_value	The new property value.
share_schema_name	The name of the share schema (defaults to uppercase of current user).
share_owner	The owner of the share. Defaults to the current schema.
Auto_commit	If TRUE (the default), the changes are automatically committed. If FALSE, the user must commit the changes. Changes are visible externally after the commit takes place.

VALIDATE_CREDENTIAL Function

Validate a credential name, converting it to canonical form first if required.

The procedure raises an exception if the name is not a valid Oracle identifier. The credential_name is returned without double quotes, as it would appear in the CREDENTIAL_NAME column of the USER_CREDENTIALS view.

Syntax

```
FUNCTION VALIDATE_CREDENTIAL
(
        credential_name      IN VARCHAR2,
        check_if_exists      IN BOOLEAN := TRUE
)
RETURN VARCHAR2;
```

Parameters

Parameter	Description
credential_name	The name of the credential to validate in standard database form, with double quotes if the name is not a simple identifier.
check_if_exists	If TRUE, then the function also confirms that the credential exists for the current user.

VALIDATE_SHARE_STORAGE Procedure

Check to see if the given storage is suitable for versioned shares.

Syntax

The procedure syntax including the `validation_results` output parameter.

```
 PROCEDURE VALIDATE_SHARE_STORAGE
(
    storage_link_name    IN VARCHAR2,
    validation_results   IN OUT NOCOPY VARCHAR2,
    run_storage_tests    IN BOOLEAN := TRUE,
    storage_link_owner   IN VARCHAR2 := NULL
);
```

The procedure syntax not including the `validation_results` output parameter.

```
PROCEDURE VALIDATE_SHARE_STORAGE
(
    storage_link_name    IN VARCHAR2,
    run_storage_tests    IN BOOLEAN := TRUE,
    storage_link_owner   IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
`storage_link_name`	The cloud storage link name.
`validation_results`	(Optional input and output) Validation result details returned in JSON form. The `validation_results` include the results for each separate test.
`run_storage_tests`	Run tests to validate the storage. If TRUE (the default), the procedure tests READ, WRITE, DELETE, and PREAUTHENTICATED REQUESTS.
`storage_link_owner`	The cloud storage link owner.

Example: Validation results

```
{
  "READ":"PASSSED",
  "WRITE":"PASSSED",
  "CREATE_PAR":"PASSSED",
  "DELETE_PAR":"PASSSED",
  "DELETE":"PASSSED"
}
```

WAIT_FOR_JOB Procedure

This procedure waits until the specified share job is complete.

Syntax

```
PROCEDURE WAIT_FOR_JOB
(
    share_job_id          IN NUMBER,
    completed             IN OUT NOCOPY BOOLEAN,
    maximum_wait          IN NUMBER := NULL
);
```

Parameters

Parameter	Description
share_job_id	The ID of the share job.
completed	Job completion indicator.
maximum_wait	The maximum wait period, in seconds. A NULL value implies no limit.

Summary of Share Consumer Subprograms

This table lists the DBMS_SHARE package procedures and functions used to consume shares.

Subprograms	Description
ASSERT_SHARING_ID Procedure	Run basic validation checks against a sharing id and return one in canonical form.
CREATE_CREDENTIALS Procedure and Function	Create a credential containing the bearer token from a delta profile.
CREATE_OR_REPLACE_SHARE_LINK Procedure	Subscribe to share from a registered share provider.
CREATE_OR_REPLACE_ORACLE_SHARE_PROVIDER Procedure	Subscribe to an Oracle share provider, with a local name.
CREATE_ORACLE_SHARE_PROVIDER Procedure	Subscribe to an Oracle share provider, with a local name.
CREATE_SHARE_LINK Procedure	Subscribe to share from a registered share provider.
CREATE_SHARE_LINK_VIEW Procedure	Create or replace a named view that gives access to a remote shared table.
CREATE_SHARE_PROVIDER Procedure	Subscribe to a delta share provider.
DISCOVER_AVAILABLE_SHARES Function	Return one SHARE_AVAILABLE_SHARES_ROW for each available table from the subscribed share providers.
DISCOVER_AVAILABLE_TABLES Function	Return one SHARE_AVAILABLE_TABLES_ROW for each available table from the subscribed share providers or from an explicit delta endpoint.
DROP_SHARE_LINK Procedure	Drop a share link created by the CREATE_SHARE_LINK procedure.
DROP_SHARE_PROVIDER Procedure	Drop a subscription to a share provider.

Subprograms	Description
ENABLE_DELTA_ENDPOINT Procedure	Create necessary ACLs that allow the specified user to connect to a delta endpoint.
FLUSH_SHARE_LINK_CACHE Procedure	Flush the cache of shares for a given share link.
FLUSH_SHARE_PROVIDER_CACHE Procedure	Flush the cache of shares for a given share provider.
GENERATE_SHARE_LINK_SELECT Procedure and Function	Generate a SELECT statement that returns data from a shared table.
GET_ORACLE_SHARE_LINK_INFO Function	Retrieve the cloud link name and namespace for an Oracle-to-Oracle share.
GET_SHARE_LINK_INFO Procedure	Get the endpoint(s), share type and share name along with any additional JSON metadata for a share link.
GET_SHARE_PROVIDER_CREDENTIAL Procedure	Get the credential name to be used by the current user when it attempts to access the given delta share provider.
GET_SHARE_PROVIDER_INFO Procedure	Get the endpoint string(s) and share type along with any additional JSON metadata for a share provider.
GET_SHARING_ID Function	Return an identifier that can be used as the `sharing_id` in the `CREATE_SHARE_RECIPIENT` procedure.
OPEN_SHARE_LINK_CURSOR Procedure	Open a cursor that returns data from a shared table.
REFRESH_BEARER_TOKEN_CREDENTIAL Procedure	Update one or more credentials created by `CREATE_BEARER_TOKEN_CREDENTIAL` or `CREATE_CREDENTIALS`.
RENAME_CLOUD_STORAGE_LINK Procedure	Rename a registered cloud storage link.
RENAME_SHARE_LINK Procedure	Rename a registered share link.
RENAME_SHARE_PROVIDER Procedure	Rename a registered share provider.
REMOVE_SHARE_SCHEMA Procedure	Remove a schema and all its contents from a share.
SET_SHARE_LINK_METADATA Procedure	Set the additional JSON metadata for the share link.
SET_SHARE_PROVIDER_CREDENTIAL Procedure	Set the credential name to access the given share provider.
SET_SHARE_PROVIDER_METADATA Procedure	Set additional JSON metadata for the share provider.
UPDATE_BEARER_TOKEN_CREDENTIAL Procedure	Modify an attribute of a credential created by `CREATE_CREDENTIALS` or `CREATE_BEARER_TOKEN_CREDENTIAL`.

- ASSERT_SHARING_ID Procedure
  Run basic validation checks against a sharing id and return one in canonical form. An exception is raised if the id is obviously invalid.
- CREATE_CREDENTIALS Procedure and Function
  Create a credential containing the bearer token from a delta sharing profile. The standard type, version 1, specifies an endpoint and a single long term bearer token.
- CREATE_OR_REPLACE_SHARE_LINK Procedure
  Subscribe to share from a registered share provider.
- CREATE_OR_REPLACE_ORACLE_SHARE_PROVIDER Procedure
  Subscribe to an Oracle share provider, with a local name.

- CREATE_ORACLE_SHARE_PROVIDER Procedure
  Subscribe to an Oracle share provider, with a local name.
- CREATE_SHARE_LINK Procedure
  Subscribe to share from a registered share provider. The available share names can be found by calling `DISCOVER_AVAILABLE_SHARES`.
- CREATE_SHARE_LINK_VIEW Procedure
  Create or replace a named view that gives access to a remote shared table.
- CREATE_SHARE_PROVIDER Procedure
  Subscribe to a delta share provider.
- DISCOVER_AVAILABLE_SHARES Function
  Return one `SHARE_AVAILABLE_SHARES_ROW` for each available table from the subscribed share providers.
- DISCOVER_AVAILABLE_TABLES Function
  Return one `SHARE_AVAILABLE_TABLES_ROW` for each available table from the subscribed share providers.
- DROP_SHARE_LINK Procedure
  Drop a share link created by the `CREATE_SHARE_LINK` procedure.
- DROP_SHARE_PROVIDER Procedure
  Drop a subscription to a share provider.
- ENABLE_DELTA_ENDPOINT Procedure
  Create necessary ACLs that allow the specified user to connect to a delta endpoint. Admin privileges are required for this procedure.
- FLUSH_SHARE_LINK_CACHE Procedure
  Flush the cache of shares for a given share link. The list of shares for the remote endpoints is fetched instead of relying on cached values.
- FLUSH_SHARE_PROVIDER_CACHE Procedure
  Flush the cache of shares for a given share provider. The list of shares for the remote endpoints is fetched instead of relying on cached values.
- GENERATE_SHARE_LINK_SELECT Procedure and Function
  Generate a SELECT statement that returns data from a shared table.
- GET_ORACLE_SHARE_LINK_INFO Function
  Retrieve the cloud link name and namespace for an Oracle-to-Oracle share.
- GET_SHARE_LINK_INFO Procedure
  Get the endpoint(s), share type and share name along with any additional JSON metadata for a share link.
- GET_SHARE_PROVIDER_CREDENTIAL Procedure
  Get the credential name to be used by the current user when it attempts to access the given delta share provider.
- GET_SHARE_PROVIDER_INFO Procedure
  Get the endpoint string(s) and share type along with any additional JSON metadata for a share provider. For ORACLE share providers, the Oracle provider ID is returned in the endpoint argument.
- GET_SHARING_ID Function
  Return an identifier that can be used as the `sharing_id` in the `CREATE_SHARE_RECIPIENT` procedure. This function can be used to share data between two users, the "provider" and the "recipient", in different databases.

- OPEN_SHARE_LINK_CURSOR Procedure
  Open a cursor that returns data from a shared table.
- REFRESH_BEARER_TOKEN_CREDENTIAL Procedure
  Update one or more credentials created by `CREATE_BEARER_TOKEN_CREDENTIAL` or `CREATE_CREDENTIALS` by calling the registered token endpoints and fetching new bearer tokens. Note this procedure is called automatically by a scheduler job, `ADP$BEARER_REFRESH_JOB`, that runs every 50 minutes.
- RENAME_CLOUD_STORAGE_LINK Procedure
  Rename a registered cloud storage link.
- RENAME_SHARE_PROVIDER Procedure
  Rename a registered share provider.
- REMOVE_SHARE_SCHEMA Procedure
  Remove a schema and all its contents from a share.
- SET_SHARE_LINK_METADATA Procedure
  Set the additional JSON metadata for the share link.
- SET_SHARE_PROVIDER_CREDENTIAL Procedure
  Set the credential name to be used by the current user when it attempts to access the given share provider.
- SET_SHARE_PROVIDER_METADATA Procedure
  Set additional JSON metadata for the share provider.
- UPDATE_BEARER_TOKEN_CREDENTIAL Procedure
  Modify an attribute of a credential created by `CREATE_CREDENTIALS` or `CREATE_BEARER_TOKEN_CREDENTIAL`.

ASSERT_SHARING_ID Procedure

Run basic validation checks against a sharing id and return one in canonical form. An exception is raised if the id is obviously invalid.

Syntax

```
PROCEDURE ASSERT_SHARING_ID
(
    sharing_id              IN OUT NOCOPY VARCHAR2,
    out_sharing_id_type     IN OUT NOCOPY VARCHAR2
);
```

Parameters

Parameter	Description
`sharing_id`	The id to check.
`out_sharing_id_type`	The type of the id, if valid. For example, `TENANCY` or `DATABASE`.

CREATE_CREDENTIALS Procedure and Function

Create a credential containing the bearer token from a delta sharing profile. The standard type, version 1, specifies an endpoint and a single long term bearer token.

Syntax

```
PROCEDURE CREATE_CREDENTIALS
(
    credential_base_name IN VARCHAR2,
    delta_profile        IN CLOB,
    out_credential_name  IN OUT NOCOPY VARCHAR2
);
```

Below is the functional version of the `create_credentials` that returns the name of the new credentials in JSON form.

```
FUNCTION CREATE_CREDENTIALS
(
   credential_base_name IN VARCHAR2,
   delta_profile        IN CLOB
)
RETURN CLOB;
```

Parameters

Parameter	Description
`credential_base_name`	The base name of the credential(s) to create.

`delta_profile`	The delta sharing profile, in JSON format, obtained from the share provider.
	
	`{`
	`  "shareCredentialsVersion": 1,`
	`  "endpoint": "https://<endpoint>",`
	`  "bearerToken": "<token>",`
	`  "expirationTime": "...",`
	`}`
	
	The profile may also, optionally, include a `tokenEndpoint` property along with a `clientID` and `clientSecret`.
	
	`{`
	`  "shareCredentialsVersion": 1,`
	`  "endpoint": "https://<endpoint>",`
	`  "bearerToken": "<token>",`
	`  "expirationTime": "...",`
	`  "tokenEndpoint": "https://<token endpoint>",`
	`  "clientID": "<client id>",`
	`  "clientSecret": "<client secret>"`
	`}`
	
	See Profile File Format and Bearer Token for further information.

`out_credential_name`	The name of the newly created bearer token credential.

CREATE_OR_REPLACE_SHARE_LINK Procedure

Subscribe to share from a registered share provider.

Syntax

```
PROCEDURE CREATE_OR_REPLACE_SHARE_LINK
(
    share_link_name        IN VARCHAR2,
    share_provider         IN VARCHAR2,
    share_name             IN VARCHAR2,
    provider_owner         IN VARCHAR2 := NULL,
    link_owner             IN VARCHAR2 := NULL,
    use_default_credential IN BOOLEAN := TRUE,
    metadata               IN SYS.JSON_OBJECT_T := NULL,
    auto_commit            IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
share_link_name	The name of the share link that will be created. The available share names can be found by calling DISCOVER_AVAILABLE_SHARES Function.
share_provider	The name of the registered share provider.
share_name	The name of the share from the share provider.
provider_owner	The owner of the share provider. Defaults to the current schema.
link_owner	The owner of the share link. Defaults to the current schema.
use_default_credential	If TRUE, use the same credentials as the share provider.
metadata	Optional metadata to associate with the share link.
auto_commit	If TRUE (the default), this procedure call commits changes that are not visible externally until the commit takes place. If FALSE, the user must COMMIT after running this call in order to make the change visible.

CREATE_OR_REPLACE_ORACLE_SHARE_PROVIDER Procedure

Subscribe to an Oracle share provider, with a local name.

It will then appear in ALL_SHARE_PROVIDERS View with `RECIPIENT_TYPE = 'ORACLE'`.

> **Note:**
>
> Use the `SET_STORAGE_CREDENTIAL` procedure to add a credential to the storage link. See SET_STORAGE_CREDENTIAL Procedure.

Syntax

```
PROCEDURE CREATE_OR_REPLACE_CLOUD_STORAGE_LINK
(
    storage_link_name    IN VARCHAR2,
```

```
    uri                 IN VARCHAR2,
    owner               IN VARCHAR2 := NULL,
    metadata            IN SYS.JSON_OBJECT_T := NULL,
    auto_commit         IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
storage_link_name	The name of the cloud storage link. The name of the link should follow standard Oracle naming conventions. See Database Object Names and Qualifiers
uri	The URI of the storage bucket. The URI should be of the form https://objectstorage.region.oraclecloud.com/n/namespace-string/b/bucket/o
owner	The owner of the storage link. Defaults to the current schema.
metadata	An optional JSON metadata document containing additional information.
auto_commit	If TRUE, the changes automatically commit after creating the link. The default is TRUE.

CREATE_ORACLE_SHARE_PROVIDER Procedure

Subscribe to an Oracle share provider, with a local name.

It will then appear in ALL_SHARE_PROVIDERS View with RECIPIENT_TYPE = 'ORACLE'.

Syntax

```
PROCEDURE CREATE_ORACLE_SHARE_PROVIDER
(
    oracle_provider_id   IN VARCHAR2,
    provider_name        IN VARCHAR2,
    owner                IN VARCHAR2 := NULL,
    metadata             IN SYS.JSON_OBJECT_T := NULL,
    auto_commit          IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
oracle_provider_id	The provider ID obtained from the ALL_AVAILABLE_ORACLE_SHARE_PROVIDERS view. See ALL_AVAILABLE_ORACLE_SHARE_PROVIDERS View.
provider_name	A local name for the provider
owner	The owner of the new share provider. Leave as NULL for the current user.
metadata	Optional JSON metadata to associate with the provider.
auto_commit	If TRUE, the changes automatically commit after creating the link. The default is TRUE.

CREATE_SHARE_LINK Procedure

Subscribe to share from a registered share provider. The available share names can be found by calling `DISCOVER_AVAILABLE_SHARES`.

Syntax

```
PROCEDURE CREATE_SHARE_LINK
(
    share_link_name      IN VARCHAR2,
    share_provider       IN VARCHAR2,
    share_name           IN VARCHAR2,
    provider_owner       IN VARCHAR2 := NULL,
    link_owner           IN VARCHAR2 := NULL,
    use_default_credential IN BOOLEAN := TRUE,
    metadata             IN SYS.JSON_OBJECT_T := NULL,
    auto_commit          IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
`share_link_name`	The name of the share link to create. This should follow standard Oracle naming conventions and does not need to be the same as the `share_name` parameter.
`share_provider`	The name of the share provider, listed in `ALL_SHARE_PROVIDERS`, that shared the data.
`share_name`	The name of the share, as defined by the share provider.
`provider_owner`	The schema that subscribed the share provider, as listed in `ALL_SHARE_PROVIDER`. Leave as null to default to the current user, which is typical.
`link_owner`	The owner of the new share link. Leave as null to default to the current user, which is typical.
`use_default_credential`	Set to TRUE to use the same credential for both the share link and the share provider.
`metadata`	Optional JSON metadata to associated with the share link.
`auto_commit`	If TRUE, the changes automatically commit after creating the link. The default is TRUE.

CREATE_SHARE_LINK_VIEW Procedure

Create or replace a named view that gives access to a remote shared table.

Syntax

```
PROCEDURE CREATE_SHARE_LINK_VIEW
(
    view_name            IN VARCHAR2,
    share_link_name      IN VARCHAR2,
    share_schema_name    IN VARCHAR2,
    share_table_name     IN VARCHAR2,
    view_owner           IN VARCHAR2 := NULL,
```

```
    share_link_owner     IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
view_name	The name of the new view.
share_link_name	The name of the share link.
share_schema_name	The name of the shared schema.
share_table_name	The name of the shared table.
view_owner	The view owner. Defaults to the current schema.
share_link_owner	The link owner. Defaults to the current schema.

CREATE_SHARE_PROVIDER Procedure

Subscribe to a delta share provider.

Syntax

```
PROCEDURE CREATE_SHARE_PROVIDER
(
  provider_name         IN VARCHAR2,
  endpoint              IN VARCHAR2,
  token_endpoint        IN VARCHAR2 := NULL,
  share_type            IN VARCHAR2 := 'DELTA',
  owner                 IN VARCHAR2 := NULL,
  metadata              IN SYS.JSON_OBJECT_T := NULL,
  auto_commit           IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
provider_name	The local name of this share provider.
endpoint	The delta endpoint, from the delta sharing profile.
token_endpoint	This parameter is ignored.
share_type	The type of share provider. Leave this as DELTA.
owner	The owner of the share provider. Defaults to the current schema.
metadata	Optional JSON metadata to associate with the share provider.
auto_commit	If TRUE (the default), this procedure call commits changes that are not visible externally until the commit takes place. If FALSE, the user must COMMIT after running this call in order to make the change visible.

DISCOVER_AVAILABLE_SHARES Function

Return one SHARE_AVAILABLE_SHARES_ROW for each available table from the subscribed share providers.

Syntax

```
FUNCTION DISCOVER_AVAILABLE_SHARES
(
    share_provider       IN VARCHAR2,
    owner                IN VARCHAR2 := NULL
) RETURN share_available_shares_tbl PIPELINED;
```

Parameters

Parameter	Description
share_provider	The name of the share provider.
owner	The owner of the share provider. Defaults to the current schema.

Example: Discover a list of shares available from a given provider

```
SQL> select available_share_name
  2    from dbms_share.discover_available_shares('share_prov')
  3   order by available_share_name;

AVAILABLE_SHARE_NAME
--------------------------------------------------------------------------------
--
BURLINGTON_EXPEDITION_2022

EGYPT_EXPEDITION_2022
```

DISCOVER_AVAILABLE_TABLES Function

Return one `SHARE_AVAILABLE_TABLES_ROW` for each available table from the subscribed share providers.

Syntax

```
FUNCTION DISCOVER_AVAILABLE_TABLES
(
  share_provider       IN VARCHAR2 := NULL,
  share_name           IN VARCHAR2 := NULL,
  owner                IN VARCHAR2 := NULL,
  endpoint             IN VARCHAR2 := NULL,
  credential_name      IN VARCHAR2 := NULL
) RETURN share_available_tables_tbl PIPELINED;
```

Parameters

Parameter	Description
share_provider	An optional share provider name. If NULL, search all subscribed share providers.
share_name	An optional share name. If NULL, search all discovered shares.
owner	The owner of the share provider. Defaults to the current schema.

Parameter	Description
endpoint	An optional delta endpoint.
credential_name	An optional bearer token credential to access the endpoint.

Example: List shares available from all subscribed share providers

```
SQL> select * from dbms_share.discover_available_tables()
  2  order by share_name, schema_name, table_name;

 PROVIDER_NAME              PROVIDER_OWNER  SHARE_NAME
------------------------- --------------- ------------------------------
SCHEMA_NAME               TABLE_NAME
------------------------- -------------------------
My Test Oracle Provider   ADP_SHARE_TEST  BURLINGTON_EXPEDITION_2022
SH                        COUNTRIES

My Test Oracle Provider   ADP_SHARE_TEST  BURLINGTON_EXPEDITION_2022
SH                        SH_COUNTRIES

My Test Oracle Provider   ADP_SHARE_TEST  EGYPT_EXPEDITION_2022
SHARED_SCHEMA             SHARED_VIEW_1

My Test Oracle Provider   ADP_SHARE_TEST  EGYPT_EXPEDITION_2022
SHARED_SCHEMA             SHARED_VIEW_2
```

Example: List tables available from an unsubscribed endpoint

```
SQL> exec dbms_cloud.create_credential('MY_CRED', 'BEARER_TOKEN', '123456')
PL/SQL procedure successfully completed.
SQL> column share_name format a13
SQL> column table_name format a20
SQL> column schema_name format a10
SQL> select share_name, schema_name, table_name
2   from dbms_share.discover_available_tables(
3     endpoint=>'https://my_endpoint',
4     credential_name=>'MY_CRED')
5   order by 1, 2, 3;

SHARE_NAME SCHEMA_NAM TABLE_NAME
------------ ---------- --------------------
DELTA_SHARING DEFAULT BOSTON-HOUSING
DELTA_SHARING DEFAULT COVID_19_NYT
DELTA_SHARING DEFAULT FLIGHT-ASA_2008
DELTA_SHARING DEFAULT LENDING_CLUB
DELTA_SHARING DEFAULT NYCTAXI_2019
DELTA_SHARING DEFAULT NYCTAXI_2019_PART
DELTA_SHARING DEFAULT OWID-COVID-DATA

7 rows selected.
```

DROP_SHARE_LINK Procedure

Drop a share link created by the `CREATE_SHARE_LINK` procedure.

See CREATE_SHARE_LINK Procedure for further information.

Syntax

```
PROCEDURE DROP_SHARE_LINK
(
    link_name            IN VARCHAR2,
    link_owner           IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
`link_name`	The name of the share link to drop.
`link_owner`	The owner of the share link to drop. Leave as null for the current schema.

DROP_SHARE_PROVIDER Procedure

Drop a subscription to a share provider.

Syntax

```
PROCEDURE DROP_SHARE_PROVIDER
(
    provider_name        IN VARCHAR2,
    owner                IN VARCHAR2 := NULL,
    drop_credentials     IN BOOLEAN := FALSE
);
```

Parameters

Parameter	Description
`provider_name`	The name of the share provider to drop.
`owner`	The owner of the share provider to drop. Defaults to the current schema.
`drop_credentials`	If TRUE, any credentials associated with the provider are dropped. If FALSE, credentials are not dropped.

ENABLE_DELTA_ENDPOINT Procedure

Create necessary ACLs that allow the specified user to connect to a delta endpoint. Admin privileges are required for this procedure.

Syntax

```
PROCEDURE ENABLE_DELTA_ENDPOINT
(
    schema_name          IN VARCHAR2,
```

```
    delta_profile         IN CLOB,
    enabled               IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
schema_name	The schema to enable or disable.
delta_profile	The delta profile. Only the endpoint and tokenEndpoint are required.
enabled	TRUE to enable and FALSE to disable.

FLUSH_SHARE_LINK_CACHE Procedure

Flush the cache of shares for a given share link. The list of shares for the remote endpoints is fetched instead of relying on cached values.

Syntax

```
PROCEDURE FLUSH_SHARE_LINK_CACHE
(
    link_name             IN VARCHAR2,
    owner                 IN VARCHAR2 := NULL,
    auto_commit           IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
link_name	The name of the share link.
owner	The owner of the share link. Defaults to the current schema.
auto_commit	If TRUE, the changes are automatically committed. The default is TRUE.

FLUSH_SHARE_PROVIDER_CACHE Procedure

Flush the cache of shares for a given share provider. The list of shares for the remote endpoints is fetched instead of relying on cached values.

Syntax

```
PROCEDURE FLUSH_SHARE_PROVIDER_CACHE
(
  provider_name         IN VARCHAR2,
  owner                 IN VARCHAR2 := NULL,
  auto_commit           IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
provider_name	The name of the share provider.

Parameter	Description
owner	The owner of the share provider. Defaults to the current schema.
auto_commit	If TRUE, the changes are automatically committed. The default is TRUE.

GENERATE_SHARE_LINK_SELECT Procedure and Function

Generate a SELECT statement that returns data from a shared table.

Syntax

Procedure version of GENERATE_SHARE_LINK_SELECT.

```
PROCEDURE GENERATE_SHARE_LINK_SELECT
(
    share_link_name      IN VARCHAR2,
    share_schema_name    IN VARCHAR2,
    share_table_name     IN VARCHAR2,
    stmt                 IN OUT NOCOPY CLOB,
    share_link_owner     IN VARCHAR2 := NULL
);
```

Function version of GENERATE_SHARE_LINK_SELECT.

```
 FUNCTION GENERATE_SHARE_LINK_SELECT
(
    share_link_name      IN VARCHAR2,
    share_schema_name    IN VARCHAR2,
    share_table_name     IN VARCHAR2,
    share_link_owner     IN VARCHAR2 := NULL
)
RETURN CLOB;
```

Parameters

Parameter	Description
share_link_name	The name of the share link.
share_schema_name	The name of the shared schema.
share_table_name	The name of the shared table.
stmt	The generated select statement.
share_link_owner	The link owner. Defaults to the current schema.

GET_ORACLE_SHARE_LINK_INFO Function

Retrieve the cloud link name and namespace for an Oracle-to-Oracle share.

Syntax

```
FUNCTION GET_ORACLE_SHARE_LINK_INFO
(
  oracle_provider_id   IN VARCHAR2,
```

```
  share_name           IN VARCHAR2,
  share_schema_name    IN VARCHAR2,
  share_table_name     IN VARCHAR2
)
RETURN CLOB;
```

Parameters

Parameter	Description
oracle_provider_id	The Oracle Provider ID from ALL_AVAILABLE_ORACLE_SHARE_PROVIDERS.
share_name	The name of a share from ALL_AVAILABLE_ORACLE_SHARE_PROVIDERS.SHARES.
share_schema_name	The name of a share schema from ALL_AVAILABLE_ORACLE_SHARE_PROVIDERS.SHARES.
share_table_name	The name of a table from ALL_AVAILABLE_ORACLE_SHARE_PROVIDERS.SHARES.

Return

The return value is a JSON object containing three properties: schema, table, and dblink. The caller can use these three properties to fetch data using a query of the following form:

```
SELECT *
FROM <schema>.<table>@<dblink>
```

GET_SHARE_LINK_INFO Procedure

Get the endpoint(s), share type and share name along with any additional JSON metadata for a share link.

Syntax

```
PROCEDURE GET_SHARE_LINK_INFO
(
    link_name            IN VARCHAR2,
    endpoint             IN OUT NOCOPY VARCHAR2,
    share_type           IN OUT NOCOPY VARCHAR2,
    share_name           IN OUT NOCOPY VARCHAR2,
    token_endpoint       IN OUT NOCOPY VARCHAR2,
    metadata             IN OUT NOCOPY BLOB,
    link_owner           IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
link_name	The name of a share link (previously created by CREATE_SHARE_LINK).
endpoint	OUT: The endpoint, if any, associated with the share link.
share_type	OUT: The type of share link (for example, DELTA).
share_name	OUT: The name of the linked share.

Parameter	Description
token_endpoint	This parameter is no longer used.
metadata	OUT: Optional metadata associated with the share link.
link_owner	The owner of the share link. Defaults to the current schema.

GET_SHARE_PROVIDER_CREDENTIAL Procedure

Get the credential name to be used by the current user when it attempts to access the given delta share provider.

Syntax

```
PROCEDURE GET_SHARE_PROVIDER_CREDENTIAL
(
    provider_name        IN VARCHAR2,
    share_credential     IN OUT NOCOPY VARCHAR2,
    get_token_credential IN OUT NOCOPY VARCHAR2,
    owner                IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
provider_name	The name of the share provider.
share_credential	OUT: The name of the credential associated with the provider. The credential name is returned without double quotes, as it would appear in the CREDENTIAL_NAME column of the USER_CREDENTIALS view. See ALL_CREDENTIALS View.
get_token_credential	This parameter is not used.
owner	The owner is the name of the schema where the share provider was registered, not the owner of the credential. Defaults to the current schema.

GET_SHARE_PROVIDER_INFO Procedure

Get the endpoint string(s) and share type along with any additional JSON metadata for a share provider. For ORACLE share providers, the Oracle provider ID is returned in the endpoint argument.

Syntax

```
PROCEDURE GET_SHARE_PROVIDER_INFO
(
    provider_name        IN VARCHAR2,
    endpoint             IN OUT NOCOPY VARCHAR2,
    share_type           IN OUT NOCOPY VARCHAR2,
    token_endpoint       IN OUT NOCOPY VARCHAR2,
    metadata             IN OUT NOCOPY BLOB,
    owner                IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
provider_name	The name of the share provider.
endpoint	The delta endpoint.
share_type	The share type: DELTA or ORACLE.
token_endpoint	This parameter is not used.
metadata	The optional metadata that was associated with the share provider.
owner	The owner of the share provider. Defaults to the current schema.

GET_SHARING_ID Function

Return an identifier that can be used as the sharing_id in the CREATE_SHARE_RECIPIENT procedure. This function can be used to share data between two users, the "provider" and the "recipient", in different databases.

See CREATE_SHARE_RECIPIENT Procedure for further information.

Syntax

```
FUNCTION GET_SHARING_ID
(
    sharing_id_type      IN VARCHAR2 := SHARING_ID_TYPE_DATABASE
)
  RETURN VARCHAR2;
```

Parameters

Parameter	Description
sharing_id_type	The type of sharing ID.

Usage

The flow is as follows:

1. The recipient calls DBMS_SHARE.GET_SHARING_ID to get a unique identifier.
2. The recipient sends this identifier (e.g. via email) to the provider.
3. The provider calls DBMS_SHARE.CREATE_SHARE_RECIPIENT, passing in the identifier as sharing_id.
4. The provider calls DBMS_SHARE.GRANT_TO_RECIPIENT to give the recipient access to shared data.

The sharing_id_type parameter is used to specify which database users can access the share following the above sequence.

- DATABASE The share will be visible to any admin user in the database where GET_SHARING_ID was called.
- COMPARTMENT The share will be visible to any admin user in any database in the same compartment where GET_SHARING_ID was called.

- TENANCY The share will be visible to any admin user in any database in the same tenancy where `GET_SHARING_ID` was called.
- REGION The share will be visible to any admin user in any database in the same region where `GET_SHARING_ID` was called.

OPEN_SHARE_LINK_CURSOR Procedure

Open a cursor that returns data from a shared table.

Syntax

```
PROCEDURE OPEN_SHARE_LINK_CURSOR
(
    share_link_name      IN VARCHAR2,
    share_schema_name    IN VARCHAR2,
    share_table_name     IN VARCHAR2,
    table_cursor         IN OUT NOCOPY SYS_REFCURSOR,
    share_link_owner     IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
`share_link_name`	The name of the share link.
`share_schema_name`	The name of the shared schema.
`share_table_name`	The name of the shared table.
`table_cursor`	The cursor.
`share_link_owner`	The link owner. Defaults to current schema.

REFRESH_BEARER_TOKEN_CREDENTIAL Procedure

Update one or more credentials created by `CREATE_BEARER_TOKEN_CREDENTIAL` or `CREATE_CREDENTIALS` by calling the registered token endpoints and fetching new bearer tokens. Note this procedure is called automatically by a scheduler job, `ADP$BEARER_REFRESH_JOB`, that runs every 50 minutes.

See CREATE_BEARER_TOKEN_CREDENTIAL Procedure or CREATE_CREDENTIALS Procedure and Function for further information.

Syntax

```
PROCEDURE REFRESH_BEARER_TOKEN_CREDENTIAL
(
    credential_name           IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
`credential_name`	The name of the credential to refresh.

RENAME_CLOUD_STORAGE_LINK Procedure

Rename a registered cloud storage link.

Syntax

```
PROCEDURE RENAME_CLOUD_STORAGE_LINK
(
    old_name             IN VARCHAR2,
    new_name             IN VARCHAR2,
    owner                IN VARCHAR2 := NULL,
    auto_commit          IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
old_name	The current name of the cloud storage link.
new_name	The new name of the cloud storage link.
owner	The owner of the cloud storage link. Defaults to the current schema.
auto_commit	If TRUE, the changes are automatically committed. The default is TRUE.

RENAME_SHARE_PROVIDER Procedure

Rename a registered share provider.

Syntax

```
PROCEDURE RENAME_SHARE_PROVIDER
(
  old_name             IN VARCHAR2,
  new_name             IN VARCHAR2,
  owner                IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
old_name	The current name of the share provider.
new_name	The new name of the share provider.
owner	The owner of the share provider. Defaults to the current schema.

REMOVE_SHARE_SCHEMA Procedure

Remove a schema and all its contents from a share.

Syntax

```
PROCEDURE REMOVE_SHARE_SCHEMA
(
```

```
    share_name           IN VARCHAR2,
    schema_name          IN VARCHAR2,
    share_owner          IN VARCHAR2 := NULL,
    auto_commit          IN BOOLEAN := FALSE
);
```

Parameters

Parameter	Description
share_name	The name of the share.
schema_name	The name of the schema to remove.
share_owner	The owner of the share.
auto_commit	If TRUE, the changes are automatically committed. The default is FALSE.

SET_SHARE_LINK_METADATA Procedure

Set the additional JSON metadata for the share link.

Syntax

```
PROCEDURE SET_SHARE_LINK_METADATA
(
    link_name            IN VARCHAR2,
    metadata             IN SYS.JSON_OBJECT_T,
    replace_existing     IN BOOLEAN := FALSE,
    link_owner           IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
link_name	The name of the share link.
metadata	The new metadata.
replace_existing	If TRUE, then all existing metadata is replaced by the new version. If FALSE, then the new value is merged with any existing metadata.
link_owner	The owner of the share link. Defaults to the current schema.

SET_SHARE_PROVIDER_CREDENTIAL Procedure

Set the credential name to be used by the current user when it attempts to access the given share provider.

Syntax

```
PROCEDURE SET_SHARE_PROVIDER_CREDENTIAL
(
    provider_name        IN VARCHAR2,
    share_credential     IN VARCHAR2,
    get_token_credential IN VARCHAR2 := NULL,
    owner                IN VARCHAR2 := NULL,
```

```
    check_if_exists      IN BOOLEAN := TRUE
);
```

Parameters

Parameter	Description
provider_name	The name of the share provider.
share_credential	The bearer token credential.
get_token_credential	This argument is ignored.
owner	The owner of the share provider. Defaults to the current schema.
check_if_exists	If TRUE (default), then validate that the credential exists.

SET_SHARE_PROVIDER_METADATA Procedure

Set additional JSON metadata for the share provider.

Syntax

```
PROCEDURE SET_SHARE_PROVIDER_METADATA
(
    provider_name        IN VARCHAR2,
    metadata             IN SYS.JSON_OBJECT_T,
    replace_existing     IN BOOLEAN := FALSE,
    owner                IN VARCHAR2 := NULL
);
```

Parameters

Parameter	Description
provider_name	The name of the share provider.
metadata	The new metadata.
replace_existing	If TRUE, then all existing metadata is replaced by the new version. If FALSE, the new value is merged with any existing metadata.
owner	The owner of the share provider. Defaults to the current schema.

UPDATE_BEARER_TOKEN_CREDENTIAL Procedure

Modify an attribute of a credential created by `CREATE_CREDENTIALS` or `CREATE_BEARER_TOKEN_CREDENTIAL`.

See CREATE_BEARER_TOKEN_CREDENTIAL Procedure or CREATE_CREDENTIALS Procedure and Function for further information.

Syntax

```
PROCEDURE UPDATE_BEARER_TOKEN_CREDENTIAL
(
     credential_name           IN VARCHAR2,
     attribute                 IN VARCHAR2,
     new_value                 IN VARCHAR2
);
```

Parameters

Parameter	Description
credential_name	The name of the credential to update.
attribute	The attribute to update. One of 'BEARER_TOKEN', 'CLIENT_ID', 'CLIENT_SECRET', 'TOKEN_REFRESH_RATE'. The token endpoint can not be changed.
new_value	The new value.

Example: Update the CLIENT_ID of a credential

```
SQL> BEGIN
  2    dbms_share.create_bearer_token_credential(
  3      credential_name=>'MY_RENEWABLE_CREDENTIAL',
  4      token_endpoint=>'https://myserver/ords/share_provider/oauth/token',
  5      client_id=>'VXGQ_44s6qJ-K4WHUNM2yQ..',
  6      client_secret=>'y9ddppgwEmZl7adDHFQndw..');
  7  END;
  8  /
PL/SQL procedure successfully completed.
SQL> select credential_name, username from user_credentials where
credential_name LIKE 'MY_RENEWABLE_CREDENTIAL%';
CREDENTIAL_NAME
----------------------------------------
USERNAME
------------------------------------
MY_RENEWABLE_CREDENTIAL
BEARER_TOKEN
MY_RENEWABLE_CREDENTIAL$TOKEN_REFRESH_CRED
ABCDEF
```

Summary of Share Producer Views

This table lists the share producer views for the DBMS_SHARE package.

View	Description
ALL_OBJECT_SHAREABILITY View	This view gives information about the shareability, or otherwise, of tables and views visible to the user.
ALL_SHAREABLE_OBJECTS View	This view lists all of the tables or views that can be added to a share.
ALL_SHARES View	This view lists all share created using CREATE_SHARE.
ALL_SHARE_DEFAULT_SETTINGS View	A view containing the default settings for shares and recipients.
ALL_SHARE_EVENTS View	This view lists all log events associated with shares.
ALL_SHARE_FILES View	A view that contains a list of all the parquet files generated by shares.
ALL_SHARE_RECIPIENTS View	A view containing all share recipients created using CREATE_SHARE_RECIPIENT.
ALL_SHARE_RECIPIENT_EVENTS View	This view contains records of endpoint request by delta share consumers.

View	Description
ALL_SHARE_RECIPIENT_GRANTS View	This view contains one row for each share/recipient combination as specified by calls to GRANT_TO_RECIPIENT.
ALL_SHARE_SCHEMAS View	This view lists all the share schemas that were added to shares using ADD_TO_SHARE.
ALL_SHARE_TABLES View	This view lists all tables or views that were added to share using ADD_TO_SHARE.
ALL_SHARE_VERSIONS View	A view listing the different versions of shares.
DBA_SHARE_DETACHED_FILES View	This view contains a list of parquet files that we generated by DBMS_SHARE but that have become detached in some way.
DBA_SHARE_EVENTS View	This view contains records of endpoint request by delta share consumers.
DBA_SHARE_FILES View	A view that contains a list of all the parquet files generated by shares.
DBA_SHARE_JOB_SEGMENTS View	This view contains a list of parquet files that we generated by DBMS_SHARE but that have become detached in some way.
DBA_SHARE_JOBS View	List share jobs associated with the current user.
DBA_SHARE_RECIPIENT_EVENTS View	This view contains records of endpoint request by delta share consumers.
DBA_SHARE_RECIPIENT_GRANTS View	This view contains one row for each share/recipient combination as specified by calls to GRANT_TO_RECIPIENT.
DBA_SHARE_RECIPIENTS View	A view containing all share recipients created using CREATE_SHARE_RECIPIENT.
DBA_SHARE_SCHEMAS View	This view lists all the share schemas that were added to shares using ADD_TO_SHARE.
DBA_SHARE_TABLES View	This view lists all tables or views that were added to share using ADD_TO_SHARE.
DBA_SHARE_VERSIONS View	A view listing the different versions of shares.
DBA_SHARES View	This view lists all share created using DBMS_SHARE.CREATE_SHARE.
USER_SHARE_DETACHED_FILES View	This view contains a list of parquet files that we generated by DBMS_SHARE but that have become detached in some way.
USER_SHARE_EVENTS View	This view contains records of endpoint request by delta share consumers.
USER_SHARE_FILES View	A view that contains a list of all the parquet files generated by shares.
USER_SHARE_JOBS View	List share jobs associated with the current user.
USER_SHARE_JOB_SEGMENTS View	List segments (tables, table partitions, or table sub-partitions) that are being actively processed by running share jobs.
USER_SHARE_RECIPIENT_EVENTS View	This view contains records of endpoint request by delta share consumers.
USER_SHARE_RECIPIENT_GRANTS View	This view contains one row for each share/recipient combination as specified by calls to GRANT_TO_RECIPIENT.
USER_SHARE_RECIPIENTS View	A view containing all share recipients created using CREATE_SHARE_RECIPIENT.

View	Description
USER_SHARE_SCHEMAS View	This view lists all the share schemas that were added to shares using ADD_TO_SHARE.
USER_SHARE_TABLES View	This view lists all tables or views that were added to share using ADD_TO_SHARE.
USER_SHARE_VERSIONS View	A view listing the different versions of shares.
USER_SHAREABLE_OBJECTS View	This view lists all of the tables or views that can be added to a share.
USER_SHARES View	This view lists all share created using DBMS_SHARE.CREATE_SHARE.

- ALL_OBJECT_SHAREABILITY View
  A view containing all share links that were created using CREATE_SHARE_LINK.
- ALL_SHAREABLE_OBJECTS View
  This view lists all of the tables or views that can be added to a share. Objects that are not shareable may be listed in ALL_OBJECT_SHAREABILITY along with a reason explaining why they are excluded.
- ALL_SHARES View
  This view lists all share created using DBMS_SHARE.CREATE_SHARE.
- ALL_SHARE_DEFAULT_SETTINGS View
  A view containing the default settings for shares and recipients.
- ALL_SHARE_EVENTS View
  This view lists all log events associated with shares.
- ALL_SHARE_FILES View
  A view that contains a list of all the parquet files generated by shares.
- ALL_SHARE_RECIPIENTS View
  A view containing all share recipients created using CREATE_SHARE_RECIPIENT.
- ALL_SHARE_RECIPIENT_EVENTS View
  This view contains records of endpoint request by delta share consumers.
- ALL_SHARE_RECIPIENT_GRANTS View
  This view contains one row for each share/recipient combination as specified by calls to GRANT_TO_RECIPIENT.
- ALL_SHARE_SCHEMAS View
  This view lists all the share schemas that were added to shares using ADD_TO_SHARE.
- ALL_SHARE_TABLES View
  This view lists all tables or views that were added to share using ADD_TO_SHARE.
- ALL_SHARE_VERSIONS View
  A view listing the different versions of shares. For VERSIONED shares, each row represents a particular set of parquet files. For LIVE shares, each row represents a different list of shared objects.
- DBA_SHARE_DETACHED_FILES View
  This view contains a list of parquet files that were generated by DBMS_SHARE but that have become detached in some way.
- DBA_SHARE_EVENTS View
  This view lists all log events associated with shares. Its columns are the same as those in the ALL_SHARE_EVENTS view.

- DBA_SHARE_FILES View
  A view that contains a list of all the parquet files generated by shares. Its columns are the same as those in the ALL_SHARE_FILES view.
- DBA_SHARE_JOB_SEGMENTS View
  List segments (tables, table partitions, or table sub-partitions) that are being actively processed by running share jobs. Its columns are the same as those in the USER_SHARE_JOB_SEGMENTS view.
- DBA_SHARE_JOBS View
  List share jobs associated with the current user. Its columns are the same as those in the USER_SHARE_JOBS view with the addition of an OWNER column.
- DBA_SHARE_RECIPIENT_EVENTS View
  This view contains records of endpoint request by delta share consumers. Its columns are the same as those in the ALL_SHARE_RECIPIENT_EVENTS view.
- DBA_SHARE_RECIPIENT_GRANTS View
  This view contains one row for each share/recipient combination as specified by calls to GRANT_TO_RECIPIENT. Its columns are the same as those in the ALL_SHARE_RECIPIENT_GRANTS view.
- DBA_SHARE_RECIPIENTS View
  A view containing all share recipients created using CREATE_SHARE_RECIPIENT. Its columns are the same as those in the ALL_SHARE_RECIPIENTS view.
- DBA_SHARE_SCHEMAS View
  This view lists all the share schemas that were added to shares using ADD_TO_SHARE. Its columns are the same as those in the ALL_SHARE_SCHEMAS view.
- DBA_SHARE_TABLES View
  This view lists all tables or views that were added to share using ADD_TO_SHARE. Its columns are the same as those in the ALL_SHARE_TABLES view.
- DBA_SHARE_VERSIONS View
  A view listing the different versions of shares. For VERSIONED shares, each row represents a particular set of parquet files. For LIVE shares, each row represents a different list of shared objects. Its columns are the same as those in the ALL_SHARE_VERSIONS view.
- DBA_SHARES View
  This view lists all share created using DBMS_SHARE.CREATE_SHARE. Its columns are the same as those in the ALL_SHARES view.
- USER_SHARE_DETACHED_FILES View
  This view contains a list of parquet files that were generated by DBMS_SHARE but that have become detached in some way.
- USER_SHARE_EVENTS View
  This view lists all log events associated with shares. Its columns are the same as those in the ALL_SHARE_EVENTS view.
- USER_SHARE_FILES View
  A view that contains a list of all the parquet files generated by shares. Its columns are the same as those in the ALL_SHARE_FILES view.
- USER_SHARE_JOBS View
  List share jobs associated with the current user.
- USER_SHARE_JOB_SEGMENTS View
  List segments (tables, table partitions, or table sub-partitions) that are being actively processed by running share jobs.

- USER_SHARE_RECIPIENT_EVENTS View
  This view contains records of endpoint request by delta share consumers. Its columns are the same as those in the ALL_SHARE_RECIPIENT_EVENTS view.
- USER_SHARE_RECIPIENT_GRANTS View
  This view contains one row for each share/recipient combination as specified by calls to GRANT_TO_RECIPIENT. Its columns are the same as those in the ALL_SHARE_RECIPIENT_GRANTS view.
- USER_SHARE_RECIPIENTS View
  A view containing all share recipients created using CREATE_SHARE_RECIPIENT. Its columns are the same as those in the ALL_SHARE_RECIPIENTS view.
- USER_SHARE_SCHEMAS View
  This view lists all the share schemas that were added to shares using ADD_TO_SHARE. Its columns are the same as those in the ALL_SHARE_SCHEMAS view.
- USER_SHARE_TABLES View
  This view lists all tables or views that were added to share using ADD_TO_SHARE. Its columns are the same as those in the ALL_SHARE_SCHEMAS view.
- USER_SHARE_VERSIONS View
  A view listing the different versions of shares. For VERSIONED shares, each row represents a particular set of parquet files. For LIVE shares, each row represents a different list of shared objects. Its columns are the same as those in the ALL_SHARE_VERSIONS view.
- USER_SHAREABLE_OBJECTS View
  This view lists all of the tables or views that can be added to a share. Objects that are not shareable may be listed in ALL_OBJECT_SHAREABILITY along with a reason explaining why they are excluded. Its columns are the same as those in the ALL_SHAREABLE_OBJECTS view.
- USER_SHARES View
  This view lists all share created using DBMS_SHARE.CREATE_SHARE.

ALL_OBJECT_SHAREABILITY View

A view containing all share links that were created using CREATE_SHARE_LINK.

See CREATE_SHARE_LINK Procedure for further information.

Column	Datatype	Description
OWNER	VARCHAR2	The owner of the object.
OBJECT_NAME	VARCHAR2	The name of the object.
OBJECT_TYPE	VARCHAR2	The type of the object, TABLE or VIEW.
NUM_ROWS	NUMBER	The number of rows, if this is a TABLE and if the statistics are available.
PRIV_CODE	VARCHAR2	A non NULL value, if you do not have the required privileges to share this object. • NO_GRANT_READ_PRIV: You have READ privilege, but not READ WITH GRANT OPTION. • PRIV_ERROR: This is an Oracle-maintained object or is otherwise unsupported

Column	Datatype	Description
SUPPORT_CODE	VARCHAR2	A non NULL value, if this object cannot be shared. • INVALID: The object is INVALID (e.g. a badly defined VIEW). • SHARDED: The table is sharded. • NESTED: This is a nested table. • EXTERNAL_PARTITIONED: This is an partitioned, external table. • EXTERNAL: This is a non-partitioned external table. • INDEX_ORGANIZED_TABLE: This is an index organized table. • TRACKING_MV: This is a tracking MV created by DBMS_SHARE. • NO_SHAREABLE_COLUMNS: None of the columns in the table or view can be shared.
SHARE_METADATA	BLOB	A JSON document that gives details of which columns are shareable and how they would be represented using parquet.

ALL_SHAREABLE_OBJECTS View

This view lists all of the tables or views that can be added to a share. Objects that are not shareable may be listed in ALL_OBJECT_SHAREABILITY along with a reason explaining why they are excluded.

Column	Datatype	Description
OWNER	VARCHAR2	The owner of the object.
OBJECT_NAME	VARCHAR2	The name of the object.
OBJECT_TYPE	VARCHAR2	The type of the object, TABLE or VIEW.
NUM_ROWS	NUMBER	The number of rows, if this is a TABLE and if the statistics are available.
SHARE_METADATA	BLOB	A JSON document that gives details of which columns are shareable and how they would be represented using parquet.
NUM_COLUMNS	VARCHAR2	The number of shareable columns in the object.

ALL_SHARES View

This view lists all share created using DBMS_SHARE.CREATE_SHARE.

See CREATE_SHARE Procedure for further information.

Column	Datatype	Description
OWNER	VARCHAR2	The owner of the share.
SHARE_NAME	VARCHAR2	The name of the share.
SHARE_ID	NUMBER	A numeric ID for this share.
METADATA_PATH	VARCHAR2	The lineage metadata path for this share.
METADATA	BLOB	A JSON document containing additional share metadata.
DESCRIPTION	CLOB	A local description of the share.

Column	Datatype	Description
SHARE_TYPE	VARCHAR2	The type is either VERSIONED or LIVE.
PUBLIC_DESCRIPTION	CLOB	An externally visible description of the share.
NUM_TABLES	NUMBER	The number of tables in the share.
NUM_RECIPIENTS	NUMBER	The number of recipients who have been granted access to the share.
NUM_ORACLE_RECIPIENTS	NUMBER	The number of Oracle recipients who have been granted access to the share.
STORAGE_FOLDER	VARCHAR2	The base URL for any generated parquet files.
ORACLE_SHARE_ID	VARCHAR2	A globally unique UUID for the share.
LAST_EXPORT_TIME	TIMESTAMP(9) WITH TIME ZONE	(VERSIONED ONLY) The time when the currently published version was generated.
LAST_NOTIFICATION	TIMESTAMP(9) WITH TIME ZONE	The time the last email notification was sent about this share..
CURRENT_VERSION	NUMBER	The currently published version number.
CREATED	TIMESTAMP(9) WITH TIME ZONE	The date this share was created.
UPDATED	TIMESTAMP(9) WITH TIME ZONE	The date this share was last modified.

ALL_SHARE_DEFAULT_SETTINGS View

A view containing the default settings for shares and recipients.

These can be modified (by ADMIN) using `UPDATE_DEFAULT_SHARE_PROPERTY` and `UPDATE_DEFAULT_RECIPIENT_PROPERTY`. Note that there is no DBA or USER version of this view.

Column	Datatype	Description
SETTING	VARCHAR2	The setting name: • JOB_CLASS • JOB_DOP • MAX_PAR_LIFETIME • MAX_TOKEN_LIFETIME • MIN_PAR_LIFETIME • PAR_LIFETIME • SPLIT_SIZE • TOKEN_LIFETIME
SETTING_VALUE	VARCHAR2	The current setting value.

ALL_SHARE_EVENTS View

This view lists all log events associated with shares.

Entries can be removed by deleting the associated share or by calling the CLEAR_SHARE_EVENTS Procedure.

Column	Datatype	Description
OWNER	VARCHAR2	The owner of the share.

Column	Datatype	Description
SHARE_NAME	VARCHAR2	The name of the share.
SHARE_ID	NUMBER	A numeric ID for the share.
STATUS	VARCHAR2	An event status, such as INFO or ERROR.
DETAILS	CLOB	Additional data for the event, typically a JSON document.
TIME	TIMESTAMP(6) WITH TIME ZONE	The time of the event.
LOG_LEVEL	NUMBER	The detail level for the event: 1-3
SHARE_JOB_ID	NUMBER	The ID of the share job (from USER_SHARE_JOBS View) that created this event.
SCHEDULER_JOB_NAME	VARCHAR2	The name of the dbms_scheduler job that created this event.
ACTION	VARCHAR2	A code for the event activity.
SHARE_VERSION	NUMBER	The share version number associated with this event, if relevant.
OBJECT_NAME	VARCHAR2	The name of the table or view being processed, if relevant.
OBJECT_OWNER	VARCHAR2	The owner of the table or view name being processed, if relevant.
PARTITION_NAME	VARCHAR2	The partition of the table being processed, if relevant.
SPLIT_ORDER	NUMBER	The split number, if a single partition is broken into splits.
SQL_STATEMENT	CLOB	A SQL statement that was executed as part of the event.

ALL_SHARE_FILES View

A view that contains a list of all the parquet files generated by shares.

Column	Datatype	Description
OWNER	VARCHAR2	The owner of the share.
SHARE_NAME	VARCHAR2	The name of the share.
SHARE_ID	NUMBER	A numeric ID for the share.
SHARE_TABLE_NAME	VARCHAR2	The share table associated with the file.
SHARE_SCHEMA_NAME	VARCHAR2	The share schema associated with the file.
SHARE_TABLE_ID	NUMBER	The numeric ID of the share table (ALL_SHARE_TABLES).
SEGMENT_NAME	VARCHAR2	The associated segment name (e.g. a partition), if appropriate.
INITIAL_SHARE_VERSION	NUMBER	The lowest share version number that requires this file.
EXPIRED_SHARE_VERSION	NUMBER	The share version number, if any, where this file was removed.
FLE_URI	VARCHAR2	The full URI for the file.
FILE_NAME	VARCHAR2	The name of the file within the cloud storage bucket.
BYTES	NUMBER	The number of bytes in the file.
SPLIT_ORDER	NUMBER	The split number within the segment.
FILE_STATS	BLOB	Delta statistics (min/max) for columns in the file.

ALL_SHARE_RECIPIENTS View

A view containing all share recipients created using `CREATE_SHARE_RECIPIENT`.

See CREATE_SHARE_RECIPIENT Procedure for further information.

Name	Datatype	Description
OWNER	VARCHAR2	The owner of the recipient object.
RECIPIENT_NAME	VARCHAR2	The name of the recipient.
RECIPIENT_ID	NUMBER	A numeric ID for the recipient
METADATA_PATH	VARCHAR2	The lineage path for the recipient.
RECIPIENT_TYPE	VARCHAR2	The type of recipient • DELTA_SHARING: recipient defined by an email • ORACLE: defined by an Oracle sharing ID • ORACLE_DELTA: defined by an email and a sharing ID
INFO	BLOB	Not currently used.
METADATA	BLOB	JSON metadata associated with this recipient.
DESCRIPTION	CLOB	A local description of the recipient.
RECIPIENT_SHARING_ID	VARCHAR2	(ORACLE) An optional sharing ID for the recipient.
MIN_PAR_INTERVAL	INTERVAL DAY(3) TO SECOND(0)	(DELTA) The shortest lifetime of pre-authenticated requests that will be sent to the consumer.
MAX_PAR_INTERVAL	INTERVAL DAY(3) TO SECOND(0)	(DELTA) The longest lifetime of pre-authenticated requests that will be sent to the consumer.
LOG_LEVEL	NUMBER	The recipient event log level: 0 - 3
CREATED	TIMESTAMP(6) WITH TIME ZONE	The date this share recipient was created.
UPDATED	TIMESTAMP(6) WITH TIME ZONE	The date this share recipient was last modified.

ALL_SHARE_RECIPIENT_EVENTS View

This view contains records of endpoint request by delta share consumers.

Name	Type	Description
OWNER	VARCHAR2	The owner of the recipient.
RECIPIENT_NAME	VARCHAR2	The name of the recipient.
RECIPIENT_ID	NUMBER	The numeric ID of the recipient.
SHARE_ID	NUMBER	The ID of the share being accessed.
STATUS	VARCHAR2	A status code, such as INFO or ERROR.
ENDPOINT	VARCHAR2	The delta endpoint being accessed.
REMOTE_ADDRESS	VARCHAR2	The external address that called the endpoint.
HTTP_USER_AGENT	VARCHAR2	The requesting user-agent, from the HTTP request.
HTTP_REFERER	VARCHAR2	The referer from the HTTP request.
REQUEST	CLOB	The request body. This is only logged if the recipient event level is high.
ERROR_MESSAGE	VARCHAR2	Any error message, if the request failed.

Name	Type	Description
ERROR_TRACE	VARCHAR2	Any error trace, if the request failed.
ELAPSED	INTERVAL DAY(9) TO SECOND(9)	The elapsed time taken to process the request.
CPU_TIME	NUMBER	The CPU time take to process the request
DETAILS	CLOB	The response details. This is only logged if the recipient event level is high.
TIME	TIMESTAMP(6) WITH TIME ZONE	The time the event was recorded.
LOG_LEVEL	NUMBER	The detail level of this event: 1-3

ALL_SHARE_RECIPIENT_GRANTS View

This view contains one row for each share/recipient combination as specified by calls to GRANT_TO_RECIPIENT.

See GRANT_TO_RECIPIENT Procedure for further information.

Name	Datatype	Description
OWNER	VARCHAR2	The owner of the share and recipient.
RECIPIENT_NAME	VARCHAR2	The name of the recipient.
RECIPIENT_ID	NUMBER	The numeric ID of the recipient.
SHARE_NAME	VARCHAR2	The name of the share.
SHARE_ID	NUMBER	The numeric ID of the share.
ORDER_NUM	NUMBER	The grant order.
CREATED	TIMESTAMP(6) WITH TIME ZONE	The time when the grant was made.
UPDATED	TIMESTAMP(6) WITH TIME ZONE	The last time when this grant was updated.

ALL_SHARE_SCHEMAS View

This view lists all the share schemas that were added to shares using ADD_TO_SHARE.

See ADD_TO_SHARE Procedure for further information.

Name	Datatype	Description
OWNER	VARCHAR2	The owner of the share.
SHARE_NAME	VARCHAR2	The name of the share.
SHARE_ID	NUMBER	The numeric ID for this share.
SCHEMA_NAME	VARCHAR2	The name of the share schema, as seen by the share consumers.
SCHEMA_ID	NUMBER	The numeric ID of the share schema.
METADATA_PATH	NUMBER	The lineage path for this share schema.
CREATED	TIMESTAMP(6) WITH TIME ZONE	The date this share schema was created.
UPDATED	TIMESTAMP(6) WITH TIME ZONE	The date this share schema was last modified.

ALL_SHARE_TABLES View

This view lists all tables or views that were added to share using `ADD_TO_SHARE`.

See ADD_TO_SHARE Procedure for further information.

Name	Datatype	Description
OWNER	VARCHAR2	The owner of the share.
SHARE_NAME	VARCHAR2	The name of the share.
SHARE_ID	NUMBER	A numeric ID for the share.
SHARE_SCHEMA_NAME	VARCHAR2	The name of the share schema, as seen by the share consumers.
SHARE_SCHEMA_ID	NUMBER	A numeric ID for the share schema.
SHARE_TABLE_NAME	VARCHAR2	The name of the share schema, as seen by the share consumers.
SHARE_TABLE_ID	NUMBER	A numeric ID for the share table.
METADATA_PATH	VARCHAR2	The lineage path for the share table.
TABLE_NAME	VARCHAR2	The local table or view being shared.
TABLE_OWNER	VARCHAR2	The owner of the local table or view being shared.
OBJECT_TYPE	VARCHAR2	The type of object being shared: TABLE or VIEW.
CREATED	TIMESTAMP(6) WITH TIME ZONE	The date this share table was created.
UPDATED	TIMESTAMP(6) WITH TIME ZONE	The date this share table was last modified.
STATUS	VARCHAR2	The STATUS of the shared object, from `ALL_OBJECTS`.
PRIV_CODE	VARCHAR2	See `ALL_OBJECT_SHAREABILITY.PRIV_CODE`.

ALL_SHARE_VERSIONS View

A view listing the different versions of shares. For VERSIONED shares, each row represents a particular set of parquet files. For LIVE shares, each row represents a different list of shared objects.

Name	Datatype	Description
OWNER	VARCHAR2	The owner of the share.
SHARE_NAME	VARCHAR2	The name of the share.
SHARE_ID	NUMBER	A numeric ID for the share.
SHARE_VERSION	NUMBER	The version of this share.
STATUS	VARCHAR2	The status of this version. • CURRENT • RETIRED • FAILED • ROLLED BACK • STOPPED • PENDING • EXPORTING
EXPORT_TIME	TIMESTAMP(9)	The time when this version was created.
LAST_UPDATE	TIMESTAMP(9)	The last time this version was modified.

Name	Datatype	Description
PROMISED_TO	TIMESTAMP(9)	The latest expiry time for any pre-authenticated request on a parquet file associated with this version.
EXPORT_SCN	NUMBER	The System Change Number current when this version was created.
NUM_FILES	NUMBER	The number of parquet files associated with this version.
NUM_NEW_FILES	NUMBER	The number of new parquet files created for this version.
NUM_PENDING_FILES	NUMBER	The number of parquet files that should have been created, but were not for some reason.
NUM_NEW_PENDING_FILES	NUMBER	The number of pending files that are new as of this version.
ERROR_MESSAGE	VARCHAR2	An error message in the case where the version failed.
TOTAL_FILE_SIZE	NUMBER	The total size of all parquet files associated with this version.

DBA_SHARE_DETACHED_FILES View

This view contains a list of parquet files that were generated by DBMS_SHARE but that have become detached in some way.

This could happen, for example, if a share or share version is dropped with the destroy_objects argument set to FALSE. It can also happen if a user with shares is dropped.

Name	Datatype	Description
OWNER	VARCHAR2	The name of the user who generated the file. This may be NULL if the user has been dropped from the database.
OWNER_ID	NUMBER	The USER_ID of the user who generated the file. This will remain valid and non-NULL even if the user was dropped.
FILE_URI	VARCHAR2	The URL of the generated parquet file.

See Also:

"USER_SHARE_DETACHED_FILES View"

DBA_SHARE_EVENTS View

This view lists all log events associated with shares. Its columns are the same as those in the ALL_SHARE_EVENTS view.

See Also:

"ALL_SHARE_EVENTS View"

DBA_SHARE_FILES View

A view that contains a list of all the parquet files generated by shares. Its columns are the same as those in the ALL_SHARE_FILES view.

See Also:

"ALL_SHARE_FILES View"

DBA_SHARE_JOB_SEGMENTS View

List segments (tables, table partitions, or table sub-partitions) that are being actively processed by running share jobs. Its columns are the same as those in the `USER_SHARE_JOB_SEGMENTS` view.

See Also:

"USER_SHARE_JOB_SEGMENTS View"

DBA_SHARE_JOBS View

List share jobs associated with the current user. Its columns are the same as those in the `USER_SHARE_JOBS` view with the addition of an `OWNER` column.

See Also:

"USER_SHARE_JOBS View"

DBA_SHARE_RECIPIENT_EVENTS View

This view contains records of endpoint request by delta share consumers. Its columns are the same as those in the `ALL_SHARE_RECIPIENT_EVENTS` view.

See Also:

"ALL_SHARE_RECIPIENT_EVENTS View"

DBA_SHARE_RECIPIENT_GRANTS View

This view contains one row for each share/recipient combination as specified by calls to `GRANT_TO_RECIPIENT`. Its columns are the same as those in the `ALL_SHARE_RECIPIENT_GRANTS` view.

See Also:

"ALL_SHARE_RECIPIENT_GRANTS View"

DBA_SHARE_RECIPIENTS View

A view containing all share recipients created using `CREATE_SHARE_RECIPIENT`. Its columns are the same as those in the `ALL_SHARE_RECIPIENTS` view.

See Also:

" ALL_SHARE_RECIPIENTS View"

DBA_SHARE_SCHEMAS View

This view lists all the share schemas that were added to shares using ADD_TO_SHARE. Its columns are the same as those in the `ALL_SHARE_SCHEMAS` view.

See Also:

"ALL_SHARE_SCHEMAS View"

DBA_SHARE_TABLES View

This view lists all tables or views that were added to share using `ADD_TO_SHARE`. Its columns are the same as those in the `ALL_SHARE_TABLES` view.

See Also:

"ALL_SHARE_TABLES View"

DBA_SHARE_VERSIONS View

A view listing the different versions of shares. For VERSIONED shares, each row represents a particular set of parquet files. For LIVE shares, each row represents a different list of shared objects. Its columns are the same as those in the `ALL_SHARE_VERSIONS` view.

See Also:

"ALL_SHARE_VERSIONS View"

DBA_SHARES View

This view lists all share created using `DBMS_SHARE.CREATE_SHARE`. Its columns are the same as those in the `ALL_SHARES` view.

See Also:

"ALL_SHARES View"

USER_SHARE_DETACHED_FILES View

This view contains a list of parquet files that were generated by `DBMS_SHARE` but that have become detached in some way.

This could happen, for example, if a share or share version is dropped with the `destroy_objects` argument set to FALSE. It can also happen if a user with shares is dropped. In that case the files would show up in `DBA_SHARE_DETACHED_FILES View`.

Name	Datatype	Description
`FILE_URI`	`VARCHAR2`	The URL of the generated parquet file.

See Also:

"DBA_SHARE_DETACHED_FILES View"

USER_SHARE_EVENTS View

This view lists all log events associated with shares. Its columns are the same as those in the `ALL_SHARE_EVENTS` view.

See Also:

"ALL_SHARE_EVENTS View"

USER_SHARE_FILES View

A view that contains a list of all the parquet files generated by shares. Its columns are the same as those in the `ALL_SHARE_FILES` view.

See Also:

"ALL_SHARE_FILES View"

USER_SHARE_JOBS View

List share jobs associated with the current user.

Share jobs are created whenever the database needs to run share processes in the background. The PUBLISH_SHARE Procedure initiates a share job.

Column	Datatype	Description
`JOB_ID`	`NUMBER`	A numeric ID for the job.
`JOB_NAME`	`VARCHAR2`	The name of the share job.
`SHARE_ID`	`NUMBER`	A numeric ID for the share associated with this job, if relevant.
`PARALLELISM`	`NUMBER`	The number of `DBMS_SCHEDULER` jobs used to run this share job.
`SHARE_NAME`	`VARCHAR2`	The name of the share associated with this job, if relevant.
`SHARE_VERSION`	`NUMBER`	The share version number associated with this job, if relevant.
`STATUS`	`VARCHAR2`	The status of this job: • PENDING • STOPPING • RUNNING • WAITING_FOR_ODI • DELETING • FAILED • GENERATING_PARS • COMPLETED • STOPPED • ERROR

USER_SHARE_JOB_SEGMENTS View

List segments (tables, table partitions, or table sub-partitions) that are being actively processed by running share jobs.

Name	Datatype	Description
`JOB_ID`	`NUMBER`	A numeric ID for the job from `USER_SHARE_JOBS`.
`JOB_NAME`	`VARCHAR2`	The name of the share job from `USER_SHARE_JOBS`.
`SHARE_ID`	`NUMBER`	A numeric ID for the share associated with this job, if relevant
`PARALLELISM`	`NUMBER`	The number of DBMS_SCHEDULER jobs used to run this share job.

Name	Datatype	Description
SHARE_NAME	VARCHAR2	The name of the share associated with this job, if relevant.
SHARE_VERSION	NUMBER	The share version number associated with this job, if relevant.
STATUS	VARCHAR2	The status of this job: • PENDING • STOPPING • RUNNING • WAITING_FOR_ODI • DELETING • FAILED • GENERATING_PARS • COMPLETED • STOPPED • ERROR
OBJECT_NAME	VARCHAR2	The name of the object that is being processed.
OBJECT_OWNER	VARCHAR2	The owner of the object that is being processed.
OBJECT_TYPE	VARCHAR2	The type of object (TABLE or VIEW) that is being processed.
SHARE_TABLE_NAME	VARCHAR2	The name of the object as it appears in the share.
SHARE_SCHEMA_NAME	VARCHAR2	The name of the object's schema as it appears in the share.
PARTITION_NAME	VARCHAR2	The name of the active partition or sub-partition, if appropriate.
SEGMENT_TYPE	VARCHAR2	The type of segment: TABLE, TABLE PARTITION, or TABLE SUBPARTITION.
BYTES	VARCHAR2	The size of the segment in bytes.

See Also:

"DBA_SHARE_JOB_SEGMENTS View"

USER_SHARE_RECIPIENT_EVENTS View

This view contains records of endpoint request by delta share consumers. Its columns are the same as those in the ALL_SHARE_RECIPIENT_EVENTS view.

See Also:

"ALL_SHARE_RECIPIENT_EVENTS View"

USER_SHARE_RECIPIENT_GRANTS View

This view contains one row for each share/recipient combination as specified by calls to GRANT_TO_RECIPIENT. Its columns are the same as those in the ALL_SHARE_RECIPIENT_GRANTS view.

See Also:

"ALL_SHARE_RECIPIENT_GRANTS View"

USER_SHARE_RECIPIENTS View

A view containing all share recipients created using `CREATE_SHARE_RECIPIENT`. Its columns are the same as those in the `ALL_SHARE_RECIPIENTS` view.

See Also:

" ALL_SHARE_RECIPIENTS View"

USER_SHARE_SCHEMAS View

This view lists all the share schemas that were added to shares using `ADD_TO_SHARE`. Its columns are the same as those in the `ALL_SHARE_SCHEMAS` view.

See Also:

"ALL_SHARE_SCHEMAS View"

USER_SHARE_TABLES View

This view lists all tables or views that were added to share using `ADD_TO_SHARE`. Its columns are the same as those in the `ALL_SHARE_SCHEMAS` view.

See Also:

"ALL_SHARE_TABLES View"

USER_SHARE_VERSIONS View

A view listing the different versions of shares. For VERSIONED shares, each row represents a particular set of parquet files. For LIVE shares, each row represents a different list of shared objects. Its columns are the same as those in the `ALL_SHARE_VERSIONS` view.

See Also:

"ALL_SHARE_VERSIONS View"

USER_SHAREABLE_OBJECTS View

This view lists all of the tables or views that can be added to a share. Objects that are not shareable may be listed in `ALL_OBJECT_SHAREABILITY` along with a reason explaining why they are excluded. Its columns are the same as those in the `ALL_SHAREABLE_OBJECTS` view.

See Also:

"ALL_SHAREABLE_OBJECTS View"

USER_SHARES View

This view lists all share created using `DBMS_SHARE.CREATE_SHARE`.

See Also:

"ALL_SHARES View"

Summary of Share Consumer Views

This table lists the DBMS_SHARE package views.

View	Description
ALL_AVAILABLE_ORACLE_SHARE_PROVIDERS View	This view lists all available Oracle share providers.
ALL_SHARE_CACHE_JOB_INFO View	This view contains information on the most recent Oracle Live share cache population job in the current POD.
ALL_SHARE_LINKS View	A view containing all share links that were created using CREATE_SHARE_LINK.
ALL_SHARE_LINK_VIEWS View	A view that contains a list of all the views created to access the shared data using CREATE_SHARE_LINK_VIEW.
ALL_SHARE_PROVIDERS View	A view that contains one row for each share provider subscription, created by CREATE_SHARE_PROVIDER.
DBA_SHARE_LINK_VIEWS View	A view that contains a list of all the views created to access the shared data using CREATE_SHARE_LINK_VIEW.
DBA_SHARE_LINKS View	A view containing all share links that were created using CREATE_SHARE_LINK.
USER_SHARE_LINKS_VIEW View	A view that contains a list of all the views created to access the shared data using CREATE_SHARE_LINK_VIEW.
USER_SHARE_LINKS View	A view containing all share links that were created using CREATE_SHARE_LINK.
USER_SHARE_PROVIDERS View	A view that contains one row for each share provider subscription, created by CREATE_SHARE_PROVIDER.

- ALL_AVAILABLE_ORACLE_SHARE_PROVIDERS View
  This view lists all available Oracle share providers.
- ALL_SHARE_CACHE_JOB_INFO View
  This view contains information on the most recent Oracle Live share cache population job in the database.
- ALL_SHARE_LINK_VIEWS View
  A view that contains a list of all the views created to access the shared data using CREATE_SHARE_LINK_VIEW.
- ALL_SHARE_LINKS View
  A view containing all share links that were created using CREATE_SHARE_LINK.
- ALL_SHARE_PROVIDERS View
  A view that contains one row for each share provider subscription, created by CREATE_SHARE_PROVIDER.
- DBA_SHARE_LINK_VIEWS View
  A view that contains a list of all the views created to access the shared data using CREATE_SHARE_LINK_VIEW. Its columns are the same as those in the ALL_SHARE_LINK_VIEWS view.

- DBA_SHARE_LINKS View
  A view containing all share links that were created using `CREATE_SHARE_LINK`. Its columns are the same as those in the `ALL_SHARE_LINKS` view.
- USER_SHARE_LINKS_VIEW View
  A view that contains a list of all the views created to access the shared data using `CREATE_SHARE_LINK_VIEW`. Its columns are the same as those in the `ALL_SHARE_LINK_VIEWS` view.
- USER_SHARE_LINKS View
  A view containing all share links that were created using `CREATE_SHARE_LINK`. Its columns are the same as those in the `ALL_SHARE_LINKS` view.
- USER_SHARE_PROVIDERS View
  A view that contains one row for each share provider subscription, created by `CREATE_SHARE_PROVIDER`. Its columns are the same as those in the `ALL_SHARE_PROVIDERS` view.

ALL_AVAILABLE_ORACLE_SHARE_PROVIDERS View

This view lists all available Oracle share providers.

Column	Datatype	Description
ORACLE_PROVIDER_ID	VARCHAR2	A UUID for the Oracle provider.
PROVIDER_NAME	VARCHAR2	A name, as specified by the provider. Note that this is not unique.
DESCRIPTION	VARCHAR2	A description, as specified by the provider.
CONTACT	VARCHAR2	A contact address, such as an email, for the provider.
METADATA	CLOB	Additional JSON metadata, as specified by the provider.
SHARES	CLOB	A JSON representation of the list of shares available from this provider.
CREATED	TIMESTAMP(6)	The time when the share provider was first registered.
UPDATED	TIMESTAMP(6)	The time when the share provider was last updated.

ALL_SHARE_CACHE_JOB_INFO View

This view contains information on the most recent Oracle Live share cache population job in the database.

Names	Datatype	Description
LAST_RUN_DATE	TIMESTAMP(6) WITH TIME ZONE	The start time of the most recent cache job execution.
NEXT_RUN_DATE	TIMESTAMP(6) WITH TIME ZONE	The scheduled start time of the next cache job execution.
RUN_STATUS	NUMBER	The status of the most recent run (0 = success). A non zero value represents the numeric code of the exception, as reported by the SQL CODE function.
MESSAGE_TEXT	VARCHAR2(4000)	A text indicating a warning or error condition encountered by the most recent run. NULL in most cases; indicates a warning if `RUN_STATUS` is 0, an error otherwise.
REPEAT_INTERVAL	VARCHAR2(4000)	The `DBMS_SCHEDULER` repeat interval for the cache job.

ALL_SHARE_LINK_VIEWS View

A view that contains a list of all the views created to access the shared data using `CREATE_SHARE_LINK_VIEW`.

See CREATE_SHARE_LINK_VIEW Procedure for further information.

Names	Datatype	Description
LINK_OWNER	VARCHAR2	The owner of the associated share link.
LINK_NAME	VARCHAR2	The name of the associated share link.
SHARE_SCHEMA_NAME	VARCHAR2	The name of the share schema, as supplied by the share provider.
SHARE_TABLE_NAME	VARCHAR2	The name of the share table, as supplied by the share provider.
VIEW_OWNER	VARCHAR2	The owner of the local view.
VIEW_NAME	VARCHAR2	The name of the local view.

ALL_SHARE_LINKS View

A view containing all share links that were created using `CREATE_SHARE_LINK`.

See CREATE_SHARE_LINK Procedure for further information.

Column	Datatype	Description
OWNER	VARCHAR2	The owner of the share link.
LINK_NAME	VARCHAR2	The name of the share link.
LINK_ID	NUMBER	A numeric ID for the share link.
SHARE_TYPE	VARCHAR2	The type of share provider: DELTA or ORACLE.
SHARE_NAME	VARCHAR2	The name of the share.
SHARE_PROVIDER_NAME	VARCHAR2	The name of the share provider.
SHARE_PROVIDER_OWNER	VARCHAR2	The owner of the share provider.
SHARE_PROVIDER_ID	NUMBER	A numeric ID for the provider.
METADATA_PATH	VARCHAR2	The lineage path of the share link.
SHARE_CREDENTIAL_NAME	VARCHAR2	(DELTA) The name of the credential used to access the endpoints.
METADATA	BLOB	Additional metadata, supplied by the user.
ORACLE_SHARE_PROVIDER_ID	VARCHAR2	(ORACLE) A globally unique GUID for the Oracle share provider.
CREATED	TIMESTAMP(6) WITH TIME ZONE	The date this share link was created.
UPDATED	TIMESTAMP(6) WITH TIME ZONE	The date this share link was last modified.

ALL_SHARE_PROVIDERS View

A view that contains one row for each share provider subscription, created by `CREATE_SHARE_PROVIDER`.

See CREATE_SHARE_PROVIDER Procedure for further information.

Column	Datatype	Description
OWNER	VARCHAR2	The user who subscribed to the share provider.
PROVIDER_NAME	VARCHAR2	The local name given to the share provider.
PROVIDER_ID	VARCHAR2	A numeric ID for the provider.
ENDPOINT	VARCHAR2	(DELTA) The delta endpoint for the provider.
SHARE_TYPE	VARCHAR2	The type of share provider: DELTA or ORACLE.
METADATA_PATH	VARCHAR2	The lineage path for the share provider.
SHARE_CREDENTIAL_NAME	VARCHAR2	(DELTA) The name of the credential used to access the delta endpoints.
METADATA	BLOB	Additional metadata, provided by the user.
ORACLE_PROVIDER_ID	VARCHAR2	(ORACLE) A globally unique GUID for the Oracle share provider.
CREATED	TIMESTAMP(6) WITH TIME ZONE	The date this share provider was created.
UPDATED	TIMESTAMP(6) WITH TIME ZONE	The date this share provider was last modified.

DBA_SHARE_LINK_VIEWS View

A view that contains a list of all the views created to access the shared data using CREATE_SHARE_LINK_VIEW. Its columns are the same as those in the ALL_SHARE_LINK_VIEWS view.

See Also:

"ALL_SHARE_LINK_VIEWS View"

DBA_SHARE_LINKS View

A view containing all share links that were created using CREATE_SHARE_LINK. Its columns are the same as those in the ALL_SHARE_LINKS view.

See Also:

"ALL_SHARE_LINKS View"

USER_SHARE_LINKS_VIEW View

A view that contains a list of all the views created to access the shared data using CREATE_SHARE_LINK_VIEW. Its columns are the same as those in the ALL_SHARE_LINK_VIEWS view.

See Also:

"ALL_SHARE_LINK_VIEWS View"

USER_SHARE_LINKS View

A view containing all share links that were created using `CREATE_SHARE_LINK`. Its columns are the same as those in the `ALL_SHARE_LINKS` view.

See Also:

"ALL_SHARE_LINKS View"

USER_SHARE_PROVIDERS View

A view that contains one row for each share provider subscription, created by `CREATE_SHARE_PROVIDER`. Its columns are the same as those in the `ALL_SHARE_PROVIDERS` view.

See Also:

"ALL_SHARE_PROVIDERS View"

DBMS_SHARE Constants

These constants are used by the `DBMS_SHARE` package.

Detached Files

Constants used by PURGE_DETACHED_FILES Procedure.

Name	Type	Value	Description
`PURGE_DROP`	`PLS_INTEGER`	1	Attempt to drop the files, using the specified credential. If the file cannot be dropped, then the file continues to be listed in the `*_SHARE_DETACHED_FILES` views.
`PURGE_DROP_FORCE`	`PLS_INTEGER`	2	Attempt to drop the files, using the specified credential. The file is removed from the `*_SHARE_DETACHED_FILES` views, even if the attempt to drop the file fails again.
`PURGE_FORGET`	`PLS_INTEGER`	3	Remove the files `*_SHARE_DETACHED_FILES` views without attempting to drop them.

Log Level

Constants used for setting the level of information collected in the log. These levels control the number of events that get logged into `ALL_SHARE_EVENTS` and `ALL_SHARE_RECIPIENTS`.

See SET_SHARE_LOG_LEVEL Procedure and SET_RECIPIENT_LOG_LEVEL Procedure.

Name	Type	Value	Description
`LOG_LEVEL_ERRORS_ONLY`	`PLS_INTEGER`	0	Only log errors.

Name	Type	Value	Description
LOG_LEVEL_BASIC	PLS_INTEGER	1	Log errors and basic information.
LOG_LEVEL_DETAIL	PLS_INTEGER	2	Log errors and additional details.
LOG_LEVEL_DEBUG	PLS_INTEGER	3	Log debug level of information.

Share Job Properties

Properties used in UPDATE_SHARE_JOB_PROPERTY Procedure.

Name	Type	Value	Description
PROP_SHARE_JOB_DOP	VARCHAR2(7)	'JOB_DOP'	The JOB_DOP property determines how many DBMS_SCHEDULER jobs are used to publish the share. This should be a number between 1 and 5.
PROP_SHARE_JOB_CLASS	VARCHAR2(9)	'JOB_CLASS'	The JOB_CLASS property determines the scheduler job class that is used to published the share. See ALL_SCHEDULER_JOB_CLASSES for a list of valid values. The default value is DEFAULT_JOB_CLASS.
PROP_SHARE_JOB_PRIORITY	VARCHAR2(12)	'JOB_PRIORITY'	The JOB_PRIORITY property determines the relative priority when two or more shares from the same user are being published at the same time. The value should be a number, with 0 as the default. Shares with a higher priority will be processed before shares with a lower priority. Shares with the same priority are processed on a first come, first serve basis. Standard values for job priority: SHARE_PRIORITY_LOW CONSTANT NUMBER := 1; SHARE_PRIORITY_DEFAULT CONSTANT NUMBER := 2; SHARE_PRIORITY_HIGH CONSTANT NUMBER := 3;

Share Properties

Constants used to indicate the properties of a share.

See UPDATE_SHARE_PROPERTY Procedure, UPDATE_DEFAULT_SHARE_PROPERTY Procedure and GET_SHARE_PROPERTY Function.

Name	Type	Value	Description
PROP_SHARE_DESC	VARCHAR2(11)	'DESCRIPTION'	A locally visible description of the share.
PROP_SHARE_PUBLIC_DESC	VARCHAR2(18)	'PUBLIC_DESCRIPTION'	An externally visible description of the share.
PROP_SHARE_SPLIT_SIZE	VARCHAR2(10)	'SPLIT_SIZE'	The SPLIT_SIZE property determine how large the generated parquet files will be. The algorithm works as follows: 1. The SQL Optimizer estimates the size in bytes of a given segment (for example, a partition of a table or the results of a view). 2. The number of bytes is divided by the SPLIT_SIZE and rounded up to produce the "number of splits". 3. The data in the segment is then chunked into this "number of splits" using some method (for example ROWID, index columns, specified columns, etc.). 4. Each chunk of data is then converted into a parquet file. Note that the parquet file will generally be smaller than the SPLIT_SIZE due to the compression. In particular, SPLIT_SIZE determines the amount of data before compression.
PROP_SHARE_LOG_LEVEL	VARCHAR2(9)	'LOG_LEVEL'	The LOG_LEVEL property determines the amount of information that will be logged in the ALL_SHARE_EVENTS log.
PROP_SHARE_VERSION_ACCESS	VARCHAR2(14)	'VERSION_ACCESS'	The versions of a published share a recipient can see.

Share Recipient PAR Type

These constants are the valid values for the property `PROP_RECIPIENT_PAR_TYPE` in UPDATE_RECIPIENT_PROPERTY Procedure and GET_RECIPIENT_PROPERTY.

Name	Type	Value	Description
PAR_TYPE_FOLDER	PLS_INTEGER	1	PARs gives access to the entire folder.
PAR_TYPE_FILE	PLS_INTEGER	2	PARs gives access to one file at a time.

Share Recipient Properties

These constants are used in GET_SHARE_PROPERTY Function, UPDATE_DEFAULT_SHARE_PROPERTY Procedure, and UPDATE_SHARE_PROPERTY Procedure.

Name	Type	Value	Description
PROP_RECIPIENT_PAR_LIFETIME	VARCHAR2(12)	'PAR_LIFETIME'	The lifetime of Pre-authenticated Request URLs sent to the recipient through the delta sharing API. The default is three hours: '00 03:00:00'
PROP_RECIPIENT_MIN_PAR_LIFETIME	VARCHAR2(16)	'MIN_PAR_LIFETIME'	The minimum guaranteed lifetime of a Pre-authenticated Request URLs sent to the recipient through the delta sharing API. If there is an existing PAR for the same parquet file, it is reused only if the remaining life exceeds this minimum value. The default is 2 1/2 hours: '00 02:30:00'
PROP_RECIPIENT_MAX_PAR_LIFETIME	VARCHAR2(16)	'MAX_PAR_LIFETIME'	An administrator setting that defines the maximum allowed par lifetime for any recipient. The default is 1 day: '01 00:00:00'
PROP_RECIPIENT_TOKEN_LIFETIME	VARCHAR2(14)	'TOKEN_LIFETIME'	A string, in INTERVAL DAY TO SECOND format, representing the lifetime of a delta sharing bearer token. The default is one hour: '01 00:00:00'.
PROP_RECIPIENT_MAX_TOKEN_LIFETIME	VARCHAR2(18)	'MAX_TOKEN_LIFETIME'	An administrator setting that defines the maximum allowed bearer token lifetime for any recipient. The default is 90 days: '90 00:00:00'
PROP_RECIPIENT_EMAIL	VARCHAR2(5)	'EMAIL'	The email of the recipient. This is only required for delta sharing recipients.
PROP_RECIPIENT_DESCRIPTION	VARCHAR2(11)	'DESCRIPTION'	A textual description of the recipient.

Name	Type	Value	Description
PROP_RECIPIENT_LOG_LEVEL	VARCHAR2(9)	'LOG_LEVEL'	The logging level for the recipient. This controls what gets logged in USER_SHARE_RECIPIENT_EVENTS. The value should be one of the LOG_LEVEL_* constants.
PROP_RECIPIENT_SHARING_ID	VARCHAR2(10)	'SHARING_ID'	The sharing ID of the recipient. This is only required for Autonomous Database recipients.
PROP_RECIPIENT_PAR_TYPE	VARCHAR2(8)	'PAR_TYPE'	The type of pre-authenticated request to create. This should be one of PAR_TYPE_FOLDER (the default) or PAR_TYPE_FILE.
PROP_RECIPIENT_VERSION_ACCESS	VARCHAR2(14)	'VERSION_ACCESS'	Specifies what versions of published shares a recipient can see. One of VERSION_ACCESS_CURRENT or VERSION_ACCESS_ANY.

Share Table Properties

Constants used to indicate share table properties. These are used by GET_SHARE_TABLE_PROPERTY Function, and UPDATE_SHARE_TABLE_PROPERTY Procedure.

Name	Type	Value	Description
PROP_SHARE_TABLE_SPLIT_METHOD	VARCHAR2(12)	'SPLIT_METHOD'	Specifies how a segment should be split into different files.
SPLIT_METHOD_AUTO	VARCHAR2(4)	'AUTO'	Autonomous Database determines the way in which segments are split.
SPLIT_METHOD_RANGE	VARCHAR2(5)	'RANGE'	Split segments based on the value ranges. The exact ranges are determined by analyzing the data histograms. This method requires one or more split_columns to be specified.
SPLIT_METHOD_ROWID	VARCHAR2(5)	'ROWID'	Split segments based on ROWID.
SPLIT_METHOD_HASH	VARCHAR2(4)	'HASH'	Split segments based on HASH values. This method requires one or more split_columns to be specified.

Name	Type	Value	Description
SPLIT_METHOD_VALUE	VARCHAR2(5)	'VALUE'	Split segments based on distinct values. This method requires one or more split_columns to be specified.
PROP_SHARE_TABLE_SPLIT_COLUMNS	VARCHAR2(13)	'SPLIT_COLUMNS'	Specifies what columns are used for splitting. These are required for RANGE, VALUE, AND HASH methods. The value should be a comma delimited set of shared columns (for example, 'state,city' or '"STATE","CITY"').
PROP_SHARE_TABLE_ORDER_COLUMNS	VARCHAR2(13)	'ORDER_COLUMNS'	Specifies what columns are used for ordering when gathering parquet. The value should be a comma delimited set of shared columns (for example, 'state,city' or '"STATE","CITY"').
PROP_SHARE_TABLE_SHARE_COLUMNS	VARCHAR2(13)	'SHARE_COLUMNS'	Specifies what columns are shared. The value is a comma delimited set of shared columns. For example, 'state,city' or '"STATE","CITY"'.
PROP_SHARE_TABLE_SPLIT_SIZE	VARCHAR2(10)	'SPLIT_SIZE'	Override share level split_size for a specific table. The value is a number.
PROP_SHARE_TABLE_GATHER_STATS	VARCHAR2(12)	'GATHER_STATS'	If Autonomous Database will gather statistics, the value of 'GATHER_STATS' is 'YES'. Otherwise the value is 'NO'.

Name	Type	Value	Description
PROP_SHARE_TABLE_SPLIT_ROWS	VARCHAR2(10)	'SPLIT_ROWS'	The SPLIT_ROWS property is an alternative way to specify the amount of data to be stored in each parquet file. 1. Autonomous Database calculates the number of rows in a given segment (for example, a partition of a table or the results of a view). 2. The number of rows is divided by the SPLIT_ROWS and rounded up to produce the "number of splits". 3. The data in the segment is then chunked into this number of splits using some method (for example, ROWID, index columns, specified columns, etc.) 4. Each chunk of data is then converted into a parquet file. It is recommended that you use the SPLIT_SIZE property to control the number of splits since this reflects the amount of data in each row. A row with 900 columns of VARCHAR2(4000), for example, is much larger than a row with a single column of type NUMBER(1). Only use SPLIT_ROWS if you have accounted for the average amount of data in each row.

Name	Type	Value	Description
PROP_SHARE_TABLE_FLASHBACK	VARCHAR2(9)	'FLASHBACK'	The FLASHBACK property determines if the table will be published using flashback queries to ensure read consistency between the different export files. The value should be YES, NO, or AUTO. The AUTO setting will use flashback if it is available. The YES setting will cause the publication to fail if flashback is not possible on the table or view. The NO setting will not use flashback, which means that parquet files could represent different snapshots of the data.

Share Types

Constants used during share creation to identify the type of share.

See CREATE_SHARE Procedure.

Name	Type	Value	Description
SHARE_TYPE_VERSIONED	DBMS_ID	'VERSIONED'	Share is used for sharing versioned data.
SHARE_TYPE_LIVE	DBMS_ID	'LIVE'	Share is used for sharing live data.

Sharing ID Type

Used in ASSERT_SHARING_ID and GET_SHARING_ID.

Name	Type	Value	Description
SHARING_ID_TYPE_TENANCY	VARCHAR2(7)	'TENANCY'	The sharing ID represents an OCI Tenancy.
SHARING_ID_TYPE_COMPARTMENT	VARCHAR2(11)	'COMPARTMENT'	The sharing ID represents an OCI Compartment.
SHARING_ID_TYPE_DATABASE	VARCHAR2(8)	'DATABASE'	The sharing ID represents an OCI Autonomous Database.
SHARING_ID_TYPE_REGION	VARCHAR2(6)	'REGION'	The sharing ID represents an OCI Region.

Version Access

Possible values for the VERSION_ACCESS share property. This is used to implement delta "time travel". These are the valid values associated with PROP_SHARE_VERSION_ACCESS and PROP_RECIPIENT_VERSION_ACCESS.

These can be used to enable the "version" property in the delta sharing query endpoint, which allows recipients to select specific versions of the data instead of always receiving the CURRENT version. You need to enable this on both the share and the recipient.

For example:

```
UPDATE_SHARE_PROPERTY(
          share_name => '...',
          share_property => DBMS_SHARE.PROP_SHARE_VERSION_ACCESS,
          new_value => DBMS_SHARE.VERSION_ACCESS_ANY);

UPDATE_SHARE_PROPERTY(
          share_name => '...',
          share_property => DBMS_SHARE.PROP_RECIPIENT_VERSION_ACCESS,
          new_value => DBMS_SHARE.VERSION_ACCESS_ANY);
```

Name	Type	Value	Description
VERSION_ACCESS_CURRENT	PLS_INTEGER	1	A recipient can see only the CURRENT version. (Default)
VERSION_ACCESS_ANY	PLS_INTEGER	2	A recipient can choose any CURRENT or RETIRED version.

CS_RESOURCE_MANAGER Package

The `CS_RESOURCE_MANAGER` package provides an interface to list and update consumer group parameters, and to revert parameters to default values.

- LIST_CURRENT_RULES Function
  This function lists the parameter values for each consumer group.
- LIST_DEFAULT_RULES Function
  This function returns the default values for all consumer groups.
- REVERT_TO_DEFAULT_VALUES Procedure
  This procedure reverts the specified resource manager's plan properties to default values.
- UPDATE_PLAN_DIRECTIVE Procedure
  Use this procedure to update the resource plan for a specified consumer group.

LIST_CURRENT_RULES Function

This function lists the parameter values for each consumer group.

Syntax

```
CS_RESOURCE_MANAGER.LIST_CURRENT_RULES
   RETURN TABLE;
```

Example

```
SELECT * FROM CS_RESOURCE_MANAGER.LIST_CURRENT_RULES();

CONSUMER_GROUP ELAPSED_TIME_LIMIT IO_MEGABYTES_LIMIT SHARES CONCURRENCY_LIMIT
DEGREE_OF_PARALLELISM
```

```
------------- ----------------- ----------------- ------ ----------------
--------------------
HIGH                                                    4
3                    3
MEDIUM                                                  2
2                    9
LOW                                                     1
900                    1
```

LIST_DEFAULT_RULES Function

This function returns the default values for all consumer groups.

Syntax

```
CS_RESOURCE_MANAGER.LIST_DEFAULT_RULES
  RETURN TABLE;
```

Usage Note

- By default the parallel degree policy value is `MANUAL` for the `TPURGENT` consumer group. The `CS_RESOURCE_MANAGER.LIST_DEFAULT_RULES` function shows no value for the default value for the `DEGREE_OF_PARALLELISM` for the `TPURGENT` consumer group.

Example

```
SELECT * FROM CS_RESOURCE_MANAGER.LIST_DEFAULT_RULES();
CONSUMER_GROUP ELAPSED_TIME_LIMIT IO_MEGABYTES_LIMIT SHARES CONCURRENCY_LIMIT
DEGREE_OF_PARALLELISM
------------- ----------------- ----------------- ------ ----------------
--------------------
HIGH                           0                 0      4
3                    1
MEDIUM                         0                 0      2
1                    1
LOW                            0                 0      1
300                    1
TP                             0                 0      8
300                    1
TPURGENT                       0                 0     12
300
```

REVERT_TO_DEFAULT_VALUES Procedure

This procedure reverts the specified resource manager's plan properties to default values.

Syntax

```
CS_RESOURCE_MANAGER.REVERT_TO_DEFAULT_VALUES(
   consumer_group         IN VARCHAR2,
   shares                 IN BOOLEAN   DEFAULT FALSE,
   concurrency_limit      IN BOOLEAN   DEFAULT FALSE);
```

Parameters

Parameter	Description
consumer_group	Specifies the consumer group to revert. Depending on the workload, valid values are: HIGH, MEDIUM, LOW, TP, or TPURGENT.
shares	When the value is TRUE, revert shares for the service to the default value.
concurrency_limit	When the value is TRUE, revert the concurrency_limit for the service to the default value. When you revert the concurrency_limit, both the concurrency_limit and the degree_of_parallelism values are set to their default values.

Usage Note

- When the workload type is Data Warehouse, the valid values for consumer_group are HIGH, MEDIUM, or LOW.

Examples

```
BEGIN
   CS_RESOURCE_MANAGER.REVERT_TO_DEFAULT_VALUES(
      consumer_group => 'MEDIUM',
      concurrency_limit => TRUE);
END;
/

BEGIN
   CS_RESOURCE_MANAGER.REVERT_TO_DEFAULT_VALUES(
      consumer_group => 'HIGH',
      shares => TRUE);
END;
/
```

UPDATE_PLAN_DIRECTIVE Procedure

Use this procedure to update the resource plan for a specified consumer group.

Syntax

```
CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(
   consumer_group         IN VARCHAR2,
   io_megabytes_limit     IN NUMBER   DEFAULT NULL,
   elapsed_time_limit     IN NUMBER   DEFAULT NULL,
   shares                 IN NUMBER   DEFAULT NULL,
   concurrency_limit      IN NUMBER   DEFAULT NULL);
```

Parameters

Parameter	Description
consumer_group	Specifies the consumer group to update. Depending on the workload, valid values are: HIGH, MEDIUM, LOW, TP, or TPURGENT.
io_megabytes_limit	Specifies the maximum megabytes of I/O that a SQL operation can issue. Specify a NULL value to clear the limit.
elapsed_time_limit	Specifies the maximum time in seconds that a SQL operation can run. Specify a NULL value to clear the limit.
shares	Specifies the shares value. A higher number of shares, relative to other consumer groups, increases the consumer group's CPU and I/O priority.
concurrency_limit	Specifies the maximum number of concurrent SQL statements that can be executed. This parameter is only valid with the MEDIUM consumer group.

Usage Notes

- When a SQL statement in the specified service runs more than the specified runtime limit (elapsed_time_limit) or does more I/O than the specified amount (io_megabytes_limit), then the SQL statement will be terminated.
- When the workload type is Data Warehouse, the valid values for consumer_group are HIGH, MEDIUM, or LOW.
- When the concurrency_limit parameter is specified, the only valid value for consumer_group is MEDIUM.

Examples

```
BEGIN
   CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(
      consumer_group => 'HIGH',
      shares => 8);
   CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(
      consumer_group => 'MEDIUM',
      shares => 2);
   CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(
      consumer_group => 'LOW',
      shares => 1);
END;
/

BEGIN
   CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(
      consumer_group => 'HIGH',
      io_megabytes_limit => null,
      elapsed_time_limit => null);
END;
/

BEGIN
```

```
   CS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE(
     consumer_group => 'MEDIUM',
     concurrency_limit => 2);
END;
/
```

CS_SESSION Package

The `CS_SESSION` package provides an interface to switch the database service and consumer group of the existing session.

When a connection is established with an Autonomous Database, that session is assigned a consumer group. For example, a session could be created using a connection to the LOW service of an Autonomous Database. You might want to switch the consumer group, for example from LOW to HIGH. The `CS_SESSION` package provides an API for switching. See Database Service Names for Autonomous Database for more information.

The consumer groups determine concurrency and degree of parallelism (DOP). For example, statements on a connection established to the LOW database service run serially. Statements on a connection established to the HIGH database service run in parallel. If you have a workload that requires serial statement processing with switching to a HIGH consumer group for a few statements, the `CS_SESSION` package enables you to switch.

- SWITCH_SERVICE Procedure
  This procedure switches the database service and consumer group of the current session.

SWITCH_SERVICE Procedure

This procedure switches the database service and consumer group of the current session.

Syntax

```
CS_SESSION.SWITCH_SERVICE(service_name IN varchar2);
```

Parameters

Parameter	Description
`service_name`	Specifies the consumer group to update. Depending on the workload, valid values are: `HIGH`, `MEDIUM`, `LOW`, `TP`, or `TPURGENT`.

Usage Notes

When called, the procedure switches the session to the specified service and the related consumer group. If the specified service does not exist in that database, an error message is provided. For example, if you specify 'TP' as the service name on a data warehouse workload, the error indicates it's not a valid service name. No error is reported if the current service and the specified service are identical.

The procedure does not reset session attributes. Anything the user set for their session before calling this procedure will continue as-is. For example, if a session parameter was modified and then later the session switched to a different service, the parameter value will stay the same.

Example

```
BEGIN
  CS_SESSION.SWITCH_SERVICE('HIGH');
END;
/
```

Security and Access

The `ADMIN` user is granted `EXECUTE` privilege on `CS_SESSION` with `GRANT OPTION`. The privilege is also granted to `DWROLE` without the `GRANT OPTION`.

Additional Security Considerations

If a user is granted `EXECUTE` privileges on this procedure and you do not want that user to switch to a specific service, you can use a `AFTER SET CONTAINER` trigger to block the operation. This is achieved by creating an `AFTER SET CONTAINER` trigger.

```
CREATE OR REPLACE TRIGGER SESS_SWITCH
AFTER SET CONTAINER ON DATABASE
BEGIN
IF SYS_CONTEXT('USERENV','SESSION_USER') = 'USER' and
   SYS_CONTEXT('USERENV','SERVICE_NAME') =
'serviceexample_low.adwc.oraclecloud.com'
THEN
    NULL;
ELSE
    RAISE_APPLICATION_ERROR(-20001, 'Denied!  You are not allowed to switch
service in the database');
END IF;
END;
/
```

Error Messages

The following table describes exceptions for `CS_SESSION`

Error	Message	Cause
20001	Invalid service name. Valid values are HIGH, MEDIUM, LOW.	For a data warehouse workload, a value other than 'HIGH', 'MEDIUM', 'LOW' was specified.
20001	Invalid service name. Valid values are HIGH, MEDIUM, LOW, TP, TPURGENT.	For a transaction processing workload, a value other than 'HIGH', 'MEDIUM', 'LOW', 'TP', 'TPURGENT' was specified.
20002	Service switch failed.	Failed to switch to the new service.

Autonomous Database for Experienced Oracle Database Users

This appendix provides information on using Autonomous Database for experienced Oracle Database users with Autonomous Database Serverless.

For equivalent information about using Oracle Database features and options with Autonomous Database on dedicated Exadata infrastructure, see Oracle Database Features in Dedicated Autonomous Database Deployments.

- About Autonomous Database for Experienced Oracle Database Users
  Autonomous Database configures and optimizes your database for you. You do not need to perform administration operations for configuring the database. SQL commands used for database administration such as `CREATE TABLESPACE` are not available. Similarly, other administrative interfaces and utilities such as `RMAN` are not available.
- Autonomous Database Views
  Autonomous Database provides several views that are not available in Oracle Database 19c. This topic lists the Autonomous Database specific views.
- Autonomous Database – Oracle Database Features
  Describes Oracle Database features available with Autonomous Database.
- Autonomous Database with Oracle Database 23ai
  When you provision or clone an Autonomous Database instance you can select either of these versions: Oracle Database 19c or Oracle Database 23ai.
- Autonomous Database RMAN Recovery Catalog
  You can use Oracle Autonomous Database as a Recovery Manager (RMAN) recovery catalog. A recovery catalog is a database schema that RMAN uses to store metadata about one or more Oracle databases.
- Notes for Users Migrating from Other Oracle Databases
  Describes information that is useful when you are migrating from other Oracle Databases to Oracle Autonomous Database.
- Database Features Unavailable in Autonomous Database
  Lists the Oracle Database features that are not available in Autonomous Database. Additionally, database features designed for administration are not available.
- Logical Partition Change Tracking and Materialized Views
  Describes information about the Logical Partition Change Tracking (LPCT) metadata framework and Query Rewrite with Logical Partition Change Tracking in Autonomous Database.

About Autonomous Database for Experienced Oracle Database Users

Autonomous Database configures and optimizes your database for you. You do not need to perform administration operations for configuring the database. SQL commands used for database administration such as `CREATE TABLESPACE` are not available. Similarly, other administrative interfaces and utilities such as `RMAN` are not available.

There are differences, depending on your workload: Data Warehouse, Transaction Processing, or JSON Database. See the following section appropriate to your workload:

- Data Warehouse Workload with Autonomous Database
  Autonomous Database configures and optimizes your database for you, based on your workload.

- Transaction Processing and JSON Database Workloads with Autonomous Database
  Autonomous Database configures and optimizes your database for you, based on your workload.

Data Warehouse Workload with Autonomous Database

Autonomous Database configures and optimizes your database for you, based on your workload.

Characteristics of a database with Data Warehouse workload:

- The default data and temporary tablespaces for the database are configured automatically. Adding, removing, or modifying tablespaces is not allowed. Autonomous Database creates one tablespace or multiple tablespaces automatically depending on the storage size.
- The database character set is Unicode `AL32UTF8`. See Choose a Character Set for Autonomous Database for more information.
- Compression is enabled by default. Autonomous Database uses Hybrid Columnar Compression for all tables by default. You can specify different compression methods for your tables using the compression clause in your `CREATE TABLE` or `ALTER TABLE` commands.
- Oracle Database Result Cache is enabled by default for all SQL statements.

Accessing a database:

- You do not have direct access to the database node. You can create and drop directories with `CREATE DIRECTORY` and `DROP DIRECTORY`, as described in Creating and Managing Directories on Autonomous Database.

  You can use `DBMS_CLOUD` procedures such as `DBMS_CLOUD.DELETE_FILE`, `DBMS_CLOUD.GET_OBJECT`, and `DBMS_CLOUD.PUT_OBJECT` with files and objects. You do not have direct access to the local file system.

Parallel Execution with Data Warehouse workload:

- Parallelism is determined by the database service. See See Database Service Names for Autonomous Database for more information. for details for parallelism support for each database service.
- When you want to disable parallel DML operations in your session, use the following SQL command:

  ```
  ALTER SESSION DISABLE PARALLEL DML;
  ```

  See *VLDB and Partitioning Guide* for more information on parallel DML operations.
- Manage DML Performance and Compression for Data Warehouse Workloads
- Create Staging Tables for Data Warehouse Workloads
  Autonomous Database supports staging tables that are optimized for loading data into a data warehouse.

Manage DML Performance and Compression for Data Warehouse Workloads

Autonomous Database with Data Warehouse workloads uses Hybrid Columnar Compression for all tables by default. This gives the best compression ratio and optimal performance for direct-path load operations like the loads done using the `DBMS_CLOUD` package. If you perform DML operations like `UPDATE` and `MERGE` on your tables these may cause the compression ratio

for the affected rows to decrease leading to larger table sizes. These operations may also perform slower compared to the same operations on an uncompressed table.

For the best compression ratio and optimal performance Oracle recommends using bulk operations like direct-path loads and `CREATE TABLE AS SELECT` statements. But, if your workload requires frequent DML operations like `UPDATE` and `MERGE` on large parts of a table, you can create those tables as uncompressed tables to achieve better DML performance. For example, the following statement creates the table `SALES` as an uncompressed table:

```
CREATE TABLE sales (
    prod_id             NUMBER          NOT NULL,
    cust_id             NUMBER          NOT NULL,
    time_id             DATE            NOT NULL,
    channel_id          NUMBER          NOT NULL,
    promo_id            NUMBER          NOT NULL,
    quantity_sold       NUMBER(10,2)    NOT NULL,
    amount_sold         NUMBER(10,2)    NOT NULL)
NOCOMPRESS;
```

At any point in time you can use the `ALTER TABLE MOVE` statement to compress these tables without impacting queries accessing them. For example, the following statement compresses the table `SALES` using Hybrid Columnar Compression.

```
ALTER TABLE sales MOVE COLUMN STORE COMPRESS FOR QUERY HIGH;
```

Create Staging Tables for Data Warehouse Workloads

Autonomous Database supports staging tables that are optimized for loading data into a data warehouse.

A staging table is a table with the `STAGING` property set. This applies the following characteristics:

- Any form of compression is explicitly turned off and disallowed on a staging table for any data load. The command `ALTER TABLE COMPRESS` is not allowed.
- Setting the `STAGING` property on an existing table does not impact the storage of existing data but does impact future data loads.
- Autonomous Database uses dynamic sampling for statistics for tables with the staging property set, and does not collect statistics on staging tables.
- Dropping staging tables immediately removes the table, bypassing the recycle bin. Setting the `recyclebin` initialization parameter to the value `ON` does not enable the recycle bin.

The characteristics of Autonomous Database partitioned staging tables includes the above, plus the following:

- Any form of compression is explicitly turned off and disallowed on all of the table's partitions and subpartitions.
- You cannot change the default attributes of the table to use compress with `ALTER TABLE MODIFY DEFAULT ATTRIBUTES`.
- You cannot perform partition maintenance operations that move data and compress the data. For example, the following are not allowed when you try to apply compression: `ALTER TABLE` with `MOVE PARTITION`, `MERGE PARTITIONS`, `SPLIT PARTITION`, or `SPLIT SUBPARTITION`.

- You cannot repartition a table with `ALTER TABLE MODIFY PARTITION` and specify any resulting partition to be compressed.

Define staging tables when you create a table or by altering an existing table as follows:

1. Create a table with the `STAGING` property.

   For example:

   ```
   CREATE TABLE staging_table (col1 number, col2 varchar2(100)) FOR STAGING;

   CREATE TABLE part_staging_table (col1 number, col2 varchar2(100))
       PARTITION BYRANGE (col1)
             (PARTITION p1 VALUESLESS THAN (100),
              PARTITION pmax VALUESLESS THAN (MAXVALUE)) FOR STAGING;
   ```

2. Change an existing table to set the `STAGING` property.

   For example:

   ```
   ALTER TABLE staging_table FOR STAGING;
   ```

3. You can verify the `STAGING` property for a table using one of the following views: `USER_TABLES`, `ALL_TABLES`, or `DBA_TABLES`.

   In these views, the `STAGING` column indicates the staging table property, a value `YES` indicates a `STAGING` table, a value `NO` is shown for all other tables.

You can alter a table to remove the `STAGING` property. For example:

```
ALTER TABLE staging_table NOT FOR STAGING;
```

Note the following for altering a table with `NOT FOR STAGING`:

- After you alter a table with `NOT FOR STAGING`, the compression attribute and existing data are not affected and are kept as uncompressed until you explicitly alter the table and specify compression. You can change the table compression and `ALTER TABLE COMPRESS` is allowed.
- Altering a table with `NOT FOR STAGING` does not trigger statistics collection. After you change the table property with `NOT FOR STAGING`, you can collect statistics, either manually or automatically.
- After you alter a table with `NOT FOR STAGING`, when the recycle bin is enabled dropping the table puts the table in the recycle bin.

Transaction Processing and JSON Database Workloads with Autonomous Database

Autonomous Database configures and optimizes your database for you, based on your workload.

Characteristics of Autonomous Database with Transaction Processing or JSON Database workloads:

- The default data and temporary tablespaces for the database are configured automatically. Adding, removing, or modifying tablespaces is not allowed. Autonomous Database creates one tablespace or multiple tablespaces automatically depending on the storage size.

- The database character set is Unicode `AL32UTF8`. See Choose a Character Set for Autonomous Database for more information.
- Compression is not enabled by default but Autonomous Database honors a compression clause if compression is specified on a table.

Accessing a database:

- You do not have direct access to the database node. You can create and drop directories with `CREATE DIRECTORY` and `DROP DIRECTORY`, as described in Create and Manage Directories.

  You can use `DBMS_CLOUD` procedures such as `DBMS_CLOUD.DELETE_FILE`, `DBMS_CLOUD.GET_OBJECT`, and `DBMS_CLOUD.PUT_OBJECT` with files and objects. You do not have direct access to the local file system.

Parallel Execution with Transaction Processing or JSON Database workloads:

- Parallelism is determined by the database service you use. See Database Service Names for Autonomous Database for details about parallelism support for each database service.
- When you want to run DML operations in parallel and the database service you are using allows this, you can enable parallel DML in your session using the following SQL command:

  ```
  ALTER SESSION ENABLE PARALLEL DML;
  ```

  See *VLDB and Partitioning Guide* for more information on parallel DML operations.
- If you create an index manually and specify the `PARALLEL` clause, the `PARALLEL` attribute remains after the index is created. In this case SQL statements can run in parallel unbeknownst to the end user.

  To specify serial execution, change the `INDEX` parallel clause to `NOPARALLEL` or set the `PARALLEL` degree attribute to 1 to specify serial execution:

  ```
  ALTER INDEX index_name NOPARALLEL;
  ```

  or

  ```
  ALTER INDEX index_name PARALLEL 1;
  ```

Autonomous Database Views

Autonomous Database provides several views that are not available in Oracle Database 19c. This topic lists the Autonomous Database specific views.

- Track Table and Partition Scan Access with Autonomous Database Views
  Oracle Autonomous Database tracks the scan count for tables and partitions. Use the table access stats data dictionary and dynamic views to retrieve scan count information.
- Track Oracle Cloud Infrastructure Resources, Cost and Usage Reports with Autonomous Database Views
  Oracle Autonomous Database tracks the Oracle Cloud Infrastructure resources, cost and usage reports. You can access these reports using the OCI views.
- Pipeline Views
  Lists details for the `DBMS_CLOUD_PIPELINE` views.

- DBMS_CLOUD_AI Views
  The DBMS_CLOUD_AI package uses the following views.
- DBMS_CLOUD_FUNCTION Views
  The DBMS_CLOUD_FUNCTION package uses the following views.
- DBMS_CLOUD_LINK Views
  The DBMS_CLOUD_LINK package uses the following views.
- DBMS_DATA_ACCESS Views
  Describes the Autonomous Database dynamic performance views you can use with the DBMS_DATA_ACCESS package to monitor Table Hyperlink URL usage.
- DBA_PROXY_SQL Views
  Describes the Autonomous Database DBA_PROXY_SQL views.
- Lockdown Errors View
  Lists details for the DBA_LOCKDOWN_ERRORS view.
- Oracle Cloud Infrastructure Logging Interface Views
  The Oracle Cloud Infrastructure (OCI) logging interface views enable authorized users to access log data through a set of views.
- Real Application Testing Capture Replay Views
  Oracle Real Application Testing capture replay views.
- Stripe Views on Autonomous Database
  Stripe is an online payment processing and credit card processing platform for businesses. Users can query views created on top of Stripe APIs using the DBMS_CLOUD package to get Stripe information such as invoices, subscriptions, and customers.
- Track DBMS_CLOUD Load Operations
  DBMS_CLOUD tracks details for load operations for each instance and for users.

Track Table and Partition Scan Access with Autonomous Database Views

Oracle Autonomous Database tracks the scan count for tables and partitions. Use the table access stats data dictionary and dynamic views to retrieve scan count information.

- GV$TABLE_ACCESS_STATS and V$TABLE_ACCESS_STATS Views
  The GV$TABLE_ACCESS_STATS and V$TABLE_ACCESS_STATS views list the scan count for tables and partitions. The scan data collection begins at instance startup time.
- ALL_TABLE_ACCESS_STATS and DBA_TABLE_ACCESS_STATS Views
  The ALL_TABLE_ACCESS_STATS and DBA_TABLE_ACCESS_STATS views list the scan count for tables and partitions. The scan data collection begins at instance startup time.
- USER_TABLE_ACCESS_STATS View
  The USER_TABLE_ACCESS_STATS view lists the scan count for the user's tables and partitions. The scan data collection begins at instance startup time.

GV$TABLE_ACCESS_STATS and V$TABLE_ACCESS_STATS Views

The GV$TABLE_ACCESS_STATS and V$TABLE_ACCESS_STATS views list the scan count for tables and partitions. The scan data collection begins at instance startup time.

Column	Datatype	Description
READ_COUNT	NUMBER	Aggregated scan count since instance startup
OBJECT_ID	NUMBER	Object ID of the table or partition

Column	Datatype	Description
INST_ID	NUMBER	Instance number where table/partition was scanned This column (INST_ID) is only shown in GV$TABLE_ACCESS_STATS
CON_ID	NUMBER	Container ID of the database

ALL_TABLE_ACCESS_STATS and DBA_TABLE_ACCESS_STATS Views

The ALL_TABLE_ACCESS_STATS and DBA_TABLE_ACCESS_STATS views list the scan count for tables and partitions. The scan data collection begins at instance startup time.

> **Note:**
>
> The ALL_TABLE_ACCESS_STATS and DBA_TABLE_ACCESS_STATS views do not list scan count information for Oracle-maintained schemas.

Column	Datatype	Description
TABLE_OWNER	VARCAR2(128)	Owner of the table
TABLE_NAME	VARCAR2(128)	Name of the table
PARTITION_NAME	VARCAR2(128)	Name of the partition A NULL value specifies a non-partitioned table
INSTANCE_ID	NUMBER	Instance number where table or partition was scanned
READ_COUNT	NUMBER	Aggregated scan count since instance startup

USER_TABLE_ACCESS_STATS View

The USER_TABLE_ACCESS_STATS view lists the scan count for the user's tables and partitions. The scan data collection begins at instance startup time.

Column	Datatype	Description
TABLE_NAME	VARCAR2(128)	Name of the table
PARTITION_NAME	VARCAR2(128)	Name of the partition A NULL value specifies a non-partitioned table
INSTANCE_ID	NUMBER	Instance number where table/partition was scanned
READ_COUNT	NUMBER	Aggregated scan count since instance startup

Track Oracle Cloud Infrastructure Resources, Cost and Usage Reports with Autonomous Database Views

Oracle Autonomous Database tracks the Oracle Cloud Infrastructure resources, cost and usage reports. You can access these reports using the OCI views.

- Prerequisite Steps to Use OCI Resource Views
  Describes the prerequisite steps you must perform to use OCI resource views on Autonomous Database.
- OCI_AUTONOMOUS_DATABASES View
  `OCI_AUTONOMOUS_DATABASES` describes all the Oracle Cloud Infrastructure Autonomous Databases in the Oracle Cloud Infrastructure tenancy obtained from the current Autonomous Database instance.
- OCI_BUDGET_ALERT_RULES View
  `OCI_BUDGET_ALERT_RULES` describes all the Oracle Cloud Infrastructure budget alert rules in the Oracle Cloud Infrastructure tenancy obtained from the current Autonomous Database instance.
- OCI_BUDGET_SUMMARY View
  `OCI_BUDGET_SUMMARY` describes all the Oracle Cloud Infrastructure budget summaries in the Oracle Cloud Infrastructure tenancy obtained from the current Autonomous Database instance.
- OCI_COST_DATA View
  `OCI_COST_DATA` describes all the Oracle Cloud Infrastructure cost data for the Oracle Cloud Infrastructure tenancy obtained from the current Autonomous Database instance.
- OCI_OBJECTSTORAGE_BUCKETS View
  `OCI_OBJECTSTORAGE_BUCKETS` describes all the Oracle Cloud Infrastructure object storage buckets in the Oracle Cloud Infrastructure tenancy obtained from the current Autonomous Database instance.
- OCI_USAGE_DATA View
  `OCI_USAGE_DATA` describes all the Oracle Cloud Infrastructure usage data for the Oracle Cloud Infrastructure tenancy obtained from the current Autonomous Database instance.

Prerequisite Steps to Use OCI Resource Views

Describes the prerequisite steps you must perform to use OCI resource views on Autonomous Database.

> **Note:**
>
> Only ADMIN user has access to the OCI resource views by default. To access these views as another user, the ADMIN must grant READ privileges.

To query an OCI resource view, do the following:

1. Create a dynamic group that includes your Autonomous Database instance and define the required policies to access a view.

   For example, the Autonomous Database instance is specified in the `resource.id` parameter with an OCID:

   ```
   resource.id = '<your_Autonomous_Database_instance_OCID>'
   ```

   Each view shows the details for the policy that you must define to query the view.

   See Perform Prerequisites to Use Resource Principal with Autonomous Database for details on creating a dynamic group and defining policies.

For example, to access all of the views, define the following policy:

```
Define tenancy usage-report as
ocid1.tenancy.oc1..aaaaaaaaned4fkpkisbwjlr56u7cj63lf3wffbilvqknstgtvzub7vhq
kggq
Endorse dynamic-group <group-name> to read objects in tenancy usage-report
Allow dynamic-group <group-name> to read buckets in tenancy
Allow dynamic-group <group-name> to read autonomous-database in tenancy
Allow dynamic-group <group-name> to read usage-budgets in tenancy
```

> **Note:**
>
> Do not replace the OCID in this policy with another OCID. This `usage-report` OCID provides the Oracle Cloud Infrastructure usage data for your tenancy.

2. Verify that resource principal is enabled for the ADMIN user on the Autonomous Database instance.

```
SELECT owner, credential_name FROM dba_credentials
   WHERE credential_name = 'OCI$RESOURCE_PRINCIPAL' AND owner = 'ADMIN';

OWNER CREDENTIAL_NAME
----- --------------------
ADMIN OCI$RESOURCE_PRINCIPAL
```

If the resource principal is not enabled, then enable the resource principal:

```
EXEC DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL();
```

See Use Resource Principal to Access Oracle Cloud Infrastructure Resources for more information.

3. Run a query on an OCI resource view.

   For example:

```
SELECT NAME, APPROXIMATESIZE FROM OCI_OBJECTSTORAGE_BUCKETS;
SELECT * FROM OCI_USAGE_DATA;
```

OCI_AUTONOMOUS_DATABASES View

`OCI_AUTONOMOUS_DATABASES` describes all the Oracle Cloud Infrastructure Autonomous Databases in the Oracle Cloud Infrastructure tenancy obtained from the current Autonomous Database instance.

To query this view you need a dynamic group that includes your Autonomous Database instance and the following policy defined on that dynamic group:

```
Allow dynamic-group <group-name> to read autonomous-database in tenancy
```

This policy lets you list all Autonomous Databases in your tenancy. Optionally you can restrict it to list Autonomous Databases in a given compartment:

```
Allow dynamic-group <group-name> to read autonomous-database in compartment
<compartment-name>
```

Column	Datatype	Description
DISPLAYNAME	VARCHAR2	The user friendly name for the Autonomous Database
REGION	VARCHAR2	Region Name
COMPARTMENTID	VARCHAR2	The OCID of the compartment
ID	VARCHAR2	The OCID of the Autonomous Database
DBNAME	VARCHAR2	The database name
LIFECYCLESTATE	VARCHAR2	The current state of the Autonomous Database
TIMECREATED	VARCHAR2	The date and time the Autonomous Database was created
DATASTORAGESIZEINTBS	VARCHAR2	The quantity of data in the database in terabytes
LICENSEMODEL	VARCHAR2	The Oracle license model that applies to the Autonomous Database
SERVICECONSOLEURL	VARCHAR2	The URL of the Service Console for the Autonomous Database
APEXDETAILS	CLOB	Information about Oracle APEX Application Development
AREPRIMARYWHITELISTEDIPSUSED	VARCHAR2	Primary White Listed IPs
AUTONOMOUSCONTAINERDATABASEID	VARCHAR2	The Autonomous Container Database OCID
AUTONOMOUSMAINTENANCESCHEDULETYPE	VARCHAR2	Maintenance Schedule Type
AVAILABLEUPGRADEVERSIONS	VARCHAR2	List of Oracle Database versions available for a database upgrade
BACKUPCONFIG	CLOB	Autonomous Database Backup Config
CONNECTIONSTRINGS	CLOB	Autonomous Database Connection Strings
CONNECTIONURLS	CLOB	Autonomous Database Connection URLs
CPUCORECOUNT	NUMBER	The number of OCPU cores to be made available to the database
CUSTOMERCONTACTS	CLOB	The Customer Contacts
DATASAFESTATUS	VARCHAR2	Status of the Data Safe registration for this Autonomous Database
DATASTORAGESIZEINGBS	NUMBER	The quantity of data in the database in gigabytes
DBVERSION	VARCHAR2	The Oracle Database version for the Autonomous Database
DATAGUARDREGIONTYPE	VARCHAR2	The Autonomous Data Guard region type of the Autonomous Database
DBWORKLOAD	VARCHAR2	The Autonomous Database workload type
DEFINEDTAGS	CLOB	Defined tags for the resource
FAILEDDATARECOVERYINSECONDS	NUMBER	Indicates the number of seconds of data loss for an Autonomous Data Guard failover
FREEFORMTAGS	CLOB	Free form tags for the resource
INFRASTRUCTURETYPE	VARCHAR2	The infrastructure type this resource belongs to
ISACCESSCONTROLENABLED	VARCHAR2	Indicates if the database level access control is enabled
ISAUTOSCALINGENABLED	VARCHAR2	Indicates if auto scaling is enabled for the Autonomous Database

Column	Datatype	Description
ISDATAGUARDENABLED	VARCHAR2	Indicates whether the Autonomous Database has a local Autonomous Data Guard enabled
ISDEDICATED	VARCHAR2	True if the database uses dedicated Exadata infrastructure
ISFREETIER	VARCHAR2	Indicates if this is an Always Free resource
ISMTLSCONNECTIONREQUIRED	VARCHAR2	Indicates whether the Autonomous Database requires mTLS connections
ISPREVIEW	VARCHAR2	Indicates if the Autonomous Database version is a preview version
ISREFRESHABLECLONE	VARCHAR2	Indicates whether the Autonomous Database is a refreshable clone
KEYHISTORYENTRY	CLOB	Key History Entry
KEYSTOREID	VARCHAR2	The OCID of the key store
KEYSTOREWALLETNAME	VARCHAR2	The wallet name for Oracle Cloud Infrastructure Vault
KMSKEYID	VARCHAR2	The OCID of the key container that is used as the master encryption key
KMSKEYLIFECYCLEDETAILS	VARCHAR2	Customer managed key lifecycle details
LIFECYCLEDETAILS	VARCHAR2	Information about the current lifecycle state
NSGIDS	CLOB	A list of the OCIDs of the network security groups NSGs
OCPUCOUNT	NUMBER	The number of OCPU cores to be made available to the database
OPENMODE	VARCHAR2	The Autonomous Database open mode
OPERATIONSINSIGHTSSTATUS	VARCHAR2	Status of OCI Ops Insights for this Autonomous Database
PEERDBIDS	VARCHAR2	The list of OCIDs of standby databases located in Autonomous Data Guard
PERMISSIONLEVEL	CLOB	The Autonomous Database permission level
PRIVATEENDPOINT	VARCHAR2	The private endpoint for the resource
PRIVATEENDPOINTIP	VARCHAR2	The private endpoint IP address for the resource
PRIVATEENDPOINTLABEL	VARCHAR2	The private endpoint label for the resource
REFRESHABLEMODE	VARCHAR2	The refresh mode of the clone
REFRESHABLESTATUS	VARCHAR2	The refresh status of the clone
ROLE	VARCHAR2	The Autonomous Data Guard role
SOURCEID	VARCHAR2	The OCID of the source Autonomous Database that was cloned
SQLWEBDEVELOPERURL	VARCHAR2	The Database Actions (SQL Developer Web) URL for the Autonomous Database
STANDBYDB	CLOB	Autonomous Database Standby Summary
STANDBYWHITELISTEDIPS	CLOB	The client IP access control list
SUBNETID	VARCHAR2	The OCID of the subnet the resource is associated with
SUPPORTEDREGIONSTOCLONETO	CLOB	The list of regions that support the creation of Autonomous Data Guard
SYSTEMTAGS	CLOB	System tags for this resource
TIMEDATAGUARDROLECHANGED	VARCHAR2	The date and time the Autonomous Data Guard role was switched
TIMEDELETIONOFFREEAUTONOMOUSDATABASE	NUMBER	Time deletion of Free Autonomous Database

Column	Datatype	Description
TIMELOCALDATAGUARDENABLED	VARCHAR2	The date and time that Autonomous Data Guard was enabled for the Autonomous Database
TIMEMAINTENANCEBEGIN	VARCHAR2	The date and time when maintenance will begin
TIMEMAINTENANCEEND	VARCHAR2	The date and time when maintenance will end
TIMEOFLASTFAILOVER	VARCHAR2	The timestamp of the last failover operation
TIMEOFLASTREFRESH	VARCHAR2	The date and time of the last refresh
TIMEOFLASTREFRESHPOINT	VARCHAR2	The refresh point timestamp
TIMEOFLASTSWITCHOVER	VARCHAR2	The timestamp of the last switchover operation for the Autonomous Database
TIMEOFNEXTREFRESH	VARCHAR2	The date and time of next refresh
TIMERECLAMATIONOFFREEAUTONOMOUSDATABASE	VARCHAR2	The date and time the Always Free database
USEDDATASTORAGESIZEINTBS	NUMBER	The amount of storage that has been used in terabytes
VAULTID	VARCHAR2	The OCID of the Oracle Cloud Infrastructure Vault
WHITELISTEDIPS	CLOB	The client IP access control list

OCI_BUDGET_ALERT_RULES View

`OCI_BUDGET_ALERT_RULES` describes all the Oracle Cloud Infrastructure budget alert rules in the Oracle Cloud Infrastructure tenancy obtained from the current Autonomous Database instance.

Queries against this view return results only if you have budgets and budget alerts created in your tenancy.

See Budgets Overview for more information.

To query this view you need a dynamic group that includes your Autonomous Database instance and the following policy defined on that dynamic group:

```
Allow dynamic-group <group-name> to read usage-budgets in tenancy
```

This policy lets you list budget summary and budget alerts in your tenancy (if you created a budget and a budget alert). Optionally you can restrict the result returned by querying the view to a given compartment:

```
Allow dynamic-group <group-name> to read usage-budgets in compartment <compartment-name>
```

Column	Datatype	Description
BUDGETID	VARCHAR2	The OCID of the budget
REGION	VARCHAR2	Region name
COMPARTMENTID	VARCHAR2	The compartment ID in which the bucket is authorized
DEFINEDTAGS	CLOB	Defined tags for the resource
DESCRIPTION	VARCHAR2	The description of the alert rule
DISPLAYNAME	VARCHAR2	The name of the alert rule
FREEFORMTAGS	CLOB	Free-form tags for the resource
ID	VARCHAR2	The OCID of the alert rule
LIFECYCLESTATE	VARCHAR2	The current state of the alert rule

Column	Datatype	Description
MESSAGE	VARCHAR2	The custom message that will be sent when the alert is triggered
RECIPIENTS	VARCHAR2	The audience that receives the alert when it triggers
THRESHOLD	NUMBER	The threshold for triggering the alert
THRESHOLDTYPE	VARCHAR2	The type of threshold
TIMECREATED	VARCHAR2	The time when the budget was created
TIMEUPDATED	VARCHAR2	The time when the budget was updated
TYPE	VARCHAR2	ACTUAL or FORECAST types of alert triggers
VERSION	NUMBER	The version of the alert rule

OCI_BUDGET_SUMMARY View

OCI_BUDGET_SUMMARY describes all the Oracle Cloud Infrastructure budget summaries in the Oracle Cloud Infrastructure tenancy obtained from the current Autonomous Database instance.

Queries against this view return results only if you have budgets created in your tenancy.

See Budgets Overview for more information.

To query this view you need a dynamic group that includes your Autonomous Database instance and the following policy defined on that dynamic group:

```
Allow dynamic-group <group-name> to read usage-budgets in tenancy
```

This policy lets you list budget summary and budget alerts in your tenancy (if you created a budget and a budget alert). Optionally you can restrict the result returned by querying the view to a given compartment:

```
Allow dynamic-group <group-name> to read usage-budgets in compartment <compartment-name>
```

Column	Datatype	Description
REGION	VARCHAR2	Region name
COMPARTMENTID	VARCHAR2	The OCID of the compartment
AMOUNT	NUMBER	The amount of the budget, expressed in the currency of a rate card
DEFINEDTAGS	CLOB	Defined tags for the resource
FREEFORMTAGS	CLOB	Free-form tags for the resource
DISPLAYNAME	VARCHAR2	The display name of the budget
LIFECYCLESTATE	VARCHAR2	The current state of the budget
ACTUALSPEND	NUMBER	The actual spend in currency for the current budget cycle
ALERTRULECOUNT	NUMBER	The total number of alert rules in the budget
BUDGETPROCESSINGPERIODSTARTOFFSET	NUMBER	The number of days offset from the first day of the month, at which the budget processing period starts
DESCRIPTION	VARCHAR2	The description of the budget
FORECASTEDSPEND	NUMBER	The forecasted spend in currency by the end of the current budget cycle
ID	VARCHAR2	The OCID of the budget

Column	Datatype	Description
RESETPERIOD	VARCHAR2	The reset period for the budget
TARGETS	CLOB	The list of targets on which the budget is applied
TARGETCOMPARTMENTID	VARCHAR2	Target compartment OCID
TARGETTYPE	VARCHAR2	The type of target on which the budget is applied
TIMECREATED	VARCHAR2	The time the budget was created
TIMESPENDCOMPUTED	VARCHAR2	The time the budget spend was last computed
TIMEUPDATED	VARCHAR2	The time the budget was updated
VERSION	VARCHAR2	The version of the budget

OCI_COST_DATA View

OCI_COST_DATA describes all the Oracle Cloud Infrastructure cost data for the Oracle Cloud Infrastructure tenancy obtained from the current Autonomous Database instance.

To query this view you need a dynamic group that includes your Autonomous Database instance and the following policy defined on that dynamic group:

```
Define tenancy usage-report as
ocid1.tenancy.oc1..aaaaaaaaned4fkpkisbwjlr56u7cj63lf3wffbilvqknstgtvzub7vhqkggq
Endorse dynamic-group <group-name> to read objects in tenancy usage-report
```

> **Note:**
>
> Do not replace the OCID in this policy with another OCID. This usage-report OCID provides the Oracle Cloud Infrastructure usage data for your tenancy.

Column	Datatype	Description
REFERENCE_NUMBER	VARCHAR2	Reference Number/Line identifier used for debugging and corrections
TENANT_ID	VARCHAR2	The identifier (OCID) for the Oracle Cloud Infrastructure tenant
INTERVAL_USAGE_START	TIMESTAMP	The start time of the usage interval for the resource in UTC
INTERVAL_USAGE_END	TIMESTAMP	The end time of the usage interval for the resource in UTC
SERVICE_NAME	VARCHAR2	The service that the resource is in
COMPARTMENT_ID	VARCHAR2	The ID of the compartment that contains the resource
COMPARTMENT_NAME	VARCHAR2	The name of the compartment that contains the resource
REGION	VARCHAR2	The region that contains the resource
AVAILABILITY_DOMAIN	VARCHAR2	The availability domain that contains the resource
RESOURCE_ID	VARCHAR2	The identifier for the resource
BILLED_QUANTITY	VARCHAR2	The quantity of the resource that has been billed over the usage interval
BILLED_QUANTITY_OVERAGE	VARCHAR2	The usage quantity for which you were billed
SUBSCRIPTION_ID	VARCHAR2	A unique identifier associated with your commitment or subscription

Column	Datatype	Description
PRODUCT_SKU	VARCHAR2	The Part Number for the resource in the line
PRODUCT_DESCRIPTION	VARCHAR2	The product description for the resource in the line
UNIT_PRICE	VARCHAR2	The cost billed to you for each unit of the resource used
UNIT_PRICE_OVERAGE	VARCHAR2	The cost per unit of usage for overage usage of a resource
MY_COST	VARCHAR2	The cost charged for this line of usage
MY_COST_OVERAGE	VARCHAR2	The cost billed for overage usage of a resource
CURRENCY_CODE	VARCHAR2	The currency code for your tenancy
BILLING_UNIT_READABLE	VARCHAR2	The unit measure associated with the usage/ billedQuantity in the line
SKU_UNIT_DESCRIPTION	VARCHAR2	The unit used for measuring billed quantity
OVERAGE_FLAG	CHAR	Flag used for overage usage
IS_CORRECTION	VARCHAR2	Used if the current line is a correction
BACK_REFERENCE_NUMBER	VARCHAR2	Data amendments and corrections reference
CREATED_BY	VARCHAR2	The user who created the service
CREATED_ON	TIMESTAMP	The time when the service was created
FREE_TIER_RETAINED	VARCHAR2	Is the service retained on free tier

OCI_OBJECTSTORAGE_BUCKETS View

OCI_OBJECTSTORAGE_BUCKETS describes all the Oracle Cloud Infrastructure object storage buckets in the Oracle Cloud Infrastructure tenancy obtained from the current Autonomous Database instance.

To query this view you need a dynamic group that includes your Autonomous Database instance and the following policy defined on that dynamic group:

```
Allow dynamic-group <group-name> to read buckets in tenancy
```

This policy lets you list object storage buckets in your tenancy. Optionally you can restrict the result returned by querying this view to a given compartment:

```
Allow dynamic-group <group-name> to read buckets in compartment <compartment-name>
```

Column	Datatype	Description
REGION	VARCHAR2	Region name
COMPARTMENTID	VARCHAR2	The compartment ID in which the bucket is authorized
NAMESPACE	VARCHAR2	The Object Storage namespace in which the bucket resides
APPROXIMATECOUNT	NUMBER	The approximate number of objects in the bucket
APPROXIMATESIZE	NUMBER	The approximate total size in bytes of all objects in the bucket
AUTOTIERING	VARCHAR2	The auto tiering status on the bucket
CREATEDBY	VARCHAR2	The OCID of the user who created the bucket
DEFINEDTAGS	CLOB	Defined tags for the resource
FREEFORMTAGS	CLOB	Free-form tags for the resource
ETAG	VARCHAR2	The entity tag (ETag) for the bucket

Column	Datatype	Description
ID	VARCHAR2	The OCID of the bucket
ISREADONLY	VARCHAR2	Whether or not this bucket is read only
KMSKEYID	VARCHAR2	The OCID of a master encryption key
METADATA	VARCHAR2	Arbitrary string keys and values for user-defined metadata
NAME	VARCHAR2	The name of the bucket
OBJECTEVENTSENABLED	VARCHAR2	Whether or not events are emitted for object state changes in this bucket
OBJECTLIFECYCLEPOLICYETAG	VARCHAR2	The entity tag (ETag) for the live object lifecycle policy on the bucket
PUBLICACCESSTYPE	VARCHAR2	The type of public access enabled on this bucket
REPLICATIONENABLED	VARCHAR2	Whether or not this bucket is a replication source
STORAGETIER	VARCHAR2	The storage tier type assigned to the bucket
TIMECREATED	VARCHAR2	The date and time the bucket was created
VERSIONING	VARCHAR2	The versioning status on the bucket

OCI_USAGE_DATA View

`OCI_USAGE_DATA` describes all the Oracle Cloud Infrastructure usage data for the Oracle Cloud Infrastructure tenancy obtained from the current Autonomous Database instance.

> **Note:**
>
> You can use OCI usage reports and `OCI_USAGE_DATA` view for information on cost and usage *before* January 31st, 2025. The OCI cost and usage reports are deprecated. You can continue to access your existing usage report CSV files until July 31, 2025. See Cost and Usage Reports for more information.

To query this view you need a dynamic group that includes your Autonomous Database instance and the following policy defined on that dynamic group:

```
Define tenancy usage-report as
ocid1.tenancy.oc1..aaaaaaaaned4fkpkisbwjlr56u7cj63lf3wffbilvqknstgtvzub7vhqkggq
Endorse dynamic-group <group-name> to read objects in tenancy usage-report
```

> **Note:**
>
> Do not replace the OCID in this policy with another OCID. This `usage-report` OCID provides the Oracle Cloud Infrastructure cost and usage data for your tenancy.

Column	Datatype	Description
REFERENCE_NUMBER	VARCHAR2	Reference Number/Line identifier used for debugging and corrections
TENANT_ID	VARCHAR2	The identifier (OCID) for the Oracle Cloud Infrastructure tenant

Column	Datatype	Description
INTERVAL_USAGE_START	TIMESTAMP	The start time of the usage interval for the resource in UTC
INTERVAL_USAGE_END	TIMESTAMP	The end time of the usage interval for the resource in UTC
SERVICE_NAME	VARCHAR2	The service that the resource is in
RESOURCE_NAME	VARCHAR2	The resource name used by the metering system
COMPARTMENT_ID	VARCHAR2	The ID of the compartment that contains the resource
COMPARTMENT_NAME	VARCHAR2	The name of the compartment that contains the resource
REGION	VARCHAR2	The region that contains the resource
AVAILABILITY_DOMAIN	VARCHAR2	The availability domain that contains the resource
RESOURCE_ID	VARCHAR2	The identifier for the resource
CONSUMED_QUANTITY	VARCHAR2	The quantity of the resource that has been consumed over the usage interval
BILLED_QUANTITY	VARCHAR2	The quantity of the resource that has been billed over the usage interval
CONSUMED_QUANTITY_UNITS	VARCHAR2	The unit for the consumed quantity and billed quantity
CONSUMED_QUANTITY_MEASURE	VARCHAR2	The measure for the consumed quantity and billed quantity
IS_CORRECTION	VARCHAR2	Used if the current line is a correction
BACK_REFERENCE_NUMBER	VARCHAR2	Data amendments and corrections reference
CREATED_BY	VARCHAR2	The user who created the service
CREATED_ON	TIMESTAMP	The time when the service was created
FREE_TIER_RETAINED	VARCHAR2	Is the service retained on free tier

Pipeline Views

Lists details for the DBMS_CLOUD_PIPELINE views.

- DBA_CLOUD_PIPELINES View
  DBA_CLOUD_PIPELINES displays the list of pipelines created in the database. The view contains one row for each pipeline.
- USER_CLOUD_PIPELINES View
  USER_CLOUD_PIPELINES displays the list of pipelines created in the database for the user. The view contains one row for each pipeline.
- DBA_CLOUD_PIPELINE_HISTORY View
  DBA_CLOUD_PIPELINE_HISTORY lists the pipeline scheduled jobs in the database. Use this view to monitor the health of the pipeline and detect failures in pipeline job runs.
- USER_CLOUD_PIPELINE_HISTORY View
  USER_CLOUD_PIPELINE_HISTORY lists the pipeline scheduled jobs in the database for the user. Use this view to monitor the health of the pipeline and detect failures in pipeline job runs.
- DBA_CLOUD_PIPELINE_ATTRIBUTES View
  DBA_CLOUD_PIPELINE_ATTRIBUTES lists the attributes of a pipelines in the database. Attributes control the configuration and behavior of the pipeline jobs.

- USER_CLOUD_PIPELINE_ATTRIBUTES View
  USER_CLOUD_PIPELINE_ATTRIBUTES lists the attributes of a pipelines in the database for the user. Attributes control the configuration and behavior of the pipeline.

DBA_CLOUD_PIPELINES View

DBA_CLOUD_PIPELINES displays the list of pipelines created in the database. The view contains one row for each pipeline.

Column	Datatype	NULL	Description
PIPELINE_ID	NUMBER	NOT NULL	Unique number assigned to the pipeline.
PIPELINE_NAME	VARCHAR2(128)	NOT NULL	Name of the pipeline.
OWNER	VARCHAR2(128)	NOT NULL	Owner schema of the pipeline
PIPELINE_TYPE	VARCHAR2(7)	NOT NULL	Type of pipeline: LOAD or EXPORT
STATUS	VARCHAR2(7)	NOT NULL	Current status of pipeline: STARTED or STOPPED
DESCRIPTION	VARCHAR2(4000)		User specified description for the pipeline.
CREATED	TIMESTAMP(6) WITH TIME ZONE	NOT NULL	Creation time for the pipeline.
LAST_MODIFIED	TIMESTAMP(6) WITH TIME ZONE	NOT NULL	Last modification time for the pipeline.
LAST_EXECUTION	TIMESTAMP(6) WITH TIME ZONE		Last successful execution time of pipeline job.
OPERATION_ID	NUMBER		Is the Operation ID for use with pipeline logs. See Monitor and Troubleshoot Pipelines for information on the USER_LOAD_OPERATIONS view for more information on data load operations.
STATUS_TABLE	VARCHAR2(128)		Status Table for pipeline progress. This is applicable for Load Pipelines.
ORACLE_MAINTAINED	CHAR(1)	NOT NULL	Indicates whether pipeline is Oracle maintained (Y) or not (N).

USER_CLOUD_PIPELINES View

USER_CLOUD_PIPELINES displays the list of pipelines created in the database for the user. The view contains one row for each pipeline.

Column	Datatype	NULL	Description
PIPELINE_ID	NUMBER	NOT NULL	Unique number assigned to the pipeline.
PIPELINE_NAME	VARCHAR2(128)	NOT NULL	Name of the pipeline.
PIPELINE_TYPE	VARCHAR2(7)	NOT NULL	Type of pipeline: LOAD or EXPORT
STATUS	VARCHAR2(7)	NOT NULL	Current status of pipeline: STARTED or STOPPED
DESCRIPTION	VARCHAR2(4000)		User specified description for the pipeline.
CREATED	TIMESTAMP(6) WITH TIME ZONE	NOT NULL	Creation time for the pipeline.
LAST_MODIFIED	TIMESTAMP(6) WITH TIME ZONE	NOT NULL	Last modification time for the pipeline.

Column	Datatype	NULL	Description
LAST_EXECUTION	TIMESTAMP(6) WITH TIME ZONE		Last successful execution time of pipeline job.
OPERATION_ID	NUMBER		Is the Operation ID for use with pipeline logs. See Monitor and Troubleshoot Pipelines for information on the USER_LOAD_OPERATIONS view for more information on data load operations.
STATUS_TABLE	VARCHAR2(128)		Status Table for pipeline progress. This is applicable for Load Pipelines.
ORACLE_MAINTAINED	CHAR(1)	NOT NULL	Indicates whether pipeline is Oracle maintained (Y) or not (N).

DBA_CLOUD_PIPELINE_HISTORY View

DBA_CLOUD_PIPELINE_HISTORY lists the pipeline scheduled jobs in the database. Use this view to monitor the health of the pipeline and detect failures in pipeline job runs.

Column	Datatype	NULL	Description
PIPELINE_ID	NUMBER	NOT NULL	Unique number assigned to the pipeline.
PIPELINE_NAME	VARCHAR2(128)	NOT NULL	Name of the pipeline
OWNER	VARCHAR2(128)	NOT NULL	Owner schema of the pipeline.
PIPELINE_TYPE	VARCHAR2(7)		Type of pipeline: LOAD or EXPORT
JOB_ID	NUMBER	NOT NULL	Unique ID for the pipeline job execution.
JOB_NAME	VARCHAR2(1044)		Unique job name for the pipeline execution.
STATUS	VARCHAR2(30)		Status of the pipeline job.
START_DATE	TIMESTAMP(9) WITH TIME ZONE		Start time of the pipeline job execution.
END_DATE	TIMESTAMP(9) WITH TIME ZON		End time of the pipeline job execution.
INSTANCE_ID	NUMBER		Instance ID where the pipeline job ran.
SESSION_ID	VARCHAR2(128)		Session ID for the job session.
ERROR_NUMBER	NUMBER		Error number in the job (if encountered).
ERROR_MESSAGE	VARCHAR2(4000)		Error message in the job (if encountered).
CPU_USED	INTERVAL DAY(3) TO SECOND(2)		Number of CPU used by the pipeline job run.
JOB_OUTPUT	CLOB		Debug output of the pipeline job run.

USER_CLOUD_PIPELINE_HISTORY View

USER_CLOUD_PIPELINE_HISTORY lists the pipeline scheduled jobs in the database for the user. Use this view to monitor the health of the pipeline and detect failures in pipeline job runs.

Column	Data Type	NULL	Description
PIPELINE_ID	NUMBER	NOT NULL	Unique number assigned to the pipeline.
PIPELINE_NAME	VARCHAR2(128)	NOT NULL	Name of the pipeline
PIPELINE_TYPE	VARCHAR2(7)		Type of pipeline: LOAD or EXPORT

Column	Data Type	NULL	Description
JOB_ID	NUMBER	NOT NULL	Unique ID for the pipeline job execution.
JOB_NAME	VARCHAR2(1044)		Unique job name for the pipeline execution.
STATUS	VARCHAR2(30)		Status of the pipeline job.
START_DATE	TIMESTAMP(9) WITH TIME ZONE		Start time of the pipeline job execution.
END_DATE	TIMESTAMP(9) WITH TIME ZON		End time of the pipeline job execution.
INSTANCE_ID	NUMBER		Instance ID where the pipeline job ran.
SESSION_ID	VARCHAR2(128)		Session ID for the job session.
ERROR_NUMBER	NUMBER		Error number in the job (if encountered).
ERROR_MESSAGE	VARCHAR2(4000)		Error message in the job (if encountered).
CPU_USED	INTERVAL DAY(3) TO SECOND(2)		Number of CPU used by the pipeline job run.
JOB_OUTPUT	CLOB		Debug output of the pipeline job run.

DBA_CLOUD_PIPELINE_ATTRIBUTES View

DBA_CLOUD_PIPELINE_ATTRIBUTES lists the attributes of a pipelines in the database. Attributes control the configuration and behavior of the pipeline jobs.

Column	Datatype	NULL	Description
PIPELINE_ID	NUMBER	NOT NULL	Unique number assigned to the pipeline.
PIPELINE_NAME	VARCHAR2(128)	NOT NULL	Name of the pipeline.
OWNER	VARCHAR2(128)	NOT NULL	Owner schema of the pipeline
PIPELINE_TYPE	VARCHAR2(7)	NOT NULL	Type of pipeline: LOAD or EXPORT
ATTRIBUTE_NAME	VARCHAR2(128)	NOT NULL	Name of the pipeline attribute.
ATTRIBUTE_VALUE	CLOB		Value of the pipeline attribute.
LAST_MODIFIED	TIMESTAMP(6) WITH TIME ZONE	NOT NULL	Last modification time for the pipeline attribute.

USER_CLOUD_PIPELINE_ATTRIBUTES View

USER_CLOUD_PIPELINE_ATTRIBUTES lists the attributes of a pipelines in the database for the user. Attributes control the configuration and behavior of the pipeline.

Column	Datatype	NULL	Description
PIPELINE_ID	NUMBER	NOT NULL	Unique number assigned to the pipeline.
PIPELINE_NAME	VARCHAR2(128)	NOT NULL	Name of the pipeline.
PIPELINE_TYPE	VARCHAR2(7)	NOT NULL	Type of pipeline: LOAD or EXPORT
ATTRIBUTE_NAME	VARCHAR2(128)	NOT NULL	Name of the pipeline attribute.
ATTRIBUTE_VALUE	CLOB		Value of the pipeline attribute.
LAST_MODIFIED	TIMESTAMP(6) WITH TIME ZONE	NOT NULL	Last modification time for the pipeline attribute.

DBMS_CLOUD_AI Views

The `DBMS_CLOUD_AI` package uses the following views.

- DBA_CLOUD_AI_PROFILES View
- USER_CLOUD_AI_PROFILES View
- DBA_CLOUD_AI_PROFILE_ATTRIBUTES View
- USER_CLOUD_AI_PROFILE_ATTRIBUTES View
- DBA_CLOUD_VECTOR _INDEXES View
- USER_CLOUD_VECTOR _INDEXES View
- DBA_CLOUD_VECTOR_INDEX_ATTRIBUTES View
- USER_CLOUD_VECTOR_INDEX_ATTRIBUTES View

> **See Also:**
>
> Use Select AI for Natural Language Interaction with your Database

DBA_CLOUD_AI_PROFILES View

The list of profiles created in the current database can be obtained by querying `DBA_CLOUD_AI_PROFILES`. The view displays the details of AI profiles in the database.

Only the ADMIN user is privileged to access `DBA_CLOUD_AI_PROFILES`. You cannot grant access to other users.

Column	Datatype	NULL	Description
`PROFILE_ID`	`NUMBER`	`NOT NULL`	Unique number assigned to the AI profile
`PROFILE_NAME`	`VARCHAR2(128)`	`NOT NULL`	Name of the AI profile
`OWNER`	`VARCHAR2(128)`	`NOT NULL`	Owner schema of the AI profile
`STATUS`	`VARCHAR2(7)`	`NOT NULL`	Current status of AI profile - `ENABLED` or `DISABLED`
`DESCRIPTION`	`VARCHAR2(4000)`	`NOT NULL`	User specified description for the AI profile
`CREATED`	`TIMESTAMP(6) WITH TIME ZONE`	`NOT NULL`	Creation time for the AI profile.
`LAST_MODIFIED`	`TIMESTAMP(6) WITH TIME ZONE`	`NOT NULL`	Last modification time for the AI profile

USER_CLOUD_AI_PROFILES View

The list of profiles created in the current database can be obtained by querying `USER_CLOUD_AI_PROFILES`. The view displays the details of AI profiles in your schema.

Only the user is privileged to access `USER_CLOUD_AI_PROFILES`. You cannot grant access to other users.

Column	Datatype	NULL	Description
PROFILE_ID	NUMBER	NOT NULL	Unique number assigned to the AI profile
PROFILE_NAME	VARCHAR2(128)	NOT NULL	Name of the AI profile
STATUS	VARCHAR2(7)	NOT NULL	Current status of AI profile - ENABLED or DISABLED
DESCRIPTION	VARCHAR2(4000)	-	User specified description for the AI profile
CREATED	TIMESTAMP(6) WITH TIME ZONE	NOT NULL	Creation time for the AI profile
LAST_MODIFIED	TIMESTAMP(6) WITH TIME ZONE	NOT NULL	Last modification time for the AI profile

The following example displays a sample output from USER_CLOUD_AI_PROFILE view:

```
select PROFILE_NAME, STATUS from USER_CLOUD_AI_PROFILES;

PROFILE_NAME    STATUS
-------------    ---------
OPENAI         ENABLED
GENAI         ENABLED
```

DBA_CLOUD_AI_PROFILE_ATTRIBUTES View

Attributes of AI profile manage the configuration and behavior of the AI profile. The view displays details of all AI profiles in the database.

Only the ADMIN user is privileged to access DBA_CLOUD_AI_PROFILE_ATTRIBUTES. You cannot grant access to other users.

Column	Datatype	NULL	Description
PROFILE_ID	NUMBER	NOT NULL	Unique number assigned to the AI profile
PROFILE_NAME	VARCHAR2(128)	NOT NULL	Name of the AI profile
OWNER	VARCHAR2(128)	NOT NULL	Owner schema of the AI profile
ATTRIBUTE_NAME	VARCHAR2(128)	NOT NULL	Name of the AI profile attribute
ATTRIBUTE_VALUE	CLOB	-	Value of the AI profile attribute
LAST_MODIFIED	TIMESTAMP(6) WITH TIME ZONE	NOT NULL	Last modification time of the AI profile attribute

USER_CLOUD_AI_PROFILE_ATTRIBUTES View

Attributes of AI profile manage the configuration and behavior of the AI profile. The view displays details of AI profiles in your schema.

Only the user is privileged to access USER_CLOUD_AI_PROFILE_ATTRIBUTES. You cannot grant access to other users.

Column	Datatype	NULL	Description
PROFILE_ID	NUMBER	NOT NULL	Unique number assigned to the AI profile
PROFILE_NAME	VARCHAR2(128)	NOT NULL	Name of the AI profile
ATTRIBUTE_NAME	VARCHAR2(128)	NOT NULL	Name of the AI profile attribute

Column	Datatype	NULL	Description
ATTRIBUTE_VALUE	CLOB	-	Value of the AI profile attribute
LAST_MODIFIED	TIMESTAMP(6) WITH TIME ZONE	NOT NULL	Last modification time of the AI profile attribute

The following is an example of a sample output from querying USER_CLOUD_AI_PROFILE_ATTRIBUTES:

```
PROFILE_ID    PROFILE_NAME    ATTRIBUTE_NAME        ATTRIBUTE_VALUE
LAST_MODIFIED
41                OPENAI    provider        openai            28-04-25
11:44:47.068780000 AM ASIA/CALCUTTA
41                OPENAI    credential_name        OPENAI_CRED
28-04-25 11:44:47.068780000 AM ASIA/CALCUTTA
41                OPENAI    object_list        [{"owner":"ADB_USER"}]28-04-25
11:44:47.068780000 AM ASIA/CALCUTTA
41                OPENAI    conversation        true             28-04-25
11:44:47.068780000 AM ASIA/CALCUTTA
33                GENAI     credential_name        GENAI_CRED
21-04-25 9:15:33.322298000 PM ASIA/CALCUTTA
33                GENAI     object_list        [{"owner":"ADB_USER"}] 21-04-25
9:15:33.322298000 PM ASIA/CALCUTTA
33                GENAI     object_list_mode    automated             21-04-25
9:15:33.322298000 PM ASIA/CALCUTTA
33                GENAI     model                gpt-3.5-turbo-1106
21-04-25 9:15:33.322298000 PM ASIA/CALCUTTA
33                GENAI     provider        openai            21-04-25
9:15:33.322298000 PM ASIA/CALCUTTA
```

DBA_CLOUD_VECTOR _INDEXES View

You can obtain a list of vector indexes created in the current database by querying DBA_CLOUD_VECTOR_INDEXES. The view displays the details of all available vector indexes in the entire database.

Only the administrator is privileged to access DBA_CLOUD_VECTOR_INDEXES. You cannot grant access to other users.

Column	Datatype	NULL	Description
INDEX_ID	NUMBER	NOT NULL	Unique number assigned to the index
INDEX_NAME	VARCHAR2(128)	NOT NULL	Name of the index
STATUS	VARCHAR2(7)	-	Current status of index: ENABLED or DISABLED
DESCRIPTION	CLOB	-	User specified description for the index
CREATED	TIMESTAMP(6) WITH TIME ZONE	NOT NULL	Creation date and time of the index
LAST_MODIFIED	TIMESTAMP(6) WITH TIME ZONE	NOT NULL	Last modification date and time of the index

USER_CLOUD_VECTOR _INDEXES View

You can obtain a list of vector indexes you created in the current database by querying `USER_CLOUD_VECTOR_INDEXES`. The view displays the details of all vector indexes that you created in your schema or a PDB.

You have the privilege to access `USER_CLOUD_VECTOR_INDEXES`. You cannot grant access to other users.

Column	Datatype	NULL	Description
INDEX_ID	NUMBER	NOT NULL	Unique number assigned to the index
INDEX_NAME	VARCHAR2(128)	NOT NULL	Name of the index
STATUS	VARCHAR2(7)	-	Current status of index: `ENABLED` or `DISABLED`
DESCRIPTION	VARCHAR2(4000)	-	User specified description for the index
CREATED	TIMESTAMP(6) WITH TIME ZONE	NOT NULL	Creation date and time of the index
LAST_MODIFIED	TIMESTAMP(6) WITH TIME ZONE	NOT NULL	Last modification date and time of the index

DBA_CLOUD_VECTOR_INDEX_ATTRIBUTES View

Query the `DBA_CLOUD_VECTOR_INDEX_ATTRIBUTES` view to display details about the attributes of vector indexes across the entire database. Attributes of the view manage the configuration and behavior of the vector index.

Only the administrator is privileged to access `DBA_CLOUD_VECTOR_INDEX_ATTRIBUTES`. You cannot grant access to other users.

Column	Datatype	NULL	Description
INDEX_ID	NUMBER	NOT NULL	Unique number assigned to the vector index
INDEX_NAME	VARCHAR2(128)	NOT NULL	Name of the vector index
OWNER	VARCHAR2(128)	NOT NULL	Owner schema of the vector index
ATTRIBUTE_NAME	VARCHAR2(128)	NOT NULL	Name of the vector index attribute
ATTRIBUTE_VALUE	CLOB	-	Value for the vector index `attribute_name`
LAST_MODIFIED	TIMESTAMP(6) WITH TIME ZONE	NOT NULL	Last modification date and time of the vector index

USER_CLOUD_VECTOR_INDEX_ATTRIBUTES View

Query the `USER_CLOUD_VECTOR_INDEX_ATTRIBUTES` view to display details about the attributes of vector indexes that you created in your schema or PDB. Attributes of the view manage the configuration and behavior of the vector index.

You have the privilege to access `USER_CLOUD_VECTOR_INDEX_ATTRIBUTES`. You cannot grant access to other users.

Column	Datatype	NULL	Description
INDEX_ID	NUMBER	NOT NULL	Unique number assigned to the vector index

Column	Datatype	NULL	Description
INDEX_NAME	VARCHAR2(128)	NOT NULL	Name of the vector index
ATTRIBUTE_NAME	VARCHAR2(128)	NOT NULL	Name of the vector index attribute
ATTRIBUTE_VALUE	CLOB	-	Value for the vector index attribute_name
LAST_MODIFIED	TIMESTAMP(6) WITH TIME ZONE	NOT NULL	Last modification date and time of the vector index

DBMS_CLOUD_FUNCTION Views

The DBMS_CLOUD_FUNCTION package uses the following views.

- DBA_CLOUD_FUNCTION View
  The DBA_CLOUD_FUNCTION view lists all the functions created in your Autonomous Database.
- DBA_CLOUD_FUNCTION_CATALOG View
  The DBA_CLOUD_FUNCTION_CATALOG view lists all the catalogs created in an Autonomous Database instance.
- USER_CLOUD_FUNCTION View
  The USER_CLOUD_FUNCTION view lists all the functions created by the user.
- USER_CLOUD_FUNCTION_CATALOG View
  The USER_CLOUD_FUNCTION_CATALOG view lists all the catalogs created by the user.
- USER_CLOUD_FUNCTION_ERRORS View
  The USER_CLOUD_FUNCTION_ERRORS view lists the errors generated during the connection validation to the remote library location.

DBA_CLOUD_FUNCTION View

The DBA_CLOUD_FUNCTION view lists all the functions created in your Autonomous Database.

Column	Datatype	Description
OWNER	VARCHAR2(128)	Owner of the catalog
CATALOG_NAME	VARCHAR2(128)	Name of the catalog
APPLICATION_ID	VARCHAR2(2000)	OCI Application ID or Azure Application ID
CLOUD_FUNCTION_NAME	VARCHAR2(128)	Remote cloud function name
FUNCTION_NAME	VARCHAR2(128)	Function Name
FUNCTION_ID	VARCHAR2(2000)	OCI Function ID or AWS Function ARN or Azure Function ID or GCP Function ID
API_TYPE	VARCHAR2(128)	API type (Function or procedure)

DBA_CLOUD_FUNCTION_CATALOG View

The DBA_CLOUD_FUNCTION_CATALOG view lists all the catalogs created in an Autonomous Database instance.

Column	Datatype	Description
OWNER	VARCHAR2(128)	Owner of the catalog

Column	Datatype	Description
CATALOG_NAME	VARCHAR2(128)	Name of the catalog
CLOUD_PROVIDER	VARCHAR2(128)	OCI or AWS or Azure or GCP

USER_CLOUD_FUNCTION View

The USER_CLOUD_FUNCTION view lists all the functions created by the user.

Column	Datatype	Description
CATALOG_NAME	VARCHAR2(128)	Name of the catalog
APPLICATION_NAME	VARCHAR2(128)	OCI Application Name
APPLICATION_ID	VARCHAR2(2000)	OCI Application ID or Azure Application ID
CLOUD_FUNCTION_NAME	VARCHAR2(128)	Remote cloud function name
FUNCTION_NAME	VARCHAR2(128)	Function Name
FUNCTION_ID	VARCHAR2(2000)	OCI Function ID or AWS Function ARN or Azure Function ID or GCP Function ID
API_TYPE	VARCHAR2(128)	API type (Function or procedure)

USER_CLOUD_FUNCTION_CATALOG View

The USER_CLOUD_FUNCTION_CATALOG view lists all the catalogs created by the user.

Column	Datatype	Description
CATALOG_NAME	VARCHAR2(128)	Name of the catalog
CLOUD_PROVIDER	VARCHAR2(128)	OCI or AWS or Azure or GCP

USER_CLOUD_FUNCTION_ERRORS View

The USER_CLOUD_FUNCTION_ERRORS view lists the errors generated during the connection validation to the remote library location.

Column	Datatype	Description
CATALOG_NAME	VARCHAR2(128)	Name of the catalog
CREATE_TIMESTAMP	TIMESTAMP	Timestamp when error was generated
log_error_msg	CLOB	The error message

DBMS_CLOUD_LINK Views

The DBMS_CLOUD_LINK package uses the following views.

- DBA_CLOUD_LINK_ACCESS and ALL_CLOUD_LINK_ACCESS Views
- DBA_CLOUD_LINK_AUTHORIZATIONS View
- DBA_CLOUD_LINK_PRIVS and USER_CLOUD_LINK_PRIVS Views
- DBA_CLOUD_LINK_REGISTRATIONS and ALL_CLOUD_LINK_REGISTRATIONS Views
- DBA_CLOUD_LINK_SERVICE_MAPPINGS View

- V$CLOUD_LINK_ACCESS_STATS and GV$CLOUD_LINK_ACCESS_STATS Views

DBA_CLOUD_LINK_ACCESS and ALL_CLOUD_LINK_ACCESS Views

Displays details of the registered data sets an Autonomous Database instance is allowed to access.

Only the ADMIN user and users with the `PDB_DBA` role are privileged to access `DBA_CLOUD_LINK_ACCESS`. You cannot grant access to other users.

Column	Datatype	NULL	Description
NAMESPACE	VARCHAR2(128)	NOT NULL	Namespace of registered data set.
NAME	VARCHAR2(128)	NOT NULL	Name of registered data set.
CREATED	TIMESTAMP(6)	NOT NULL	Time/Date when the data set was created.
LAST_MODIFIED	TIMESTAMP(6)	NOT NULL	Time when metadata of data set was last modified
AUTHORIZATION_REQUIRED	VARCHAR2(1)	NOT NULL	Data set requires additional authorization. Possible values are: "T" or "F".
DATA_SET_OWNER	VARCHAR2(128)	NULL	Data set owner.

DBA_CLOUD_LINK_AUTHORIZATIONS View

Displays details of which databases are authorized to access data sets. This applies to data sets that require additional authorization.

Only the ADMIN user and users with the `PDB_DBA` role are privileged to access `DBA_CLOUD_LINK_AUTHORIZATIONS`. You cannot grant access to other users.

Column	Datatype	NULL	Description
NAMESPACE	VARCHAR2(128)	NOT NULL	Namespace of registered data set.
NAME	VARCHAR2(128)	NOT NULL	Name of registered data set.
CLOUD_LINK_DATABASE_ID	VARCHAR2(40)	NOT NULL	Database ID of database that has been authorized for access.

DBA_CLOUD_LINK_PRIVS and USER_CLOUD_LINK_PRIVS Views

Displays details of the Cloud Link specific privileges, `REGISTER`, `READ`, `AUTHORIZE`, granted to specific users or to the current user.
`USER_CLOUD_LINK_PRIVS` columns are the same as those in `DBA_CLOUD_LINK_PRIVS`. (except for `GRANTEE`).

Only the ADMIN user and users with the `PDB_DBA` role are privileged to read from `DBA_CLOUD_LINK_PRIVS`. You cannot grant access to other users.

The privilege to read from `USER_CLOUD_LINK_PRIVS` view is granted to the ADMIN user, to the `PDB_DBA` role, and to the role `DWROLE`.

Column	Datatype	NULL	Description
GRANTEE	VARCHAR2(128)	NOT NULL	Name of user to whom privilege is granted. This column is not available in `USER_CLOUD_LINK_PRIVS`.

Column	Datatype	NULL	Description
PRIVILEGE	VARCHAR2(9)	NOT NULL	Privilege, one of: REGISTER, READ, AUTHORIZE, granted to user.
SCOPE	VARCHAR2(14)	NOT NULL	Scope of the privilege.

DBA_CLOUD_LINK_REGISTRATIONS and ALL_CLOUD_LINK_REGISTRATIONS Views

Displays details of all the data set registrations on an Autonomous Database instance.

Only the ADMIN user and users with the PDB_DBA role are privileged to access DBA_CLOUD_LINK_REGISTRATIONS. You cannot grant access to other users.

Column	Datatype	NULL	Description
NAMESPACE	VARCHAR2(128)	NOT NULL	Namespace of registered data set.
NAME	VARCHAR2(128)	NOT NULL	Name of registered data set.
ID	RAW(16)	NOT NULL	Unique ID of registered data set.
OWNER	VARCHAR2(128)	NOT NULL	Owner of object (table or view owner).
OBJECT_NAME	VARCHAR2(128)	NOT NULL	Name of object.
SCOPE	CLOB	NOT NULL	JSON document with details of who is allowed to access the registered data set.
DESCRIPTION	CLOB	NOT NULL	Description of registered data set.
CREATED	TIMESTAMP(6)	NOT NULL	Time/Date when the data set was created.
LAST_MODIFIED	TIMESTAMP(6)	NOT NULL	Time when metadata of data set was last modified
AUTHORIZATION_REQUIRED	VARCHAR2(1)	NOT NULL	Data set requires additional authorization. Possible values are: "T" or "F".
DATA_SET_OWNER	VARCHAR2(128)	NULL	Data set owner.

DBA_CLOUD_LINK_SERVICE_MAPPINGS View

Displays details of all service mappings for Cloud Links consumer databases. Each service mapping consists of a Cloud Link database ID and a database service.

Only the ADMIN user and users with the PDB_DBA role are privileged to access DBA_CLOUD_LINK_SERVICE_MAPPINGS. You cannot grant access to other users.

Column	Datatype	NULL	Description
CLOUD_LINK_DATABASE_ID	VARCHAR2(40)	NOT NULL	Database Identifier of consumer database or "ANY".
SERVICE_NAME	VARCHAR2(10)	NULL	Database service name (HIGH, MEDIUM, or LOW).

V$CLOUD_LINK_ACCESS_STATS and GV$CLOUD_LINK_ACCESS_STATS Views

Displays access to each registered data set owned by an Autonomous Database instance.

V$CLOUD_LINK_ACCESS_STATS tracks the elapsed time, DB CPU time, the number of rows retrieved, the number of executions, and the unique database identifier of the database accessing a registered data set.

In addition to the V$CLOUD_LINK_ACCESS_STATS information, the GV$CLOUD_LINK_ACCESS_STATS view contains an extra column named INST_ID of data type NUMBER. The INST_ID column displays the instance from which the associated V$ view information was obtained. The INST_ID column can be used as a filter to retrieve V$CLOUD_LINK_ACCESS_STATS information from a subset of available instances.

The unique database identifier does not disclose the OCIDs (Tenancy, Compartment, Database). This information can be used for auditing an Autonomous Database.

Column	Datatype	NULL	Description
INST_ID	NUMBER	NOT NULL	Instance ID
OWNER	VARCHAR2(128)	NOT NULL	Owner of object(table or view owner)
OBJECT_NAME	VARCHAR2(128)	NOT NULL	Name of object
DATA_SET_NAMESPACE	VARCHAR2(128)	NOT NULL	Namespace of registered data set
DATA_SET_NAME	VARCHAR2(128)	NOT NULL	Name of registered data set
SQL_ID	VARCHAR2(13)	NOT NULL	SQL ID of local SQL accessing a registered data set
LAST_ACTIVE_TIME	DATE	NOT NULL	Time when local SQL was last active
CLOUD_LINK_DATABASE_ID	VARCHAR2(512)	NOT NULL	Database identifier of database accessing a registered data set
ELAPSED_TIME	NUMBER	NOT NULL	Elapsed time (in microseconds) of local SQL
CPU_TIME	NUMBER	NOT NULL	CPU time (in microseconds) of local SQL
USER_IO_WAIT_TIME	NUMBER	NOT NULL	User I/O Wait time (in microseconds) of local SQL
EXECUTIONS	NUMBER	NOT NULL	Number of executions of local SQL
FETCHES	NUMBER	NOT NULL	Number of fetches of local SQL
ROWS_PROCESSED	NUMBER	NOT NULL	Number of rows retrieved by local SQL
CON_ID	NUMBER	NOT NULL	Container ID

DBMS_DATA_ACCESS Views

Describes the Autonomous Database dynamic performance views you can use with the DBMS_DATA_ACCESS package to monitor Table Hyperlink URL usage.

- V$DATA_ACCESS_URL_STATS and GV$DATA_ACCESS_URL_STATS Views

V$DATA_ACCESS_URL_STATS and GV$DATA_ACCESS_URL_STATS Views

Displays access to each Table Hyperlink owned by an Autonomous Database instance.

When a Table Hyperlink is accessed, CPU time, User I/O wait time, are updated in the V$DATA_ACCESS_URL_STATS view.

The GV$DATA_ACCESS_URL_STATS view contains the additional column named INST_ID of data type NUMBER. The INST_ID column displays the instance from which the associated V$ view information was obtained. The INST_ID column can be used as a filter to retrieve V$DATA_ACCESS_URL_STATS information from a subset of available instances.

Column	Datatype	NULL	Description
INST_ID	NUMBER	NOT NULL	Instance ID
SQL_ID	VARCHAR2(13)	NOT NULL	SQL ID of SQL accessing data when a Table Hyperlink is accessed
PRE_AUTH_URL_ID	VARCHAR2(128)	NOT NULL	Table Hyperlink identifier
ELAPSED_TIME	NUMBER	NOT NULL	Elapsed time (in microseconds) of local SQL
CPU_TIME	NUMBER	NOT NULL	CPU time (in microseconds) of local SQL
USER_IO_WAIT_TIME	NUMBER	NOT NULL	User I/O Wait time (in microseconds) of local SQL
EXECUTIONS	NUMBER	NOT NULL	Number of executions of local SQL
FETCHES	NUMBER	NOT NULL	Number of fetches of local SQL
ROWS_PROCESSED	NUMBER	NOT NULL	Number of rows retrieved by local SQL
USER_ID	VARCHAR2(128)		Application User ID defined for Table Hyperlink
CON_ID	NUMBER	NOT NULL	Container ID

DBA_PROXY_SQL Views

Describes the Autonomous Database DBA_PROXY_SQL views.

- DBA_PROXY_SQL_ACTIONS View
- DBA_PROXY_SQL_MODULES View

DBA_PROXY_SQL_ACTIONS View

Displays the action names for which query offload is configured. You must query this view from a session that is not enabled for query offload on an Autonomous Database instance.

Column	Datatype	NULL	Description
ACTION_NAME	VARCHAR2(4000)		List of actions

DBA_PROXY_SQL_MODULES View

Displays the module names for which query offload is configured. You must query this view from a session that is not enabled for query offload on an Autonomous Database instance.

Column	Datatype	NULL	Description
MODULE_NAME	VARCHAR2(4000)		List of modules

Lockdown Errors View

Lists details for the DBA_LOCKDOWN_ERRORS view.

- DBA_LOCKDOWN_ERRORS View
  The DBA_LOCKDOWN_ERRORS view provides information about SQL statements where lockdown profile errors were generated.

DBA_LOCKDOWN_ERRORS View

The DBA_LOCKDOWN_ERRORS view provides information about SQL statements where lockdown profile errors were generated.

Column	Datatype	NULL	Description
LAST_OCCURRENCE	TIMESTAMP(6)		Timestamp of latest execution failure.
OCCURRENCE_COUNT	NUMBER		Total occurrences of execution failure.
ERROR_NUM	NUMBER		Error number
ERROR_MSG	VARCHAR2(4000)		Error message
RULE_TYPE	VARCHAR2(128)		Lockdown Profile Rule type: STATEMENT, OPTIONS, FEATURE, EXECUTION
SQL_TEXT	CLOB		SQL text
SQL_ID	VARCHAR2(13)		SQL_ID of statement that hit lockdown error.
SESSION_ID	NUMBER		Session identifier.
SESSION_SERIAL	NUMBER		Session serial number.
SCHEMA	VARCHAR2(128)		User's current schema.
MODULE	VARCHAR2(64)		Name of the currently executing module.
ACTION	VARCHAR2(64)		Name of the currently executing action
CLIENT_ID	VARCHAR2(64)		Client identifier of the session.
USER#	NUMBER	NOT NULL	ID number of the user.
USERNAME	VARCHAR2(128)	NOT NULL	Name of the user.
LUSER#	NUMBER	NOT NULL	ID number of the Logon user
LUSERNAME	VARCHAR2(128)	NOT NULL	Name of the logon user.

Oracle Cloud Infrastructure Logging Interface Views

The Oracle Cloud Infrastructure (OCI) logging interface views enable authorized users to access log data through a set of views.

Oracle Cloud Infrastructure (OCI) logging is the central logging solution for OCI services, including VCN flow logs, load balancer logs, object storage logs, and others within the added capability to include custom logs. The logs are not stored in Autonomous Database. Instead, the logs are dynamically retrieved based on various predicates provided by the user, including date range, log group, log name, and others.

- Oracle Cloud Infrastructure Logging Interface Overview
- OCI_LOG_LIST View
- OCI_LOG_DATA View
- OCI_VCN_FLOWLOGS View
- OCI_LBLOG_ACCESS View
- OCI_LBLOG_ERRORS View
- Oracle Cloud Infrastructure Logging Interface Overview
  Providing access to the log data inside an Autonomous Database in relational format provides a useful access method, along with the ability to get the logging data in relational

or JSON format and enrich the analytics by joining with other data they may have in the Autonomous Database.

- OCI_LOG_LIST View
  Displays a list of all logs in the tenant. The view does not show log data, but provides the metadata about the logs in a tenancy. This metadata is used to properly specify parameter values for other logging view. If there is no logs under a log group, that log group will not be included in this view.
- OCI_LOG_DATA View
  The common view that provides access to all log types supported by the OCI Logging Service. The log content is in the `DATA` column.
- OCI_VCN_FLOWLOGS View
  View that provides details about traffic that passes through the user's VCN. This log enables users to audit traffic and troubleshoot security lists. Each flow log record reflects logged traffic in one direction of a connection between two endpoints. For example, a single TCP connection, you may have two records in the capture window: one for ingress traffic, and one for egress traffic.
- OCI_LBLOG_ACCESS View
  View that provides load balancer access logs capturing detailed information about requests sent to a load balancer. Each entry contains the time the request was received, client, intermediate HTTP proxy IP addresses, and times used by at the load balancer and backend to process the request.
- OCI_LBLOG_ERRORS View
  View that provides load balancer error logs capturing detailed information about requests related to troubleshooting and monitoring. Each entry contains information such as the time the request was received, error type, and additional details of the specific error.

Oracle Cloud Infrastructure Logging Interface Overview

Providing access to the log data inside an Autonomous Database in relational format provides a useful access method, along with the ability to get the logging data in relational or JSON format and enrich the analytics by joining with other data they may have in the Autonomous Database.

Oracle logging service provides access to logs from Oracle Cloud Infrastructure (OCI) resources. These logs include critical diagnostic information that describes how resources are performing and being accessed. There are varying log types for each Oracle service. To learn more about the logging service and supported OCI services, see Oracle Cloud Infrastructure Logging Overview.

Through the Oracle Cloud Infrastructure Logging Interface, log data can now be accessed through an Autonomous Database in relational format. Users can query different log data in OCI across all compartments and regions. The implemented OCI views use the `OCI$RESOURCE_PRINCIPAL` credential.

To use these views, the following steps need to be done.

1. An Administrator enables use of the `OCI$RESOURCE_PRINCIPAL` credential in the database.
2. Users create a new dynamic group for the `OCI$RESOURCE_PRINCIPAL` credential.

To learn more about enabling the credential, setting up a dynamic group, and creating policy statements, see Use Resource Principal to Access Oracle Cloud Infrastructure Resources.

As an Administrator, enable the resource principal.

```
EXEC DBMS_CLOUD_ADMIN.ENABLE_RESOURCE_PRINCIPAL();
```

The dynamic group must have the following rules that include the resource ID of the instance.

```
resource.id = '<resource.id>'
```

There are two policy statements required to access logging views.

```
-- For OCI Logging Views
Allow dynamic-group <group-name> to use logging-family in tenancy
Allow dynamic-group <group-name> to use compartments in tenancy
```

All the views have mandatory and optional predicate values (column names).

Predicate values are as follows:

> **Note:**
>
> `OCI_LOG_LIST` view does not have mandatory columns. If the `REGION` value is not provided, the view returns data for the home region.

- `REGION` (Mandatory)
- `COMPARTMENT_ID` (Mandatory)
- `LOG_GROUP_ID` (Mandatory)
- `LOG_ID` (Optional) - If the value is provided, query results will use the `LOG_ID`value. If no value provided, the value will be returned from the server.
- `START_TIME` (Optional) - If no value provided, return values are for past 5 minutes in GMT.
- `END_TIME` (Optional) - If no value provided, return values are for past 5 minutes in GMT.
- `SEARCH_CRITERIA` (Optional) - If no value provided, null will be returned.

Since there can be different log types within a single log group, it's recommended to specify `LOG_ID` in a predicate when using `OCI_VCN_FLOWLOGS`, `OCI_LBLOG_ACCESS`, and `OCI_LBLOG_ERRORS` views.

Notes and Restrictions:

- The common view `OCI_LOG_DATA` supports all log types, so `COMPARTMENT_ID`, `LOG_GROUP_ID`, and `REGION` are sufficient. Likewise, if a user uses the `LOG_ID` of a load balancer with the `OCI_VCN_FLOWLOGS` view, it will either not return correct data or any data at all.
- For the `OCI_LOG_LIST` view, the `select * from OCI_LOG_LIST` will shows data only for the home region. If you need data for another region, you will have to use a `REGION` equality predicate in the WHERE clause.
- If the `OCI_LOG_DATA`, `OCI_VCN_FLOWLOGS`, `OCI_LBLOG_ACCESS`, and `OCI_LBLOG_ERRORS`, the `select * from <view>` gives an `ORA-20000: compartment_id,log_group_id and region`

`are mandatory predicates, please provide valid values for them as equality predicate in the WHERE clause.` provide the required predicate values.

- The framework supports querying up to past 7 days (from current time) log data to query from the supported views. Any values for the predicates `START_TIME` and `END_TIME` predicates outside this range will result in `ORA-20000: start_time and end_time should be from current_timestamp to current_timestamp-7.`
- If you do not provide the listed mandatory predicates, you will see an `ORA-20000: compartment_id,log_group_id and region are mandatory predicates, please provide valid values for them as equality predicate in the WHERE clause`. Other columns can be used as predicates of any type (`=`, `IN`). The mandatory predicates can be only used once in a query, and they cannot be used multiple times with `AND` or `NOT`.
- Values for `START_TIME` and `END_TIME` have to be given in `DD-MM-YY HH24:MI:SS` format. If not provided in this format, the query will not return any data.

OCI_LOG_LIST View

Displays a list of all logs in the tenant. The view does not show log data, but provides the metadata about the logs in a tenancy. This metadata is used to properly specify parameter values for other logging view. If there is no logs under a log group, that log group will not be included in this view.

Column	Datatype	Description
REGION	VARCHAR2	OCI region name
COMPARTMENT_ID	VARCHAR2	Compartment OCID
LOG_GROUP_ID	VARCHAR2	`LOG_GROUP` OCID
LOG_GROUP_NAME	VARCHAR2	Log group name.
LOG_ID	VARCHAR2	`LOG` OCID
LOG_NAME	VARCHAR2	Log name
LOG_SERVICE	VARCHAR2	Log service

OCI_LOG_DATA View

The common view that provides access to all log types supported by the OCI Logging Service. The log content is in the `DATA` column.

Column	Datatype	Description
DATE_TIME	TIMESTAMP	Same as the oracle.ingestedtime field
ID	VARCHAR2	A random UUID for each log entry
SOURCE	VARCHAR2	Resource name that generated the log message
SPEC_VERSION	NUMBER	Version of the `CloudEvents` specification this log message uses
LOG_TIME	TIMESTAMP	Log message generation time
TYPE	VARCHAR2	Log message type
COMPARTMENT_ID	VARCHAR2	Compartment OCID
INGESTED_TIME	TIMESTAMP	Log message ingest time

Column	Datatype	Description
LOG_GROUP_ID	VARCHAR2	Log group OCID
LOG_ID	VARCHAR2	Log OCID
TENANT_ID	VARCHAR2	Tenant OCID of log object owner
DATA	CLOB	Log message
ORACLE_DETAILS	CLOB	Oracle-specific metadata

OCI_VCN_FLOWLOGS View

View that provides details about traffic that passes through the user's VCN. This log enables users to audit traffic and troubleshoot security lists. Each flow log record reflects logged traffic in one direction of a connection between two endpoints. For example, a single TCP connection, you may have two records in the capture window: one for ingress traffic, and one for egress traffic.

Column	Datatype	Description
REGION	VARCHAR2	OCI region name
DATE_TIME	TIMESTAMP	Same as the oracle.ingestedtime field
ID	VARCHAR2	A random UUID unique to each log entry
ACTION	VARCHAR2	Possible values of **ACCEPT** or **REJECT**
FLOW_ID	VARCHAR2	Hash of key fields (source and destination addresses, ports, and protocol)
STATUS	VARCHAR2	Possible values of **OK**, **NODATA**, or **SKIPDATA**
END_TIME	TIMESTAMP	Capture window end time
PACKETS	NUMBER	Number of packets recorded in capture window
VERSION	VARCHAR2	Flow log record schema number
BYTES_OUT	NUMBER	Number of bytes recorded in capture window
PROTOCOL	NUMBER	IANA protocol number
START_TIME	TIMESTAMP	Capture window start time
SOURCE_PORT	NUMBER	Source IANA port number
PROTOCOL_NAME	VARCHAR2	IANA protocol name
SOURCE_ADDRESS	VARCHAR2	IP address of the source in IPv4 or IPv6 notation
DESTINATION_PORT	NUMBER	Destination IANA port number
DESTINATION_ADDRESS	VARCHAR2	IP address of the destination in IPv4 or IPv6 notation
TIME	TIMESTAMP	Same as START_TIME
TYPE	VARCHAR2	Log category
LOG_ID	VARCHAR2	Log OCID
TENANT_ID	VARCHAR2	Tenant OCID
VNIC_OCID	VARCHAR2	VNIC OCID

Column	Datatype	Description
LOG_GROUP_ID	VARCHAR2	Log group OCID
INGESTED_TIME	TIMESTAMP	Time log ingested by OCI Logging
COMPARTMENT_ID	VARCHAR2	Log group compartment OCID
VNIC_SUBNET_OCID	VARCHAR2	VNIC subnet OCID
VNIC_COMPARTMENT_ID	VARCHAR2	VNIC compartment OCID
SOURCE	VARCHAR2	Resource name that generated log message
SPEC_VERSION	NUMBER	Version of the `CloudEvents` specification this log message uses

OCI_LBLOG_ACCESS View

View that provides load balancer access logs capturing detailed information about requests sent to a load balancer. Each entry contains the time the request was received, client, intermediate HTTP proxy IP addresses, and times used by at the load balancer and backend to process the request.

Column	Datatype	Description
REGION	VARCHAR2	OCI region name
DATE_TIME	TIMESTAMP	Same as the oracle.ingestedtime field
ID	VARCHAR2	A random UUID unique to each log entry
BACKEND_ADDRESS	VARCHAR2	IP address and port number of the backend server which processed the client request
BACKEND_CONNECT_TIME	NUMBER	Time spent (in seconds, with millisecond precision) to establish backend server connection
BACKEND_PROCESSING_TIME	NUMBER	Total time taken (in seconds, with millisecond precision) from the load balancer establishing a connection to a backend until it completes
BACKEND_STATUS_CODE	NUMBER	Status code of the response from the target
CLIENT_ADDRESS	VARCHAR2	IP address and port number of the requesting client
FORWARDED_FOR_ADDRESS	VARCHAR2	IP address of the client and http proxies between client and load balancer
HOST	VARCHAR2	Domain name which resolves to IP address assigned to the load balancer
LB_STATUS_CODE	NUMBER	Load balancer status code
LISTENER_NAME	VARCHAR2	Listener which received the incoming traffic request on the load balancer's IP address
RECEIVED_BYTES	NUMBER	Total size of the request (in bytes) received from the client
REQUEST	VARCHAR2	Request line received from the client

Column	Datatype	Description
REQUEST_PROCESSING_TIME	NUMBER	Total time taken (in seconds, with millisecond precision) from the load balancer receiving the request from the client until load balancer completes sending response to the client.
ROUTING_RULES_ENGINE_ERRORS	NUMBER	Routing rule engine error during policy evaluation of the request with a 0 (no error) or 1 (error). If an error occurred, requests are forwarded to the default backend attached to the listener.
ROUTING_RULES_MATCHED_RULE	VARCHAR2	Routing policy rule name, which was matched for this specific request
ROUTING_RULES_RULE_HITS	NUMBER	Number of routing rules evaluated to true for the request
ROUTING_RULES_RULE_MISSES	NUMBER	Number of routing rules evaluated to false for the request
SENT_BYTES	NUMBER	Total size of the request (in bytes) sent to the client
SSL_CIPHER	VARCHAR2	Negotiated SSL cipher between the client and the load balancer
SSL_PROTOCOL	VARCHAR2	Negotiated SSL protocol between the client and the load balancer
TIMESTAMP	TIMESTAMP	Log entry generation time
USER_AGENT	VARCHAR2	User agent used to send the request to the load balancer
TIME	TIMESTAMP	Log entry generation time
TYPE	VARCHAR2	Log category
LOG_ID	VARCHAR2	Log OCID
TENANT_ID	VARCHAR2	Tenant OCID
LOG_GROUP_ID	VARCHAR2	Log group OCID
RESOURCE_ID	VARCHAR2	Resource OCID
INGESTED_TIME	TIMESTAMP	Time log ingested by OCI Logging
COMPARTMENT_ID	VARCHAR2	Log group compartment OCID
SOURCE	VARCHAR2	Resource name that generated log message
SUBJECT	VARCHAR2	Subject of the log
SPEC_VERSION	NUMBER	Version of the `CloudEvents` specification this log message uses

OCI_LBLOG_ERRORS View

View that provides load balancer error logs capturing detailed information about requests related to troubleshooting and monitoring. Each entry contains information such as the time the request was received, error type, and additional details of the specific error.

Column	Datatype	Description
REGION	VARCHAR2	OCI region name

Column	Datatype	Description
DATE_TIME	TIMESTAMP	Same as the oracle.ingestedtime field
ID	VARCHAR2	A random UUID unique to each log entry
ERROR_LOG_TYPE	VARCHAR2	Log type
ERROR_LOG_ERROR_DETAILS	VARCHAR2	Detailed description of the error message
DATA_TIMESTAMP	TIMESTAMP	Log entry generation time
TIME	TIMESTAMP	Log entry generation time
TYPE	VARCHAR2	Log category
LOG_ID	VARCHAR2	Log OCID
TENANT_ID	VARCHAR2	Tenant OCID
LOG_GROUP_ID	VARCHAR2	Log group OCID
RESOURCE_ID	VARCHAR2	Resource OCID
INGESTED_TIME	TIMESTAMP	Time log ingested by OCI Logging
COMPARTMENT_ID	VARCHAR2	Log group compartment OCID
SOURCE	VARCHAR2	Resource name that generated log message
SUBJECT	VARCHAR2	Subject of the log
SPEC_VERSION	NUMBER	Version of the CloudEvents specification this log message uses

Real Application Testing Capture Replay Views

Oracle Real Application Testing capture replay views.

- DBA_CAPTURE_REPLAY_HISTORY View
  The DBA_CAPTURE_REPLAY_HISTORY view provides historical information for all your workload captures and replays.
- DBA_CAPTURE_REPLAY_STATUS View
  The DBA_CAPTURE_REPLAY_STATUS view provides the most recent status for your workload. The values in the view are updated when you execute FINISH_WORKLOAD_CAPTURE, CANCEL_WORKLOAD_CAPTURE or REPLAY_WORKLOAD.

DBA_CAPTURE_REPLAY_HISTORY View

The DBA_CAPTURE_REPLAY_HISTORY view provides historical information for all your workload captures and replays.

Column	Datatype	Description
NAME	VARCHAR2(32)	Name of the workload capture or replay
UNIQUE_NAME	VARCHAR2(64)	Internal unique name for capture or replay
STATUS	VARCHAR2(8)	Workload capture or replay status
DATABASE_NAME	VARCHAR2(128)	Name of the database
DATABASE_GUID	RAW(16)	Unique id of the database
TYPE	VARCHAR2(10)	Type (capture or replay)
START_TIME	DATE	Start time of the workload capture or replay.

Column	Datatype	Description
END_TIME	DATE	End time of the workload capture or replay.
REPORT_URL	VARCHAR2(1000)	PAR URL path for capture or replay report.

DBA_CAPTURE_REPLAY_STATUS View

The DBA_CAPTURE_REPLAY_STATUS view provides the most recent status for your workload. The values in the view are updated when you execute FINISH_WORKLOAD_CAPTURE, CANCEL_WORKLOAD_CAPTURE or REPLAY_WORKLOAD.

Column	Datatype	Description
STATE	VARCHAR2(32)	Provides the current status for a workload capture or replay. For a workload capture, status value could be one of these: • CAPTURE *capture_name* STARTED • FINISHING CAPTURE • UPLOADING CAPTURE FILES • CANCELING CAPTURE • CANCEL CAPTURE *capture_name* COMPLETED • FINISH CAPTURE *capture_name* COMPLETED For a workload replay, status value could be one of these: • DOWNLOADING CAPTURE FILES • PROCESSING CAPTURE • INITIALIZING REPLAY • PREPARING REPLAY • STARTING REPLAY • UPLOADING REPLAY FILES • REPLAY *replay_name* COMPLETED
PROGRESS	VARCHAR2(8)	Percentage completed for download or upload.

Stripe Views on Autonomous Database

Stripe is an online payment processing and credit card processing platform for businesses. Users can query views created on top of Stripe APIs using the DBMS_CLOUD package to get Stripe information such as invoices, subscriptions, and customers.

For more information on Stripe, see the Stripe website.

- Prerequisites to Use Stripe Views
  To use Stripe with Autonomous Database, set Network ACL for accessing Stripe and create a credential that allows access to Stripe.
- STRIPE_ACCOUNTS View
  STRIPE_ACCOUNTS lists the Stripe account information for the authenticated caller.
- STRIPE_COUPONS View
  STRIPE_COUPONS lists the coupons issued in the Stripe account.
- STRIPE_CUSTOMERS View
  STRIPE_CUSTOMERS lists the customers defined in the Stripe account.

- STRIPE_INVOICES View
  STRIPE_INVOICES lists the invoices defined in the Stripe account.
- STRIPE_PLANS View
  STRIPE_PLANS lists the pricing plans defined in the Stripe account.
- STRIPE_PRODUCTS View
  STRIPE_PRODUCTS lists the query information about products defined in the Stripe account.
- STRIPE_SUBSCRIPTIONS View
  STRIPE_SUBSCRIPTIONS lists the subscriptions managed in the Stripe account.

Prerequisites to Use Stripe Views

To use Stripe with Autonomous Database, set Network ACL for accessing Stripe and create a credential that allows access to Stripe.

To use Stripe API:

1. Use the `DBMS_NETWORK_ACL_ADMIN` package to set Network ACL for accessing Stripe.

```
BEGIN
    DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
         host => 'stripe.com',
         ace  => xs$ace_type(privilege_list => xs$name_list('http'),
                             principal_name => 'FOO',
                             principal_type => xs_acl.ptype_db)
   );
END;
/
```

2. Only an ADMIN has access to the Stripe Views. A user must be granted READ access to use them.

   For Example:

```
GRANT READ on STRIPE_COUPONS to TYLER;
```

3. Create a Credential for Stripe APIs. There are two ways to perform this task:
   - (Option 1: Create a credential object with the name `STRIPE$CRED`.) The username is `STRIPE_TOKEN` and the password is the Stripe API token. This token can be found in the Stripe Developers Dashboard. See Stripe API keys for additional information.

```
BEGIN
    DBMS_CLOUD.CREATE_CREDENTIAL(
        credential_name => 'STRIPE$CRED',
        username        => 'STRIPE_TOKEN',
        password        => 'bearer_token' );
END;
/
```

   - (Option 2: Create a credential object with a user defined name.) The username is `STRIPE_TOKEN` and the password is the Stripe API token. This token can be found in the Stripe Developers Dashboard. Before selecting from the Stripe Views, the Session

parameter `DEFAULT_CREDENTIAL` should be set using the owner and credential object name. See Stripe API keys for additional information.

```
BEGIN
    DBMS_CLOUD.CREATE_CREDENTIAL(
        credential_name => 'MY_STRIPE_CRED',
        username        => 'STRIPE_TOKEN',
        password        => 'bearer_token');
END;
/
ALTER SESSION SET default_credential = 'MY_SCHEMA.MY_STRIPE_CRED';
```

4. Test by querying a Stripe View:

```
SELECT name, percent_off, duration, times_redeemed FROM STRIPE_COUPONS;

NAME              PERCENT_OFF DURATION        TIMES_REDEEMED
--------------- ------------- --------------- --------------
Promotion Feb           10.00 once                         0
Seasonal disc            5.00 once                         0
Firstpurchase            5.00 forever                      1
```

STRIPE_ACCOUNTS View

STRIPE_ACCOUNTS lists the Stripe account information for the authenticated caller.

See Account for additional information.

Column	Datatype	Description
ID	VARCHAR2(4000)	Unique identifier for the object.
OBJECT	VARCHAR2(4000)	String representing the objects type.
BUSINESS_TYPE	VARCHAR2(4000)	The business type.
BUSINESS_PROFILE	CLOB	Business information about the account.
CAPABILITIES	CLOB	A hash containing the set of capabilities that was requested for this account and their associated states.
CHARGES_ENABLED	VARCHAR2(4000)	Whether the account can create live charges.
COMPANY	CLOB	Information about the company or business.
CONTROLLER	CLOB	The controller of the account.
COUNTRY	VARCHAR2(4000)	The account country.
CREATED	VARCHAR2(4000)	Time at which the account was connected.
DEFAULT_CURRENCY	VARCHAR2(4000)	Three-letter ISO currency code representing the default currency for the account.
DETAILS_SUBMITTED	VARCHAR2(4000)	Whether account details have been submitted.
EMAIL	VARCHAR2(4000)	An email address associated with the account.
EXTERNAL_ACCOUNTS	CLOB	External accounts currently attached to this account.
FUTURE_REQUIREMENTS	VARCHAR2(4000)	Information about the upcoming new requirements.
INDIVIDUAL	CLOB	Information about the person represented by the account.

Column	Datatype	Description
METADATA	VARCHAR2(4000)	Set of key-value pairs that you can attach to an object.
PAYOUTS_ENABLED	VARCHAR2(4000)	Whether Stripe can send payouts to this account.
REQUIREMENTS	CLOB	Information about the requirements for the account.
SETTINGS	CLOB	Options for customizing how the account functions within Stripe.
TOS_ACCEPTANCE	CLOB	Details on the acceptance of the Stripe Services Agreement.
TYPE	VARCHAR2(4000)	The Stripe account type.

STRIPE_COUPONS View

STRIPE_COUPONS lists the coupons issued in the Stripe account.

See Discounts for subscriptions for additional information.

Column	Datatype	Description
ID	VARCHAR2(4000)	Unique identifier for the object.
OBJECT	VARCHAR2(4000)	String representing the object type.
AMOUNT_OFF	NUMBER	Amount that will be taken off the subtotal of any invoices for this customer.
CREATED	VARCHAR2(4000)	Time at which the object was created.
CURRENCY	VARCHAR2(4000)	If amount_off has been set, the three-letter ISO code for the currency of the amount to take off.
DURATION	VARCHAR2(4000)	Describes how long a customer who applies this coupon will get the discount.
DURATION_IN_MONTHS	NUMBER	If duration is repeating, the number of months the coupon applies.
LIVEMODE	VARCHAR2(4000)	Has the value true if the object exists in live mode or the value false if the object exists in test mode.
MAX_REDEMPTIONS	NUMBER	Maximum number of times this coupon can be redeemed, in total, across all customers, before it is no longer valid.
METADATA	CLOB	Set of key-value pairs that you can attach to an object.
NAME	VARCHAR2(4000)	Name of the coupon displayed to customers on for instance invoices or receipts.

Column	Datatype	Description
PERCENT_OFF	NUMBER	Percent that will be taken off the subtotal of any invoices for this customer for the duration of the coupon.
REDEEM_BY	VARCHAR2(4000)	Date after which the coupon can no longer be redeemed.
TIMES_REDEEMED	NUMBER	Number of times this coupon has been applied to a customer.
VALID	VARCHAR2(4000)	Taking account of the above properties, whether this coupon can still be applied to a customer.

STRIPE_CUSTOMERS View

STRIPE_CUSTOMERS lists the customers defined in the Stripe account.

See Customers for additional information.

Column	Datatype	Description
ID	VARCHAR2(4000)	Unique identifier for the customer.
OBJECT	VARCHAR2(4000)	String representing the object type.
ADDRESS	CLOB	The customers address.
BALANCE	NUMBER	Current balance, if any, being stored on the customer.
CREATED	VARCHAR2(4000)	Time at which the object was created.
CURRENCY	VARCHAR2(4000)	Three-letter ISO code for the currency the customer can be charged in for recurring billing purposes.
DEFAULT_SOURCE	VARCHAR2(4000)	ID of the default payment source for the customer.
DELETED	VARCHAR2(4000)	True if the customer is marked as deleted.
DELINQUENT	VARCHAR2(4000)	When the customers latest invoice is billed by charging automatically, delinquent is true if the invoices latest charge failed.
DESCRIPTION	VARCHAR2(4000)	An arbitrary string attached to the object.
DISCOUNT	CLOB	Describes the current discount active on the customer, if there is one.
EMAIL	VARCHAR2(4000)	The customer email address.
INVOICE_PREFIX	VARCHAR2(4000)	The prefix for the customer used to generate unique invoice numbers.
INVOICE_SETTINGS	CLOB	The customer default invoice settings.

Column	Datatype	Description
LIVEMODE	VARCHAR2(4000)	Has the value true if the object exists in live mode or then value false if the object exists in test mode.
METADATA	CLOB	Set of key-value pairs that you can attach to an object.
NAME	VARCHAR2(4000)	The customer full name or business name.
NEXT_INVOICE_SEQUENCE	NUMBER	The suffix of the customers next invoice number.
PHONE	VARCHAR2(4000)	The customers phone number.
PREFERRED_LOCALES	CLOB	The customers preferred locales, ordered by preference.
SHIPPING	CLOB	Mailing and shipping address for the customer.
TAX_EXEMPT	VARCHAR2(4000)	Describes the customers tax exemption status.
TAX_IDS	CLOB	The customer tax IDs.
TEST_CLOCK	CLOB	ID of the test clock this customer belongs to.

STRIPE_INVOICES View

STRIPE_INVOICES lists the invoices defined in the Stripe account.

See Invoicing for additional information.

Column	Datatype	Description
ID	VARCHAR2(4000)	Unique identifier for the object
OBJECT	VARCHAR2(4000)	String representing the objects type.
ACCOUNT_COUNTRY	VARCHAR2(4000)	The country of the business associated with this invoice.
ACCOUNT_NAME	VARCHAR2(4000)	The public name of the business associated with this invoice.
ACCOUNT_TAX_IDS	CLOB	The account tax IDs associated with the invoice.
AMOUNT_DUE	NUMBER	Final amount due at this time for this invoice.
AMOUNT_PAID	NUMBER	The amount, in cents, that was paid.
AMOUNT_REMAINING	NUMBER	The difference between amount_due and amount_paid, in cents.
AMOUNT_SHIPPING	NUMBER	This is the sum of all the shipping amounts.
APPLICATION	VARCHAR2(4000)	ID of the Connect Application that created the invoice.

Column	Datatype	Description
APPLICATION_FEE_AMOUNT	NUMBER	The fee in cents that will be applied to the invoice and transferred to the application owners Stripe account when the invoice is paid.
ATTEMPT_COUNT	NUMBER	Number of payment attempts made for this invoice, from the perspective of the payment retry schedule.
ATTEMPTED	VARCHAR2(4000)	Whether an attempt has been made to pay the invoice.
AUTO_ADVANCE	VARCHAR2(4000)	Controls whether Stripe will perform automatic collection of the invoice.
AUTOMATIC_TAX	CLOB	Settings and latest results for automatic tax lookup for this invoice.
BILLING_REASON	VARCHAR2(4000)	Indicates the reason why the invoice was created.
CHARGE	CLOB	ID of the latest charge generated for this invoice, if any.
COLLECTION_METHOD	VARCHAR2(4000)	Either charge_automatically, or send_invoice.
CREATED	VARCHAR2(4000)	Time at which the object was created.
CURRENCY	VARCHAR2(4000)	Three-letter ISO currency code, in lowercase.
CUSTOM_FIELDS	CLOB	Custom fields displayed on the invoice.
CUSTOMER	VARCHAR2(4000)	The ID of the customer who will be billed.
CUSTOMER_ADDRESS	CLOB	The customers address.
CUSTOMER_EMAIL	VARCHAR2(4000)	The customers email.
CUSTOMER_NAME	VARCHAR2(4000)	The customers name.
CUSTOMER_PHONE	VARCHAR2(4000)	The customers phone number.
CUSTOMER_SHIPPING	CLOB	The customers shipping information.
CUSTOMER_TAX_EXEMPT	VARCHAR2(4000)	The customers tax exempt status.
CUSTOMER_TAX_IDS	CLOB	The customers tax IDs.
DEFAULT_PAYMENT_METHOD	VARCHAR2(4000)	ID of the default payment method for the invoice.
DEFAULT_SOURCE	VARCHAR2(4000)	ID of the default payment source for the invoice.
DEFAULT_TAX_RATES	CLOB	The tax rates applied to this invoice, if any.
DESCRIPTION	VARCHAR2(4000)	An arbitrary string attached to the object.
DISCOUNT	CLOB	Describes the current discount applied to this invoice, if there is one.

Column	Datatype	Description
DISCOUNTS	CLOB	The discounts applied to the invoice.
DUE_DATE	VARCHAR2(4000)	The date on which payment for this invoice is due.
ENDING_BALANCE	NUMBER	Ending customer balance after the invoice is finalized.
FOOTER	VARCHAR2(4000)	Footer displayed on the invoice.
FROM_INVOICE	VARCHAR2(4000)	Details of the invoice that was cloned.
HOSTED_INVOICE_URL	VARCHAR2(4000)	The URL for the hosted invoice page.
INVOICE_PDF	VARCHAR2(4000)	The link to download the PDF for the invoice.
LAST_FINALIZATION_ERROR	CLOB	The error encountered during the previous attempt to finalize the invoice.
LATEST_REVISION	VARCHAR2(4000)	The ID of the most recent non-draft revision of this invoice.
LINES	CLOB	The individual line items that make up the invoice.
LIVEMODE	VARCHAR2(4000)	Has the value true if the object exists in live mode or the value false if the object exists in test mode.
METADATA	CLOB	Set of key-value pairs that you can attach to an object.
NEXT_PAYMENT_ATTEMPT	VARCHAR2(4000)	The time at which payment will next be attempted.
NUMBER_	VARCHAR2(4000)	A unique, identifying string that appears on emails sent to the customer for this invoice.
ON_BEHALF_OF	VARCHAR2(4000)	The account (if any) for which the funds of the invoice payment are intended.
PAID	VARCHAR2(4000)	Whether payment was successfully collected for this invoice.
PAID_OUT_OF_BAND	VARCHAR2(4000)	Returns true if the invoice was manually marked paid, returns false if the invoice hasn't been paid yet or was paid on Stripe.
PAYMENT_INTENT	CLOB	The PaymentIntent associated with this invoice.
PAYMENT_SETTINGS	CLOB	Configuration settings for the PaymentIntent that is generated when the invoice is finalized.
PERIOD_END	VARCHAR2(4000)	End of the usage period during which invoice items were added to this invoice.
PERIOD_START	VARCHAR2(4000)	Start of the usage period during which invoice items were added to this invoice.

Column	Datatype	Description
POST_PAYMENT_CREDIT_NOTES_AMOUNT	NUMBER	Total amount of all post-payment credit notes issued for this invoice.
PRE_PAYMENT_CREDIT_NOTES_AMOUNT	NUMBER	Total amount of all pre-payment credit notes issued for this invoice.
QUOTE	VARCHAR2(4000)	The quote this invoice was generated from.
RECEIPT_NUMBER	VARCHAR2(4000)	This is the transaction number that appears on email receipts sent for this invoice.
RENDERING_OPTIONS	CLOB	Options for invoice PDF rendering.
SHIPPING_COST	CLOB	The details of the cost of shipping, including the ShippingRate applied on the invoice.
SHIPPING_DETAILS	CLOB	Shipping details for the invoice.
STARTING_BALANCE	NUMBER	Starting customer balance before the invoice is finalized.
STATEMENT_DESCRIPTOR	VARCHAR2(4000)	Extra information about an invoice for the customers credit card statement.
STATUS	VARCHAR2(4000)	The status of the invoice, one of draft, open, paid, uncollectible, or void.
STATUS_TRANSITIONS	CLOB	The timestamps at which the invoice status was updated.
SUBSCRIPTION	VARCHAR2(4000)	The subscription that this invoice was prepared for, if any.
SUBTOTAL	NUMBER	Total of all subscriptions, invoice items, and prorations on the invoice before any invoice level discount or exclusive tax is applied.
SUBTOTAL_EXCLUDING_TAX	NUMBER	The integer amount in cents representing the subtotal of the invoice before any invoice level discount or tax is applied.
TAX	NUMBER	The amount of tax on this invoice.
THRESHOLD_REASON	CLOB	If billing_reason is set to subscription_threshold this returns more information on which threshold rules triggered the invoice.
TEST_CLOCK	VARCHAR2(4000)	ID of the test clock this invoice belongs to.
TOTAL	NUMBER	Total after discounts and taxes.
TOTAL_DISCOUNT_AMOUNTS	CLOB	The aggregate amounts calculated per discount across all line items.

Column	Datatype	Description
TOTAL_EXCLUDING_TAX	NUMBER	The integer amount in cents representing the total amount of the invoice including all discounts but excluding all tax.
TOTAL_TAX_AMOUNTS	CLOB	The aggregate amounts calculated per tax rate for all line items.
TRANSFER_DATA	CLOB	The account (if any) the payment will be attributed to for tax reporting, and where funds from the payment will be transferred to for the invoice.
WEBHOOKS_DELIVERED_AT	VARCHAR2(4000)	This field tracks the time when webhooks for this invoice were successfully delivered.

STRIPE_PLANS View

STRIPE_PLANS lists the pricing plans defined in the Stripe account.

See Plans for additional information.

Column	Datatype	Description
ID	VARCHAR2(4000)	Unique identifier for the plan.
OBJECT	VARCHAR2(4000)	String representing the object type.
ACTIVE	VARCHAR2(4000)	Whether the plan is currently available for purchase.
AGGREGATE_USAGE	VARCHAR2(4000)	Specifies a usage aggregation strategy for plans of usage_type=metered.
AMOUNT	NUMBER	The unit amount in cents to be charged, represented as a whole integer.
AMOUNT_DECIMAL	VARCHAR2(4000)	The unit amount in cents to be charged, represented as a decimal string with at most 12 decimal places.
BILLING_SCHEME	VARCHAR2(4000)	Describes how to compute the price per period.
CREATED	VARCHAR2(4000)	Time at which the plan was created.
CURRENCY	VARCHAR2(4000)	Three-letter ISO currency code, in lowercase.
DELETED	VARCHAR2(4000)	True if the plan is marked as deleted.
INTERVAL	VARCHAR2(4000)	The frequency at which a subscription is billed.
INTERVAL_COUNT	NUMBER	The number of intervals between subscription billings.

Column	Datatype	Description
LIVEMODE	VARCHAR2(4000)	Has the value true if the plan exists in live mode or the value false if the plan exists in test mode.
METADATA	CLOB	Set of key-value pairs that you can attach to an plan.
NICKNAME	VARCHAR2(4000)	A brief description of the plan, hidden from customers.
PRODUCT	VARCHAR2(4000)	ID of the product whose pricing this plan determines.
TIERS	CLOB	Each element represents a pricing tier.
TIERS_MODE	VARCHAR2(4000)	Defines if the tiering price should be graduated or volume based.
TRANSFORM_USAGE	CLOB	Apply a transformation to the reported usage or set quantity before computing the amount billed.
TRIAL_PERIOD_DAYS	NUMBER	Default number of trial days when subscribing a customer to this plan using trial_from_plan=true.
USAGE_TYPE	VARCHAR2(4000)	Configures how the quantity per period should be determined.

STRIPE_PRODUCTS View

STRIPE_PRODUCTS lists the query information about products defined in the Stripe account.

See Manage products and prices for additional information.

Column	Datatype	Description
ID	VARCHAR2(4000)	Unique identifier for the object.
OBJECT	VARCHAR2(4000)	String representing the object type.
ACTIVE	VARCHAR2(4000)	Whether the product is currently available for purchase.
CREATED	VARCHAR2(4000)	Time at which the object was created.
DEFAULT_PRICE	VARCHAR2(4000)	The ID of the Price object that is the default price for this product.
DESCRIPTION	VARCHAR2(4000)	The product description, meant to be displayable to the customer.
IMAGES	CLOB	A list of up to 8 URLs of images for this product, meant to be displayable to the customer.
LIVEMODE	VARCHAR2(4000)	Has the value true if the object exists in live mode or the value false if the object exists in test mode.
METADATA	CLOB	Set of key-value pairs that you can attach to an object.

Column	Datatype	Description
NAME	VARCHAR2(4000)	The product name, meant to be displayable to the customer.
PACKAGE_DIMENSIONS	CLOB	The dimensions of this product for shipping purposes.
SHIPPABLE	VARCHAR2(4000)	Whether this product is shipped.
STATEMENT_DESCRIPTOR	VARCHAR2(4000)	Extra information about a product which will appear on your customers credit card statement.
TAX_CODE	VARCHAR2(4000)	A tax code ID.
UNIT_LABEL	VARCHAR2(4000)	A label that represents units of this product.
UPDATED	VARCHAR2(4000)	Time at which the object was last updated.
URL	VARCHAR2(4000)	A URL of a publicly-accessible webpage for this product.

STRIPE_SUBSCRIPTIONS View

STRIPE_SUBSCRIPTIONS lists the subscriptions managed in the Stripe account.

See Subscriptions for additional information.

Column	Datatype	Subscription
ID	VARCHAR2(4000)	Unique identifier for the object.
OBJECT	VARCHAR2(4000)	String representing the object type.
APPLICATION	VARCHAR2(4000)	ID of the Connect Application that created the subscription.
APPLICATION_FEE_PERCENT	NUMBER	A non-negative decimal between 0 and 100, with at most two decimal places.
AUTOMATIC_TAX	CLOB	Automatic tax settings for this subscription.
BILLING_CYCLE_ANCHOR	VARCHAR2(4000)	Determines the date of the first full invoice, and, for plans with month or year intervals, the day of the month for subsequent invoices.
BILLING_THRESHOLDS	CLOB	Define thresholds at which an invoice will be sent, and the subscription advanced to a new billing period.
CANCEL_AT	VARCHAR2(4000)	A date in the future at which the subscription will automatically get canceled.
CANCEL_AT_PERIOD_END	VARCHAR2(4000)	If the subscription has been canceled with the at_period_end flag set to true, cancel_at_period_end on the subscription will be true.

Column	Datatype	Subscription
CANCELED_AT	VARCHAR2(4000)	If the subscription has been canceled, the date of that cancellation.
COLLECTION_METHOD	VARCHAR2(4000)	Either charge_automatically, or send_invoice.
CREATED	VARCHAR2(4000)	Time at which the object was created.
CURRENCY	VARCHAR2(4000)	Three-letter ISO currency code, in lowercase.
CURRENT_PERIOD_END	VARCHAR2(4000)	End of the current period that the subscription has been invoiced for.
CURRENT_PERIOD_START	VARCHAR2(4000)	Start of the current period that the subscription has been invoiced for.
CUSTOMER	VARCHAR2(4000)	ID of the customer who owns the subscription.
DAYS_UNTIL_DUE	NUMBER	Number of days a customer has to pay invoices generated by this subscription.
DEFAULT_PAYMENT_METHOD	VARCHAR2(4000)	ID of the default payment method for the subscription.
DEFAULT_SOURCE	VARCHAR2(4000)	ID of the default payment source for the subscription.
DEFAULT_TAX_RATES	CLOB	The tax rates that will apply to any subscription item that does not have tax_rates set.
DESCRIPTION	VARCHAR2(4000)	The subscription description, meant to be displayable to the customer.
DISCOUNT	CLOB	Describes the current discount applied to this subscription, if there is one.
ENDED_AT	VARCHAR2(4000)	If the subscription has ended, the date the subscription ended.
ITEMS	CLOB	List of subscription items, each with an attached price.
LATEST_INVOICE	VARCHAR2(4000)	The most recent invoice this subscription has generated.
LIVEMODE	VARCHAR2(4000)	Has the value true if the object exists in live mode or the value false if the object exists in test mode.
METADATA	CLOB	Set of key-value pairs that you can attach to an object.
NEXT_PENDING_INVOICE_ITEM_INVOICE	VARCHAR2(4000)	Specifies the approximate timestamp on which any pending invoice items will be billed acc to the schedule provided at pending_invoice_item_interval.
ON_BEHALF_OF	VARCHAR2(4000)	The account the charge was made on behalf of for charges associated with this subscription.

Column	Datatype	Subscription
PAUSE_COLLECTION	CLOB	If specified, payment collection for this subscription will be paused.
PAYMENT_SETTINGS	VARCHAR2(4000)	Payment settings passed on to invoices created by the subscription.
PENDING_INVOICE_ITEM_INTERVAL	CLOB	Specifies an interval for how often to bill for any pending invoice items.
PENDING_SETUP_INTENT	VARCHAR2(4000)	For collecting user authentication when creating a subscription without immediate payment or updating a subscriptions payment method, allowing you to optimize for off-session payments.
PENDING_UPDATE	CLOB	If specified, pending updates that will be applied to the subscription once the latest_invoice has been paid.
SCHEDULE	CLOB	The schedule attached to the subscription.
START_DATE	VARCHAR2(4000)	Date when the subscription was first created.
STATUS	VARCHAR2(4000)	Possible values are incomplete, incomplete_expired, trialing, active, past_due, canceled, or unpaid.
TEST_CLOCK	VARCHAR2(4000)	ID of the test clock this subscription belongs to.
TRANSFER_DATA	CLOB	The account the subscriptions payments will be attributed to for tax reporting, and where funds from each payment will be transferred to for each of the subscriptions invoices.
TRIAL_END	VARCHAR2(4000)	If the subscription has a trial, the end of that trial.
TRIAL_SETTINGS	VARCHAR2(4000)	Settings related to subscription trials.
TRIAL_START	VARCHAR2(4000)	If the subscription has a trial, the beginning of that trial.

Track DBMS_CLOUD Load Operations

DBMS_CLOUD tracks details for load operations for each instance and for users.

- DBA_LOAD_OPERATIONS and USER_LOAD_OPERATIONS Views
  USER_LOAD_OPERATIONS and DBA_LOAD_OPERATIONS views show the load operations in your schema or for all users. The USER_LOAD_OPERATIONS view does not display the USERNAME column.

DBA_LOAD_OPERATIONS and USER_LOAD_OPERATIONS Views

`USER_LOAD_OPERATIONS` and `DBA_LOAD_OPERATIONS` views show the load operations in your schema or for all users. The `USER_LOAD_OPERATIONS` view does not display the `USERNAME` column.

Column	Datatype	NULL	Description
ID	NUMBER	NOT NULL	Specifies the operation ID.
TYPE	VARCHAR2(128)	NOT NULL	Shows the operation type.
SID	NUMBER	NOT NULL	Session ID of the session.
SERIAL#	NUMBER	NOT NULL	
USERNAME	VARCHAR2(128)	NOT NULL	User ID of the user performing the operation. The `USER_LOAD_OPERATIONS` view does not display the `USERNAME` column.
START_TIME	TIMESTAMP(6) WITH TIME ZONE		Start time of the task.
UPDATE_TIME	TIMESTAMP(6) WITH TIME ZONE		Last update time for the task.
STATUS	VARCHAR2(9)		Status of the operation. Valid values are: `COMPLETED`, `FAILED`, `PENDING`, `RUNNING`, `SKIPPED`.
OWNER_NAME	VARCHAR2(128)		Schema that owns the operation.
TABLE_NAME	VARCHAR2(128)		
PARTITION_NAME	VARCHAR2(128)		
SUBPARTITION_NAME	VARCHAR2(128)		
FILE_URI_LIST	VARCHAR2(4000)		This value shows either a comma-delimited list of source file URIs or one or more directories and source files that were specified in the `FILE_URI_LIST` parameter.
ROWS_LOADED	NUMBER		Number of rows loaded in the load operation.
LOGFILE_TABLE	VARCHAR2(128)		Shows the name of the table you can query to look at the log of a load operation.
BADFILE_TABLE	VARCHAR2(128)		Shows the name of the table you can query to look at the rows that got errors during loading.
STATUS_TABLE	VARCHAR2(128)		Shows the name of the status table for the load operation.
TEMPEXT_TABLE	VARCHAR2(128)		
CREDENTIAL_NAME	VARCHAR2(128)		
EXPIRATION_TIME	TIMESTAMP(9) WITH TIME ZONE		

Autonomous Database – Oracle Database Features

Describes Oracle Database features available with Autonomous Database.

Autonomous Database includes features that:

- Automate index management tasks, such as creating, rebuilding, and dropping indexes based on changes in the application workload.

  See Manage Automatic Indexing on Autonomous Database for more information.

  > **Note:**
  >
  > There are restrictions for Automatic Indexing when you use JSON data with Autonomous Database. See SODA Notes for more information.

- Gather real-time statistics automatically while a conventional DML workload is running. Because statistics can go stale between stats gathering jobs, online statistics gathering for conventional DML helps the optimizer generate more optimal plans. Online statistics aim to reduce the possibility of the optimizer being misled by stale statistics.

  See Real-Time Statistics for more information.

- Gather statistics automatically on a more frequent basis. High-Frequency Automatic Optimizer Statistics Collection complements the standard statistics collection job. By default, the collection occurs every 15 minutes, meaning that statistics have less time in which to be stale. High-Frequency Automatic Optimizer Statistics Collection is enabled by default.

  See Configuring High-Frequency Automatic Optimizer Statistics Collection for more information.

- Quarantine execution plans for SQL statements, for example, statements that are terminated by the Resource Manager for consuming excessive system resources in an Oracle Database. Automatic SQL Quarantine based on Resource Manager consumption limit violations is disabled by default but any manually quarantined SQL statement will be honored.

  See Quarantine for Execution Plans for SQL Statements Consuming Excessive System Resources for more information.

- Automatically assess the opportunity for SQL plan changes to improve the performance for known statements.

  See Managing the SPM Evolve Advisor Task for more information.

- Apache ORC format is supported in Autonomous Database for loading and querying data in object store.

  See Create Credentials and Load Data Pump Dump Files into an Existing Table and Query External Data with ORC, Parquet, or Avro Source Files for more information.

- Complex types are supported in Autonomous Database for ORC, Avro, and Parquet structured files.

  See DBMS_CLOUD Package Avro, ORC, and Parquet Complex Types for more information.

Autonomous Database with Oracle Database 23ai

When you provision or clone an Autonomous Database instance you can select either of these versions: Oracle Database 19c or Oracle Database 23ai.

See Provision an Autonomous Database Instance and Clone an Autonomous Database Instance for information on provisioning and cloning Autonomous Database.

Topics

- Autonomous Database with Oracle Database 23ai Features
  Autonomous Database with Oracle Database 23ai offers many new innovative autonomous features and provides developer-oriented functionality.
- Autonomous Database with Oracle Database 23ai Notes
  If you are using Autonomous Database with Oracle Database 23ai, the following Oracle Database 23ai functionality is not supported:

Autonomous Database with Oracle Database 23ai Features

Autonomous Database with Oracle Database 23ai offers many new innovative autonomous features and provides developer-oriented functionality.

If you use Autonomous Database with Oracle Database 23ai, then many of the concepts and features of this service are further documented here:

- Oracle Database 23ai
- *Oracle Database New Features*

Autonomous Database with Oracle Database 23ai Notes

If you are using Autonomous Database with Oracle Database 23ai, the following Oracle Database 23ai functionality is not supported:

- Autonomous Database with Oracle Database 23ai in the Paid tier is available in all commercial public cloud regions.

  See Provision an Autonomous Database Instance for more information.
- Always Free Autonomous Database with Oracle Database 23ai is available in all commercial public cloud regions *except* the following regions: Colombia Central: Bogota (BOG), Saudi Arabia Central (RUH), Singapore West: Singapore (XSP)
  See Provision an Autonomous Database Instance for more information.

Autonomous Database RMAN Recovery Catalog

You can use Oracle Autonomous Database as a Recovery Manager (RMAN) recovery catalog. A recovery catalog is a database schema that RMAN uses to store metadata about one or more Oracle databases.

- Use Autonomous Database as an RMAN Recovery Catalog
  Recovery Manager (RMAN) recovery catalog is preinstalled in Autonomous Database in schema `RMAN$CATALOG`. The preinstalled catalog version is based on the latest version of Oracle Database and is compatible with all supported Oracle database versions.

Use Autonomous Database as an RMAN Recovery Catalog

Recovery Manager (RMAN) recovery catalog is preinstalled in Autonomous Database in schema `RMAN$CATALOG`. The preinstalled catalog version is based on the latest version of Oracle Database and is compatible with all supported Oracle database versions.

The recovery catalog contains metadata about RMAN operations for each registered target database. When RMAN is connected to a recovery catalog, RMAN obtains its metadata exclusively from the catalog.

> **Note:**
>
> Autonomous Database is not supported as an RMAN target database. An RMAN target database is an Oracle Database to which RMAN is connected with the `TARGET` keyword. A target database is a database on which RMAN is performing backup and recovery operations. See Backup and Restore Autonomous Database Instances for information on Autonomous Database backup and recovery operations.

Access to RMAN Recovery Catalog

Access to the recovery catalog is provided through predefined user `RMAN$CATALOG` with the appropriate access to the recovery catalog only. The `RMAN$CATALOG` user is locked by default.

You can either proxy to the predefined user `RMAN$CATALOG` through the ADMIN user or explicitly unlock the preinstalled schema:

- ADMIN user proxy into `RMAN$CATALOG` using ADMIN user's password:

  ```
  connect admin[rman$catalog]/password@connect_string
  ```

- ADMIN user can set a password for `RMAN$CATALOG`. Then the `RMAN$CATALOG` user can directly connect:

  ```
  connect admin/password@connect_string
  alter user rman$catalog identified by password account unlock;
  connect rman$catalog/password@connect_string
  ```

Use the RMAN Recovery Catalog

You can use the RMAN recovery catalog by connecting RMAN to the preinstalled recovery catalog. Registering a target database in the recovery catalog maintains the database's records in the recovery catalog. For example, to register a target database:

```
RMAN> connect catalog rman$catalog/password@connect_string;

connected to recovery catalog database
recovery catalog schema version 21.01.00.00. is newer than RMAN version

RMAN> register database;
database registered in recovery catalog
starting full resync of recovery catalog
```

To use your Autonomous Database as a recovery catalog, it is recommended to connect with the `LOW` service.

See Registering a Database in the Recovery Catalog for more details about using the RMAN recovery catalog.

Notes for Users Migrating from Other Oracle Databases

Describes information that is useful when you are migrating from other Oracle Databases to Oracle Autonomous Database.

- Initialization Parameters
  Autonomous Database configures database initialization parameters automatically when you provision a database. You do not need to set any initialization parameters to start using your service. But, you can modify some parameters if you need to.
- SQL Commands
  Autonomous Database allows most of the SQL commands available in Oracle Database. To ensure the security and the performance of Autonomous Database, some SQL commands are restricted.
- Data Types
  Autonomous Database allows most of the data types available in Oracle Database. To ensure the security and the performance of Autonomous Database, some data types are restricted.
- PL/SQL Package Notes for Autonomous Database
  Notes for Oracle Database PL/SQL packages in Autonomous Database.
- ADMIN User Roles and Privileges
  In Autonomous Database the predefined administrative user is ADMIN and this account has privileges to manage users and to manage the database.
- Oracle XML DB
  Describes Autonomous Database support for Oracle XML DB features. To ensure the security and the performance of your Autonomous Database, some Oracle XML DB features are restricted.
- Oracle Text
  Describes Autonomous Database support for Oracle Text features. To ensure the security and the performance of your Autonomous Database, some Oracle Text features are restricted.
- Oracle Flashback
  Oracle Flashback Technology is a group of Oracle Database features that let you view past states of database objects or to return database objects to a previous state without using point-in-time media recovery.
- Oracle Database Real Application Security
  Oracle Database Real Application Security is a database authorization model that: supports declarative security policies, enables end-to-end security for multitier applications, provides an integrated solution to secure database and application resources, and advances the security architecture of Oracle Database to meet existing and emerging demands of applications developed for the Internet.
- Oracle LogMiner
  Describes restrictions for Oracle LogMiner on Autonomous Database.
- Choose a Character Set for Autonomous Database
- Query the Alert Log
  You can query the database alert log using `V$DIAG_ALERT_EXT` to search for log data.

Initialization Parameters

Autonomous Database configures database initialization parameters automatically when you provision a database. You do not need to set any initialization parameters to start using your service. But, you can modify some parameters if you need to.

Modifiable Initialization Parameters

The following table shows the initialization parameters that are only modifiable with `ALTER SESSION`.

Only Modifiable with ALTER SESSION	More Information
CONSTRAINTS	
CONTAINER	
CURRENT_SCHEMA	
CURSOR_INVALIDATION	CURSOR_INVALIDATION
DEFAULT_COLLATION	
DEFAULT_CREDENTIAL	
EDITION	
ISOLATION_LEVEL	
JSON_BEHAVIOR	This parameter is only applicable with Oracle Database 23ai. See JSON_BEHAVIOR for more information.
JSON_EXPRESSION_CHECK	JSON_EXPRESSION_CHECK
OPTIMIZER_SESSION_TYPE	OPTIMIZER_SESSION_TYPE
OPTIMIZER_USE_INVISIBLE_INDEXES	OPTIMIZER_USE_INVISIBLE_INDEXES
READ_ONLY	
SQL_TRACE	See Perform SQL Tracing on Autonomous Database for details
SQL_TRANSLATION_PROFILE	
SQL_TRANSPILER	This parameter is only applicable with Oracle Database 23ai. See SQL_TRANSPILER for more information.
STATISTICS_LEVEL	STATISTICS_LEVEL
TIME_ZONE	For more information on `TIME_ZONE`, see *Oracle Database SQL Language Reference*.
XML_PARAMS	This parameter is only applicable with Oracle Database 23ai. See XML_PARAMS for more information.

The following table shows the initialization parameters that are only modifiable with `ALTER SYSTEM`.

Only Modifiable with ALTER SYSTEM	More Information
BLANK_TRIMMING	BLANK_TRIMMING
FIXED_DATE	FIXED_DATE
JOB_QUEUE_PROCESSES	JOB_QUEUE_PROCESSES
LOCKDOWN_ERRORS	See LOCKDOWN_ERRORS for details

Only Modifiable with ALTER SYSTEM	More Information
MAX_IDLE_BLOCKER_TIME	MAX_IDLE_BLOCKER_TIME With a value higher than 5, the parameter acts as if it was set to 5
MAX_IDLE_TIME	MAX_IDLE_TIME By default MAX_IDLE_TIME is set 1440. Setting the MAX_IDLE_TIME value to 0 or to a value higher than 1440, the parameter acts as if it was set to 1440.
SESSION_EXIT_ON_PACKAGE_STATE_ERROR	SESSION_EXIT_ON_PACKAGE_STATE_ERROR

The following table shows the initialization parameters that are modifiable with either ALTER SESSION or ALTER SYSTEM.

Modifiable with ALTER SESSION or ALTER SYSTEM	More Information
APPROX_FOR_AGGREGATION	APPROX_FOR_AGGREGATION
APPROX_FOR_COUNT_DISTINCT	APPROX_FOR_COUNT_DISTINCT
APPROX_FOR_PERCENTILE	APPROX_FOR_PERCENTILE
CLIENT_PREFETCH_ROWS	See CLIENT_PREFETCH_ROWS
CONTAINER_DATA	CONTAINER_DATA
CURSOR_SHARING	CURSOR_SHARING
DDL_LOCK_TIMEOUT	DDL_LOCK_TIMEOUT
GROUP_BY_POSITION	
GROUP_BY_POSITION_ENABLED	This parameter is only applicable with Oracle Database 23ai. See GROUP_BY_POSITION_ENABLED for more information
HEAT_MAP	HEAT_MAP
IGNORE_SESSION_SET_PARAM_ERRORS	IGNORE_SESSION_SET_PARAM_ERRORS
LDAP_DIRECTORY_ACCESS	LDAP_DIRECTORY_ACCESS
LOAD_WITHOUT_COMPILE	
MAX_STRING_SIZE	See Data Types for details
NLS_CALENDAR	NLS_CALENDAR
NLS_COMP	NLS_COMP
NLS_CURRENCY	NLS_CURRENCY
NLS_DATE_FORMAT	NLS_DATE_FORMAT
NLS_DATE_LANGUAGE	NLS_DATE_LANGUAGE
NLS_DUAL_CURRENCY	NLS_DUAL_CURRENCY
NLS_ISO_CURRENCY	NLS_ISO_CURRENCY
NLS_LANGUAGE	NLS_LANGUAGE
NLS_LENGTH_SEMANTICS	NLS_LENGTH_SEMANTICS
NLS_NCHAR_CONV_EXCP	NLS_NCHAR_CONV_EXCP
NLS_NUMERIC_CHARACTERS	NLS_NUMERIC_CHARACTERS
NLS_SORT	NLS_SORT
NLS_TERRITORY	NLS_TERRITORY
NLS_TIME_FORMAT	

Modifiable with ALTER SESSION or ALTER SYSTEM	More Information
NLS_TIME_TZ_FORMAT	
NLS_TIMESTAMP_FORMAT	NLS_TIMESTAMP_FORMAT
NLS_TIMESTAMP_TZ_FORMAT	NLS_TIMESTAMP_TZ_FORMAT
OPTIMIZER_CAPTURE_SQL_QUARANTINE	OPTIMIZER_CAPTURE_SQL_QUARANTINE
OPTIMIZER_IGNORE_HINTS	For more information on OPTIMIZER_IGNORE_HINTS, see Manage Optimizer Statistics on Autonomous Database.
OPTIMIZER_IGNORE_PARALLEL_HINTS	For more information on OPTIMIZER_IGNORE_PARALLEL_HINTS, see Manage Optimizer Statistics on Autonomous Database.
OPTIMIZER_MODE	OPTIMIZER_MODE
OPTIMIZER_REAL_TIME_STATISTICS	OPTIMIZER_REAL_TIME_STATISTICS
OPTIMIZER_USE_SQL_QUARANTINE	OPTIMIZER_USE_SQL_QUARANTINE
PLSCOPE_SETTINGS	PLSCOPE_SETTINGS
PLSQL_CCFLAGS	PLSQL_CCFLAGS
PLSQL_DEBUG	PLSQL_DEBUG
PLSQL_IMPLICIT_CONVERSION_BOOL	This parameter is only applicable with Oracle Database 23ai. See PLSQL_IMPLICIT_CONVERSION_BOOL
PLSQL_OPTIMIZE_LEVEL	PLSQL_OPTIMIZE_LEVEL
PLSQL_WARNINGS	PLSQL_WARNINGS
QUERY_REWRITE_INTEGRITY	QUERY_REWRITE_INTEGRITY
RECYCLEBIN	RECYCLEBIN
REMOTE_DEPENDENCIES_MODE	REMOTE_DEPENDENCIES_MODE
RESULT_CACHE_INTEGRITY	See RESULT_CACHE_INTEGRITY
RESULT_CACHE_MODE	See RESULT_CACHE_MODE
SKIP_UNUSABLE_INEDEXES	SKIP_UNUSABLE_INDEXES
SQL_HISTORY_ENABLED	This parameter is available starting with Oracle Database 23ai. See SQL_HISTORY_ENABLED for more information.
SYSDATE_AT_DBTIMEZONE	See SYSDATE_AT_DBTIMEZONE Select a Time Zone for SYSDATE on Autonomous Database
XML_CLIENT_SIDE_DECODING	See XML_CLIENT_SIDE_DECODING

For more information on initialization parameters see *Oracle Database Reference*.

- SESSION_EXIT_ON_PACKAGE_STATE_ERROR
  SESSION_EXIT_ON_PACKAGE_STATE_ERROR enables or disables special handling for stateful PL/SQL packages running in a session.
- SYSDATE_AT_DBTIMEZONE Select a Time Zone for SYSDATE on Autonomous Database
- CLIENT_PREFETCH_ROWS
- JOB_QUEUE_PROCESSES
- LOCKDOWN_ERRORS
  Set the LOCKDOWN_ERRORS parameter to control whether Autonomous Database lockdown profile errors are raised or ignored.

- RESULT_CACHE_INTEGRITY
  Set the RESULT_CACHE_INTEGRITY parameter to specify whether the result cache considers queries using possibly non-deterministic constructs as candidates for result caching.
- RESULT_CACHE_MODE
  Set the RESULT_CACHE_MODE parameter to specify which queries are eligible to store result sets in the result cache. Only query execution plans with the result cache operator will attempt to read from or write to the result cache.

SESSION_EXIT_ON_PACKAGE_STATE_ERROR

SESSION_EXIT_ON_PACKAGE_STATE_ERROR enables or disables special handling for stateful PL/SQL packages running in a session.

Property	Description
Parameter type	Boolean
Default Value	FALSE
Modifiable	ALTER SYSTEM
Range of values	TRUE | FALSE

SESSION_EXIT_ON_PACKAGE_STATE_ERROR specifies the handling for a stateful PL/SQL package running in a session. When such a package undergoes modification, such as during planned maintenance for Oracle-supplied objects, the sessions that have an active instantiation of the package receive the following error when they attempt to run the package:

```
ORA-4068 existing state of package has been discarded
```

However, the application code that receives the ORA-4068 error may not be equipped to handle this error with its retry logic.

Setting SESSION_EXIT_ON_PACKAGE_STATE_ERROR to TRUE provides different handling for this case. When SESSION_EXIT_ON_PACKAGE_STATE_ERROR is TRUE, instead of just raising the ORA-4068 error when the package state is discarded, the session immediately exits. This can be advantageous because many applications are able to handle session termination by automatically and transparently re-establishing the connection.

SYSDATE_AT_DBTIMEZONE Select a Time Zone for SYSDATE on Autonomous Database

SYSDATE_AT_DBTIMEZONE enables special handling in a session for the date and time value returned in calls to SYSDATE and SYSTIMESTAMP.

Depending on the value of SYSDATE_AT_DBTIMEZONE, you see either the date and time based on the default Autonomous Database time zone, Coordinated Universal Time (UTC), or based on the time zone that you set in your database.

Property	Description
Parameter type	Boolean
Default Value	FALSE
Modifiable	ALTER SESSION, ALTER SYSTEM
Range of values	TRUE | FALSE

Default Autonomous Database Time Zone

The default Autonomous Database time zone is Coordinated Universal Time (UTC) and by default calls to `SYSDATE` and `SYSTIMESTAMP` return the date and time in UTC.

In order to change database time zone, you can run the following statement.

This example sets the database time zone to `UTC-5`.

```
ALTER DATABASE SET TIME_ZONE='-05:00';
```

This example sets the database time zone for the US east coast: `America/New_York` and the time is automatically adjusted at DST switches:

```
ALTER DATABASE SET TIME_ZONE='America/New_York';
```

> **Note:**
>
> You must restart the Autonomous Database instance for the change to take effect.

After you set the database time zone, by default `SYSDATE` and `SYSTIMESTAMP` continue to return date and time in UTC (`SYSDATE_AT_DBTIMEZONE` is `FALSE` by default). If you set `SYSDATE_AT_DBTIMEZONE` to `TRUE` in a session, `SYSDATE` and `SYSTIMESTAMP` return the database time zone.

See *Setting the Database Time Zone* in Oracle Database 19c or Oracle Database 23ai for more information on using the `SET TIME_ZONE` clause with `ALTER DATABASE`.

Using SYSDATE_AT_DBTIMEZONE in a Session

When `SYSDATE_AT_DBTIMEZONE` is `FALSE` in a session, calls to `SYSDATE` and `SYSTIMESTAMP` return values based on the default Autonomous Database time zone, Coordinated Universal Time (UTC).

When `SYSDATE_AT_DBTIMEZONE` is `TRUE` in a session, calls to `SYSDATE` or `SYSTIMESTAMP` return the date and time based on the database time zone.

> **Note:**
>
> Setting `SYSDATE_AT_DBTIMEZONE` to `TRUE` only affects the use of `SYSDATE` and `SYSTIMESTAMP` as operators in application SQL (for example, in queries, DML, and CTAS operations). When using this parameter, it is recommended that your client/session timezone matches your database timezone.

Example

The following example returns dates and times for two different time zones, based on the `SYSDATE_AT_DBTIMEZONE` parameter value:

```
SQL> SELECT DBTIMEZONE FROM DUAL;
```

```
DBTIMEZONE

_____________
-05:00

SQL> ALTER SESSION SET SYSDATE_AT_DBTIMEZONE=FALSE;

Session altered.

SQL> SELECT SYSTIMESTAMP FROM DUAL;

SYSTIMESTAMP

_____________________________________
27-JAN-22 06.59.45.708082000 PM GMT

SQL> ALTER SESSION SET SYSDATE_AT_DBTIMEZONE=TRUE;

Session altered.

SQL> SELECT SYSTIMESTAMP FROM DUAL;

SYSTIMESTAMP

_____________________________________
27-JAN-22 02.14.47.578946000 PM -05:00
```

> **Note:**
>
> When a `SYSDATE` or `SYSTIMESTAMP` query is executed in SQL Worksheet of Database Actions, the time and date value that is returned is in UTC (when `SYSDATE_AT_DBTIMEZONE` parameter is set to `TRUE` or `FALSE`). To obtain the database time zone when working in Database Actions, use `TO_CHAR()` as follows:
>
> ```
> SQL> SELECT TO_CHAR(SYSTIMESTAMP,'YYYY-MM-DD"T"HH24:MI:SS TZH":"TZM')
> FROM DUAL;
>
> TO_CHAR(SYSTIMESTAMP,'YYYY-MM-DD"T"HH24:MI:SSTZH":"TZM')
> ____________________________________________________________
> 2022-01-27T14:15:00 -05:00
> ```

CLIENT_PREFETCH_ROWS

Set the `CLIENT_PREFETCH_ROWS` parameter to enable clients to reduce the number of roundtrips required while fetching rows of a query result set.

`CLIENT_PREFETCH_ROWS` specifies the number of rows to be prefetched by the Oracle client driver, without making any changes to the client application. The client driver buffers the prefetched rows after each successful query execution and for each subsequent fetch request sent to the database.

This parameter applies only to clients that use Oracle Call Interface (OCI) to connect to the database.

This parameter applies only with Oracle Instant Client/Oracle Database Client 19.17 (or later) and 21.8 (or later), for all platforms.

Property	Description
Parameter type	Integer
Syntax	`CLIENT_PREFETCH_ROWS = integer`
Default Value	`0` (Only client-side settings apply)
Modifiable	`ALTER SYSTEM`, `ALTER SESSION`
Range of values	`0` to `UB4MAXVAL` (`4294967295`)
Basic	No

The `CLIENT_PREFETCH_ROWS` parameter can be set with `ALTER SESSION` or `ALTER SYSTEM`. If the parameter value changes using `ALTER SESSION`, the new value becomes effective for that specific session on subsequent resultset fetches. If the parameter value changes using `ALTER SYSTEM`, the new value takes effect for the statements that run on connections created after the `ALTER SYSTEM` command.

For example, if `CLIENT_PREFETCH_ROWS` is set to 100 and a client application asks to fetch 10 rows, a total of 110 rows are returned to the client driver. The first 10 rows out of the 110 rows are given to the application, and the client driver internally buffers the remaining 100 rows. The next 10 row fetches from the client application, each with 10 rows per fetch iteration can be fulfilled from the 100 rows that are internally buffered by the client driver. This process reduces the number of required network roundtrips to and from the database. In this example, on the 11th fetch, a new network roundtrip is incurred and the database returns the next batch of 110 rows, as long as the result set is not exhausted, and the cycle repeats.

Notes for setting `CLIENT_PREFETCH_ROWS`:

- When `CLIENT_PREFETCH_ROWS` is set to a non-zero value, its value takes precedence over the default `OCI_ATTR_PREFETCH_ROWS` value for prefetch row count.
- If the `OCI_ATTR_PREFETCH_ROWS` value is set to a non-default value, then the `CLIENT_PREFETCH_ROWS` value is ignored for the prefetch row count.
- Using `CLIENT_PREFETCH_ROWS` with `OCIAttrSet()`:

  `OCI_ATTR_PREFETCH_ROWS` sets the number of top-level rows to be prefetched. The default value is 1 row. However, if `CLIENT_PREFETCH_ROWS` is set, the number of top-level rows to be prefetched is determined by the following precedence:

  1. If you set the `OCI_ATTR_PREFETCH_ROWS` attribute using `OCIAttrSet()` function or `oraaccess.xml` as the value '1', then the database initialization parameter `CLIENT_PREFETCH_ROWS` value takes precedence and determines the number of top-level rows to be prefetched.
  2. If you set the `OCI_ATTR_PREFETCH_ROWS` attribute using `OCIAttrSet()` function or `oraaccess.xml` as the value '*x*' other than 1, then '*x*' number of top-level rows will be prefetched, and the database initialization parameter `CLIENT_PREFETCH_ROWS` is ignored.
  3. If you do not set an `OCI_ATTR_PREFETCH_ROWS` value using `OCIAttrSet()` or `oraaccess.xml`, then the database initialization parameter `CLIENT_PREFETCH_ROWS` value takes precedence and determines the number of top-level rows to be prefetched.
- Using `CLIENT_PREFETCH_ROWS` with `OCIAttrGet()`:

The function `OCIAttrGet()` returns the effective prefetch row value set from `OCI_ATTR_PREFETCH_ROWS`, `oraaccess.xml` and the database initialization parameter `CLIENT_PREFETCH_ROWS`. If the `OCI_ATTR_PREFETCH_MEMORY` value is set, the value returned by `OCIAttrGet()` might not be the final prefetch rows value and may be restricted to the maximum number of rows allowed by the memory value specified by the `OCI_ATTR_PREFETCH_MEMORY` attribute.

See Also:

- *Oracle Call Interface Developer's Guide* for more information about fetching results and setting the prefetch count.
- OCIAttrSet()
- OCI_ATTR_PREFETCH_ROWS

JOB_QUEUE_PROCESSES

Set the `JOB_QUEUE_PROCESSES` parameter to specify the maximum number of job workers that can be created to run Oracle Scheduler (`DBMS_SCHEDULER`) jobs.

Setting the value to 0 disables non-Oracle supplied Scheduler jobs.

Property	Description
Parameter type	Integer
Syntax	`JOB_QUEUE_PROCESSES = integer`
Default Value	The default value depends on the **ECPU count** (**OCPU count** if your database uses OCPUs) and the setting for **ECPU auto scaling** (**OCPU auto scaling** if your database uses OCPUs). See the following table for details.
Modifiable	`ALTER SYSTEM`
Range of values	Minimum value: `0` Maximum value depends on the **ECPU count** (**OCPU count** if your database uses OCPUs) and the setting for **ECPU auto scaling** (**OCPU auto scaling** if your database uses OCPUs), as shown in the following table.

The default and maximum values for `JOB_QUEUE_PROCESSES` differ depending on the compute model you use. See Compute Models in Autonomous Database for more information.

Compute Model	Default and Maximum Value with Auto Scaling Disabled	Default and Maximum Value with Auto Scaling Enabled
ECPU	7.5 x ECPU count	22.5 x ECPU count
OCPU	30 x OCPU count	90 x OCPU count

Oracle Scheduler job coordinator and job workers are controlled by the `JOB_QUEUE_PROCESSES` parameter. The actual number of job workers created for Oracle Scheduler jobs is auto-tuned by the Scheduler depending on several factors, including available resources, Resource Manager settings, and currently running jobs.

The default value for JOB_QUEUE_PROCESSES provides a compromise between quality of service for applications and reasonable use of system resources. However, it is possible that the default value does not suit every environment.

Setting the value of JOB_QUEUE_PROCESSES to 0 disables non-Oracle supplied Scheduler jobs. When JOB_QUEUE_PROCESSES is set to 0 this does not disable any internal jobs for Oracle-supplied users (service related jobs run by Oracle-supplied users continue to be scheduled).

Oracle-supplied users are users marked as ORACLE_MAINTAINED with value Y. Non Oracle-supplied users are users marked as ORACLE_MAINTAINED with value N.

See ALL_USERS for more information.

LOCKDOWN_ERRORS

Set the LOCKDOWN_ERRORS parameter to control whether Autonomous Database lockdown profile errors are raised or ignored.

Autonomous Database uses lockdown profiles to restrict certain database functionality and features. By default the database raises an error whenever you attempt to run a SQL command that is restricted by the Autonomous Database lockdown profile. The LOCKDOWN_ERRORS parameter allows you to suppress lockdown profile errors. Additionally, when LOCKDOWN_ERRORS is set to the value IGNORE, upon ignoring an error the database logs the details regarding the SQL statement, user, and timestamp for the command.

> **Note:**
>
> Carefully evaluate your existing scripts before you set the value of this parameter to IGNORE. When you set the value to IGNORE, any existing scripts you use will no longer raise lockdown errors.

You can use DBA_LOCKDOWN_ERRORS view to obtain information on lockdown profile errors. See Lockdown Errors View for more information.

Property	Description
Parameter type	String
Syntax	LOCKDOWN_ERRORS = IGNORE | RAISE
Default Value	The default value is RAISE.
Modifiable	ALTER SYSTEM
Valid values	RAISE: Raise lockdown profile errors. IGNORE: Suppress lockdown profile errors.

RESULT_CACHE_INTEGRITY

Set the RESULT_CACHE_INTEGRITY parameter to specify whether the result cache considers queries using possibly non-deterministic constructs as candidates for result caching.

Property	Description
Parameter type	String
Syntax	RESULT_CACHE_INTEGRITY = { ENFORCED | TRUSTED }

Property	Description
Default Value	For Autonomous Database the default value is: `ENFORCED`
Modifiable	`ALTER SYSTEM`

Values

- `ENFORCED`: Irrespective of the setting of `RESULT_CACHE_MODE` or specified hints, only deterministic constructs are eligible for result caching. For example, queries using PL/SQL functions that are not declared as deterministic are never cached (unless the functions are declared as deterministic the query results with such functions will not be cached).
- `TRUSTED`: The database honors the setting of `RESULT_CACHE_MODE` and specified hints and will consider queries using possibly non-deterministic constructs as candidates for result caching. For example, queries using PL/SQL functions that are not declared as deterministic can be cached. Note, however, that results that are known to be non-deterministic are not cached (for example `SYSDATE` or constructs involving `SYSDATE`).

RESULT_CACHE_MODE

Set the `RESULT_CACHE_MODE` parameter to specify which queries are eligible to store result sets in the result cache. Only query execution plans with the result cache operator will attempt to read from or write to the result cache.

Property	Description
Parameter type	String
Syntax	`RESULT_CACHE_MODE = { MANUAL | MANUAL_TEMP | FORCE | FORCE_TEMP }`
Default Value	For Autonomous Database with workload type set to Data Warehouse: `FORCE` For workload types Transaction Processing, JSON, or APEX: `MANUAL`
Modifiable	`ALTER SESSION`, `ALTER SYSTEM`

See RESULT_CACHE_MODE for more information.

- Using SQL Result Cache Hints
  Use result cache hints at the application level to control caching behavior. The SQL result cache hints take precedence over the result cache mode and result cache table annotations.

Using SQL Result Cache Hints

Use result cache hints at the application level to control caching behavior. The SQL result cache hints take precedence over the result cache mode and result cache table annotations.

You can use SQL result cache hints in the following ways:

- Using the `RESULT_CACHE` Hint
- Using the `NO_RESULT_CACHE` Hint
- Using the `RESULT_CACHE` Hint in Views

Using the RESULT_CACHE Hint

When the result cache mode is MANUAL, the /*+ RESULT_CACHE */ hint instructs the database to cache the results of a query block and to use the cached results in future executions.

See Using the RESULT_CACHE Hint and RESULT_CACHE Hint for more information.

Using the NO_RESULT_CACHE Hint

The /*+ NO_RESULT_CACHE */ hint instructs the database not to cache the results in either the server or client result caches.

See Using the NO_RESULT_CACHE Hint and RESULT_CACHE Hint for more information.

Using the RESULT_CACHE Hint in Views

The RESULT_CACHE hint applies only to the query block in which the hint is specified. If the hint is specified only in a view, then only these results are cached.

See Using the RESULT_CACHE Hint in Views and RESULT_CACHE Hint for more information.

SQL Commands

Autonomous Database allows most of the SQL commands available in Oracle Database. To ensure the security and the performance of Autonomous Database, some SQL commands are restricted.

This section provides a list of SQL command limitations that are required to protect security and for the performance integrity of Autonomous Databases. Most of the standard SQL and PL/SQL syntax and constructs available with Oracle Database work in Autonomous Databases.

> **Note:**
>
> If you try to use a restricted SQL command the system reports:
>
> ```
> ORA-01031: insufficient privileges
> ```
>
> This error indicates that you are not allowed to run the SQL command in Autonomous Database.

The following SQL statements are not available in Autonomous Database:

- ADMINISTER KEY MANAGEMENT: By default Autonomous Database uses Oracle-managed encryption keys. Using Oracle-managed keys, Autonomous Database creates and manages the encryption keys that protect your data and Oracle handles rotation of the TDE master key.

  If you want customer-managed keys, a master encryption key in the Oracle Cloud Infrastructure Vault is used to generate the TDE master key on Autonomous Database. See Managing Encryption Keys on Autonomous Database for more information.
- CREATE TABLESPACE, ALTER TABLESPACE, and DROP TABLESPACE: Autonomous Database automatically configures default data and temporary tablespaces for the database. Adding,

removing, or modifying tablespaces is not allowed. Autonomous Database creates one tablespace or multiple tablespaces automatically depending on the storage size.

- `CREATE DATABASE LINK`

  Use `DBMS_CLOUD_ADMIN.CREATE_DATABASE_LINK` to create database links in Autonomous Database. See Use Database Links with Autonomous Database for more information.

- `CREATE LIBRARY`

- `DROP DATABASE LINK`

  Use `DBMS_CLOUD_ADMIN.DROP_DATABASE_LINK` to drop database links in Autonomous Database. See Use Database Links with Autonomous Database for more information.

Roles and Views Restrictions for Data Dictionary

Granting `SELECT ANY DICTIONARY` does not provide access to the `SYS`/`SYSTEM` schemas. You can grant `SELECT_CATALOG_ROLE` to allow `SELECT` privileges on all data dictionary views, if needed.

SQL Statements with Restrictions in Autonomous Database

The following DDL statements are available in Autonomous Database with some restrictions:

SQL Command	Restrictions

`ALTER PLUGGABLE DATABASE` and `ALTER DATABASE`	Only the following clauses are allowed:
	`DEFAULT EDITION`
	`SET TIME_ZONE`
	`SET CMU_WALLET`

`ALTER PROFILE`	Using `ALTER PROFILE`, there are restrictions for a user defined `PASSWORD_VERIFY_FUNCTION`. See Manage Password Complexity on Autonomous Database for more information. Using `ALTER PROFILE`, the optional `CONTAINER` clause is ignored if specified. See Create Users on Autonomous Database for information on the password parameter values defined in the default profile.

`ALTER SESSION`	Only the following clauses are allowed:
	`ADVISE COMMIT`, `ADVISE ROLLBACK`, `ADVISE NOTHING`
	`CLOSE DATABASE LINK`
	`ENABLE COMMIT IN PROCEDURE`, `DISABLE COMMIT IN PROCEDURE`
	`ENABLE PARALLEL <QUERY|DDL|DML>`, `DISABLE PARALLEL <QUERY|DDL|DML>`, `FORCE PARALLEL <QUERY|DDL|DML>`
	`ENABLE RESUMABLE`, `DISABLE RESUMABLE`
	`SET CONSTRAINTS`
	`SET CURRENT_SCHEMA`
	`SET DEFAULT_COLLATION`
	`SET EDITION`
	`SET ISOLATION_LEVEL`
	`SET OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES`
	`SET ROW ARCHIVAL VISIBILITY`
	`SET STATISTICS_LEVEL`
	`SET TIME_ZONE`

SQL Command	Restrictions
ALTER SYSTEM	ALTER SYSTEM is not allowed except ALTER SYSTEM SET and ALTER SYSTEM KILL SESSION SET can only be used to set parameters listed in Initialization Parameters.
ALTER USER	The following clause is ignored: DEFAULT TABLESPACE The IDENTIFIED with the EXTERNALLY clause is not supported. The IDENTIFIED BY VALUES clause is not allowed.
ALTER TABLE	For restrictions, see ALTER TABLE Restrictions.
CREATE PROFILE	PASSWORD_VERIFY_FUNCTION See Manage Password Complexity on Autonomous Database for more information. Using ALTER PROFILE, the optional CONTAINER clause is ignored if specified. See Create Users on Autonomous Database for information on the password parameter values defined in the default profile.
CREATE TABLE	For restrictions, see CREATE TABLE Restrictions.
CREATE OR REPLACE TRIGGER	The AFTER STARTUP and BEFORE SHUTDOWN events are not supported with CREATE OR REPLACE TRIGGER.
CREATE USER	The following clause is ignored: • DEFAULT TABLESPACE IDENTIFIED with the EXTERNALLY clause is not supported. The IDENTIFIED BY VALUES clause is not allowed.

CREATE TABLE Restrictions

XMLType tables using XML schema-based storage are not allowed. See Oracle XML DB for more information.

The clauses not in this list are allowed.

Clause	Comment
cluster	Ignored
ilm_clause	Ignored
inmemory_table_clause	Ignored
LOB_storage_clause	The LOB_compression_clause is recognized. Other LOB_storage_clause parameters are ignored. See LOB_compression_clause for more information.
logging_clause	Ignored
organization external	Ignored
organization index	Creates a regular table with a primary key. Using the organization index clause does not create an index-organized table. You should test and verify the performance of the generated table for your application.
physical_properties	Ignored

> **Note:**
>
> For more information on CREATE TABLE, see *Database SQL Language Reference*.

ALTER TABLE Restrictions

The clauses not in this list are allowed.

Clause	Comment
`allocate_extent_clause`	Ignored
`alter_iot_clauses`	Ignored
`deallocate_unused_clause`	Ignored
`ilm_clause`	Ignored
`inmemory_table_clause`	Ignored
`logging_clause`	Ignored
`modify_LOB_storage_clause`	Ignored
`physical_attributes_clause`	Ignored
`shrink_clause`	Ignored

> **Note:**
>
> For more information on ALTER TABLE, see *Database SQL Language Reference*.

Data Types

Autonomous Database allows most of the data types available in Oracle Database. To ensure the security and the performance of Autonomous Database, some data types are restricted.

The following data types are not supported or have limited support in Autonomous Database:

- Large Object (LOB) data types: only SecureFiles LOB storage is supported. BasicFiles LOBs are automatically converted to SecureFiles LOBs.
- Media types are not supported (Oracle Multimedia is desupported)

Checking and Setting MAX_STRING_SIZE

By default Autonomous Database uses extended data types and the value of `MAX_STRING_SIZE` is set to the value `EXTENDED`. With this setting you can specify a maximum size of 32767 bytes for the `VARCHAR2`, `NVARCHAR2`, and `RAW` data types. The default, `EXTENDED`, is the recommended setting and allows Autonomous Database to take full advantage of database capabilities.

Use `DBMS_MAX_STRING_SIZE` subprograms to check usage of extended data types and to change the database to revert to the older style `STANDARD`, supporting a maximum size of 4000 bytes for `VARCHAR2`, `NVARCHAR2`, and `RAW` data types.

> **Note:**
>
> Using `DBMS_MAX_STRING_SIZE.MODIFY_MAX_STRING_SIZE` is a one-way change that cannot be reverted. After a database is switched back to the `STANDARD` style of supporting a maximum length of 4000 bytes for the `VARCHAR2`, `NVARCHAR2`, and `RAW` data types, you cannot re-enable `EXTENDED` data types.
>
> The ADMIN user is granted `EXECUTE` privilege `WITH GRANT OPTION` clause on `DBMS_MAX_STRING_SIZE`. Oracle recommends that you do not `GRANT EXECUTE` on this package to other users.

1. Check whether your environment can be reverted to the old style, `STANDARD` behavior:

   ```
   SELECT * FROM
   TABLE(DBMS_MAX_STRING_SIZE.CHECK_MAX_STRING_SIZE('STANDARD'));
   ```

   See CHECK_MAX_STRING_SIZE Function for more information.
2. Check and correct all reported violations from Step 1, if applicable.
3. After fixing any reported violations found in Step 1, if you want to revert to a maximum length of 4000 bytes for `VARCHAR2`, `NVARCHAR2`, and `RAW` data types, use `DBMS_MAX_STRING_SIZE.MODIFY_MAX_STRING_SIZE` as follows:

   ```
   EXEC DBMS_MAX_STRING_SIZE.MODIFY_MAX_STRING_SIZE('STANDARD');
   ```

   See MODIFY_MAX_STRING_SIZE Procedure for more information.

See MAX_STRING_SIZE for more information.

See Extended Data Types for details on extended data types.

For a list of Oracle data types see *Oracle Database SQL Language Reference*.

PL/SQL Package Notes for Autonomous Database

Notes for Oracle Database PL/SQL packages in Autonomous Database.

Unavailable PL/SQL Packages

- `DBMS_DEBUG_JDWP`
- `DBMS_DEBUG_JDWP_CUSTOM`

DBMS_LDAP PL/SQL Package Notes

Provides notes for the `DBMS_LDAP` package:

- Specifying an IP address in the host name is not allowed.
- The only allowed port is 636.
- The `SSLWRL` and `SSLWALLETPASSWD` arguments to the `OPEN_SSL` procedure are ignored. The default value for the `SSLWRL` property is set to the wallet that is used by `UTL_HTTP` and `DBMS_CLOUD` for making outbound web requests on Autonomous Database.

- The `DBMS_LDAP.SIMPLE_BIND_S` and `DBMS_LDAP.BIND_S` subprograms perform authentication to the directory server.

  The `DBMS_LDAP.SIMPLE_BIND_S` and `DBMS_LDAP.BIND_S` subprograms are modified to accept credential objects as an argument.

  Following are the usage notes and examples of these modified subprograms:

  – The modified `SIMPLE_BIND_S` and `BIND_S` subprograms enable you to pass credential objects to set directory server authentication. Credential objects are schema objects, hence they can be accessed only by privileged users and enable you to configure schema-level privileges to access control the credentials. Passing scheduler credentials is an appropriate and secure way to store and manage username/password/keys for authentication.

  – The modified `SIMPLE_BIND_S` and `BIND_S` subprograms are a secure and convenient alternative to previously existed `SIMPLE_BIND_S` and `BIND_S` subprogram.

    See FUNCTION simple_bind_s and FUNCTION bind_s for more information.

  – The `CREDENTIAL` argument of the `SIMPLE_BIND_S` and `BIND_S` functions is used to perform credential based authentication to the directory server.

  – For example:

    * Create a credential object:

```
BEGIN DBMS_CLOUD.CREATE_CREDENTIAL (
    credential_name => 'LDAP_CRED',
    username        => 'web_app_user',
    password        => 'password' );
END;
```

      This creates a credential object which creates a stored username/password pair.

      See CREATE_CREDENTIAL Procedure for more information.

      See Specifying Scheduler Job Credentials for more information.

    * Invoke `DBMS_LDAP.SIMPLE_BIND_S`:

```
DECLARE
      l_mail_conn DBMS_LDAP.INIT;
    BEGIN
      l_ldap_conn   := DBMS_LDAP.INIT('ldap.example.com', 636);
      l_auth_result := DBMS_LDAP.SIMPLE_BIND_S(l_ldap_conn,
'LDAP_CRED');
      ...
END;
```

      The code in this example first invokes the `DBMS_LDAP.INIT` function which initializes a session with an LDAP server and establishes a connection with the LDAP server `ldap.example.com` at port number `636`. The value `l_ldap_conn` in the `SIMPLE_BIND_S` function is the LDAP session handle and `LDAP_CRED` is the credentials name.

* The `function bind_s` performs complex authentication to the directory server. For example:

```
DECLARE
      l_mail_conn DBMS_LDAP.INIT;
    BEGIN
      l_ldap_conn   := DBMS_LDAP.INIT('ldap.example.com', 636);
      l_auth_result := DBMS_LDAP.BIND_S(l_ldap_conn, 'LDAP_CRED',
METH => DBMS_LDAP.AUTH_SIMPLE);
      ...
END;
```

The code in this example first invokes the `DBMS_LDAP.INIT` function which initializes a session with an LDAP server and establishes a connection with the LDAP server `ldap.example.com` at port number `636`. The value `l_ldap_conn` in the `BIND_S` function is the LDAP session handle and `LDAP_CRED` is the credentials name. `METH` is the authentication method. The only valid value is `DBMS_LDAP_UTL.AUTH_SIMPLE`.

– The `EXECUTE` privileges on `DBMS_CLOUD` or `DWROLE` is required to create scheduler credentials.

– The passed credential must be present in the current schema or have a synonym pointing to a credential in the schema and be in the enabled state.

– A public or private synonym that points to a credential in a different user schema can be supplied as a value for the `CREDENTIAL` parameter provided you have the `EXECUTE` privilege on the base credential object pointed to by the synonym. See Overview of Synonyms for more information.

- SSL/TLS is enforced for all communication happening between LDAP server and Autonomous Database.
- When your Autonomous Database instance is configured with a private endpoint, set the `ROUTE_OUTBOUND_CONNECTIONS` database parameter to '`PRIVATE_ENDPOINT`' to specify that all outgoing LDAP connections are subject to the Autonomous Database instance private endpoint VCN's egress rules. See Enhanced Security for Outbound Connections with Private Endpoints for more information.
- To use `DBMS_LDAP` for a connection on a private endpoint, use `DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE` and specify the `private_target` parameter with value `TRUE`.

> **Note:**
>
> If you set `ROUTE_OUTBOUND_CONNECTIONS` to `PRIVATE_ENDPOINT`, setting the `private_target` parameter to `TRUE` is not required in this API. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

- **`DBMS_LDAP` Error**

Error Message	Potential Cause
`ORA-31400: Missing or invalid scheduler credential`	Passed credentials are NULL or invalid.

See DBMS_LDAP in *PL/SQL Packages and Types Reference* for more information.

DBMS_NETWORK_ACL_ADMIN PL/SQL Package Notes

Provides notes for the `DBMS_NETWORK_ACL_ADMIN` package:

- Granting ACL privileges on IP addresses is not allowed.
- The `HTTP_PROXY` ACL is allowed on private endpoints.

See DBMS_NETWORK_ACL_ADMIN in *PL/SQL Packages and Types Reference* for more information.

UTL_HTTP PL/SQL Package Notes

Provides notes for the `UTL_HTTP` package:

- Connections through IP addresses are not allowed.
- Only `HTTPS` is allowed when the Autonomous Database instance is on a public endpoint. When the Autonomous Database instance is on a private endpoint, both `HTTPS` and `HTTP_PROXY` connections are allowed (`HTTP` connections are disallowed for both public endpoints and private endpoints).
- The `UTL_HTTP.set_proxy` API is allowed when the Autonomous Database instance is on a private endpoint.
- When the Autonomous Database instance is on a private endpoint and you use `HTTP_PROXY` or the `UTL_HTTP.SET_PROXY` API:
  - `DBMS_CLOUD` requests do not honor the proxy server you set with `UTL_HTTP.SET_PROXY`. This includes `DBMS_CLOUD.SEND_REQUEST` and all object storage access for `DBMS_CLOUD` external tables that you define with `DBMS_CLOUD.CREATE_EXTERNAL_TABLE`, `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE`, or `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE`.
  - `APEX_WEB_SERVICE` requests do not honor the proxy server you set with `UTL_HTTP.SET_PROXY`.
- All web services must be secured. The only allowed port is 443 when the Autonomous Database instance is on a public endpoint. When the Autonomous Database instance is on a private endpoint this restriction does not apply.

  Your Autonomous Database instance is preconfigured with an Oracle Wallet that contains more than 90 of the most commonly trusted root and intermediate SSL certificates. The Oracle Wallet is centrally managed. You can configure `UTL_HTTP` to use a wallet for a site that is protected using self-signed SSL certificates. See Use a Customer-Managed Wallet for External Calls with UTL_HTTP for more information.
- The `SET_AUTHENTICATION_FROM_WALLET` procedure is disallowed.
- The `WALLET_PATH` and `WALLET_PASSWORD` arguments for the `CREATE_REQUEST_CONTEXT`, `REQUEST`, and `REQUEST_PIECES` procedures are ignored.
- The `CREDENTIAL` argument of the `SET_CREDENTIAL` procedure is used to pass the credential object as an input to the procedure. See Specifying Scheduler Job Credentials and CREATE_CREDENTIAL Procedure for more information.
- The `EXECUTE` privileges on `DBMS_CLOUD` or `DWROLE` is required to create credential objects.
- The passed credentials must be present in the current user schema and be in the enabled state.

- A public or private synonym that points to a credential in a different user schema can be supplied as a value for the `CREDENTIAL` parameter provided you have the `EXECUTE` privilege on the base credential object pointed to by the synonym. See Overview of Synonyms for more information.
- Oracle Wallet configuration cannot be altered. All arguments for `SET_WALLET` procedure are ignored.
- When your Autonomous Database instance is configured with a private endpoint, set the `ROUTE_OUTBOUND_CONNECTIONS` database parameter to '`PRIVATE_ENDPOINT`' to specify that all outgoing `UTL_HTTP` connections are subject to the Autonomous Database instance private endpoint VCN's egress rules. See Enhanced Security for Outbound Connections with Private Endpoints for more information.
- **`UTL_HTTP` Errors**

  The following table shows error messages and possible causes for these error messages when using `UTL_HTTP`:

Error Message	Potential Cause
`ORA-12545: Connect failed because target host or object does not exist`	Target host or object does not exist or it is private.
`ORA-24247: network access denied by access control list (ACL)`	Access control list (ACL) for the specified host could not be found.
`ORA-29024: Certificate validation failure`	Certificate of the host does not exist or is not among the supported certificates.
`ORA-29261: Bad argument`	Passed credentials are invalid or disabled or the user does not have sufficient privileges on the credential.

See UTL_HTTP in *PL/SQL Packages and Types Reference* for more information.

UTL_INADDR PL/SQL Package Notes

Provides notes for the `UTL_INADDR` package:

- The `UTL_INADDR` package is available for use on an Autonomous Database instance with a private endpoint (PE).
- The **`GET_HOST_ADDRESS`** function is available.
- The **`GET_HOST_NAME`** function is not available.

See UTL_INADDR in *PL/SQL Packages and Types Reference* for more information.

UTL_SMTP PL/SQL Package Notes

Provides notes for the `UTL_SMTP` package:

- Public SMTP endpoints can be used as long as routing is done through customer reverse connection endpoint (RCE).
- Mail with an IP address in the host name is not allowed.
- The only allowed ports are 25 and 587.
- The `CREDENTIAL` argument of the `SET_CREDENTIAL` function is used to pass the scheduler credentials object as an input to the function. See Specifying Scheduler Job Credentials and CREATE_CREDENTIAL Procedure for more information.

- The `EXECUTE` privileges on `DBMS_CLOUD` or `DWROLE` is required to create credential objects.
- The `CREDENTIAL` argument of the `SET_CREDENTIAL` procedure is used to pass the credential objects object as an input to the procedure. See Specifying Scheduler Job Credentials for more information.
- The passed credentials must be present in the current user schema and be in the enabled state.
- A public or private synonym that points to a credential in a different user schema can be supplied as a value for the `CREDENTIAL` parameter provided you have the `EXECUTE` privilege on the base credential object pointed to by the synonym. See Overview of Synonyms for more information.
- When your Autonomous Database instance is configured with a private endpoint, set the `ROUTE_OUTBOUND_CONNECTIONS` database parameter to '`PRIVATE_ENDPOINT`' to specify that all outgoing `UTL_SMTP` connections are subject to the Autonomous Database instance private endpoint VCN's egress rules. See Enhanced Security for Outbound Connections with Private Endpoints for more information.
- **`UTL_SMTP` Error**

Error Message	Potential Cause
`ORA-29261: Bad argument`	Passed credentials are invalid or disabled or the user does not have sufficient privileges on the credential.

See UTL_SMTP in *PL/SQL Packages and Types Reference* for more information.

UTL_TCP PL/SQL Package Notes

Provides notes for the `UTL_TCP` package:

- The IP address is not allowed in the host name.
- The only allowed ports are: 443 (HTTP) 25 and 587 (SMTP).
- For port 443, only HTTPS URLs are allowed.
- The `WALLET_PATH` and `WALLET_PASSWORD` arguments for the `OPEN_CONNECTION` procedure are ignored. The default value for the `WALLET_PATH` and `WALLET_PASSWORD` property are set to the wallet that is used by `UTL_HTTP` and `DBMS_CLOUD` for making outbound web requests on Autonomous Database.
- SSL/TLS is enforced for all communication happening over TCP/IP connections.
- When your Autonomous Database instance is configured with a private endpoint, set the `ROUTE_OUTBOUND_CONNECTIONS` database parameter to '`PRIVATE_ENDPOINT`' to specify that all outgoing `UTL_TCP` connections are subject to the Autonomous Database instance private endpoint VCN's egress rules. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

See UTL_TCP in *PL/SQL Packages and Types Reference* for more information.

ADMIN User Roles and Privileges

In Autonomous Database the predefined administrative user is ADMIN and this account has privileges to manage users and to manage the database.

It is recommended to use the ADMIN user to create accounts and grant privileges for users to connect to the database. See Manage Users for more information.

To maintain security, the privilege SYSDBA is not granted to the ADMIN user; a limited number of system privileges are granted to ADMIN. The following query shows privileges that are granted to SYS that are not granted to the ADMIN user:

```
SELECT privilege FROM dba_sys_privs WHERE grantee='SYS' MINUS
    SELECT privilege FROM dba_sys_privs WHERE grantee = 'ADMIN' ORDER BY 1;
```

ADMIN User GRANT ANY Restrictions

The `GRANT ANY PRIVILEGE` system privilege is not available for the ADMIN user. Instead use `GRANT ANY OBJECT PRIVILEGE`, `GRANT ANY SCHEMA PRIVILEGE`, or `GRANT ANY ROLE`.

System Privilege	Description
`GRANT ANY OBJECT PRIVILEGE`	This allows granting object privileges on objects including those owned by SYS with a few exceptions. In Autonomous Database, `GRANT ANY OBJECT PRIVILEGE` can only be exercised against customer created user schemas and can't be exercised against Oracle maintained schema such as SYSTEM or SYS, and many others, to ensure security. `ORA-1031/942` errors indicate that a privilege cannot be granted by ADMIN.
`GRANT ANY PRIVILEGE`	This allows granting of all system privileges excluding administrative privileges such as SYSDBA. The ADMIN user does not have SYSDBA privileges (instead a list of system privileges is granted). Use the following query to list the ADMIN privileges: `SELECT privilege FROM dba_sys_privs WHERE grantee = 'ADMIN' ORDER BY 1;`
`GRANT ANY ROLE`	This allows granting of roles to users and user roles. In Autonomous Database, `GRANT ANY ROLE` can only be exercised against customer created roles. The privileges that are NOT granted with `WITH ADMIN OPTION` to ADMIN user, such as `DBA`, `EXP_FULL_DATABASE`, and others can't be granted by the ADMIN user. `ORA-1031` errors indicate that a rol cannot be granted by ADMIN.

Roles and Views Restrictions for Data Dictionary

Granting `SELECT ANY DICTIONARY` does not provide access to the `SYS`/`SYSTEM` schemas. You can grant `SELECT_CATALOG_ROLE` to allow `SELECT` privileges on all data dictionary views, if needed.

Oracle XML DB

Describes Autonomous Database support for Oracle XML DB features. To ensure the security and the performance of your Autonomous Database, some Oracle XML DB features are restricted.

The following is supported, in addition to the features listed:

- Full support for XMLQuery, XMLTable, and other SQL/XML standard functions
- Indexing schema including functional indexes using SQL/XML expressions, Structured XMLIndex and XQuery Full Text Index

> **Note:**
>
> If you migrate tables containing `XMLType` columns to Autonomous Database using Oracle Data Pump, you need to convert to Non-Schema Binary XML prior to using Oracle Data Pump Export (expdp).

Area	XML DB Feature	Supported in Autonomous Database	More Information
Repository	XML DB Protocol	No	Repository Access Using Protocols
Repository	XML DB Resources	No	Oracle XML DB Repository Resources
Repository	XML DB ACLs	No	Repository Access Control
Storage	XML Schema Registration	No [1]	XML Schema Registration with Oracle XML DB
Storage	CLOB	No	Deprecated
Storage	Object Relational	No	XML Schema and Object-Relational XMLType
Storage	Binary XML	Yes (Non schema-based only)	XMLType Storage Models
Index	Structured XML Index	Yes	XMLIndex Structured Component
Index	XQuery Full Text Index	Yes	Indexing XML Data for Full-Text Queries
Index	Unstructured XMLIndex	No	XMLIndex Unstructured Component
Packages	XML DOM package	Yes	PL/SQL DOM API for XMLType (DBMS_XMLDOM)
Packages	XML Parser Package	Yes	PL/SQL Parser API for XMLType (DBMS_XMLPARSER)
Packages	XSL Processor (DBMS_XSLPROCESSOR)	Yes	PL/SQL XSLT Processor for XMLType (DBMS_XSLPROCESSOR)

[1] While you cannot do XML schema registration, runtime validation of XML documents is available through DBMS_XMLSCHEMA_UTIL.

For details on Oracle XML DB, see Oracle XML DB Developer's Guide.

Oracle Text

Describes Autonomous Database support for Oracle Text features. To ensure the security and the performance of your Autonomous Database, some Oracle Text features are restricted.

Oracle Text Feature	Supported in Autonomous Database	More Information
All logging, and APIs which perform logging such as `ctx_report.query_log_summary`	Not Supported	QUERY_LOG_SUMMARY
File Datastore	Not Supported (see replacement `DIRECTORY_DATASTORE`)	Datastore Types

Oracle Text Feature	Supported in Autonomous Database	More Information
URL Datastore	Not Supported (see replacement `NETWORK_DATASTORE`)	Datastore Types
`CREATE INDEX` with `BIG_IO` option	Supported if you grant the privilege to create a trigger to the user (`GRANT CREATE TRIGGER`).	Improved Response Time Using the BIG_IO Option of CONTEXT Index
`OPTIMIZE_INDEX` in rebuild mode	Supported if you grant the privilege to create a trigger to the user (`GRANT CREATE TRIGGER`).	OPTIMIZE_INDEX

For details on Oracle Text, see Oracle Text Application Developer's Guide.

Oracle Flashback

Oracle Flashback Technology is a group of Oracle Database features that let you view past states of database objects or to return database objects to a previous state without using point-in-time media recovery.

To restore and recover your database to a point in time, see Restore and Recover your Autonomous Database.

Oracle Flashback Feature	Supported in Autonomous Database
`DBMS_FLASHBACK`	Yes except the procedure: `DBMS_FLASHBACK.TRANSACTION_BACKOUT`
Flashback Data Archive	Yes
Flashback Drop	Yes
Flashback Query	Yes
Flashback Table	Yes
Flashback Transaction	No
Flashback Transaction Query	Yes
Flashback Version Query	Yes

See Track Table Changes with Flashback Time Travel for information on using Flashback features.

Oracle Database Real Application Security

Oracle Database Real Application Security is a database authorization model that: supports declarative security policies, enables end-to-end security for multitier applications, provides an integrated solution to secure database and application resources, and advances the security architecture of Oracle Database to meet existing and emerging demands of applications developed for the Internet.

See Introducing Oracle Database Real Application Security more information.

Real Application Security works the same on Autonomous Database as on an on-premises Oracle Database except you need to perform the following ADMIN tasks before using Real Application Security on Autonomous Database:

- To create Real Application Security users/roles, you need the `PROVISION` system privilege. As the ADMIN user run the following command to grant this privilege to a database user:

  ```
  SQL> EXEC
  XS_ADMIN_CLOUD_UTIL.GRANT_SYSTEM_PRIVILEGE('PROVISION','DB_USER');
  ```

  In this example, *`DB_USER`* is a database user.

  Running this command on Autonomous Database replaces the following on-premise database command (note the `_CLOUD_` is not in the following package name):

  ```
  SQL> EXEC SYS.XS_ADMIN_UTIL.GRANT_SYSTEM_PRIVILEGE('PROVISION', 'DB_USER',
  XS_ADMIN_UTIL.PTYPE_XS);
  ```

  See General Procedures for Creating Application User Accounts for more information.

- To create Real Application Security data controls, you need the `ADMIN_ANY_SEC_POLICY` privilege. As the ADMIN user run the following command to grant this privilege:

  ```
  EXEC
  XS_ADMIN_CLOUD_UTIL.GRANT_SYSTEM_PRIVILEGE('ADMIN_ANY_SEC_POLICY','DB_USER'
  );
  ```

  In this example, *`DB_USER`* is a database user.

  Running this command on Autonomous Database replaces the following on-premise database command (note the `_CLOUD_` is not in the following package name):

  ```
  SQL> EXEC
  SYS.XS_ADMIN_UTIL.GRANT_SYSTEM_PRIVILEGE('ADMIN_ANY_SEC_POLICY','DB_USER');
  ```

  See Creating Roles and Application Users for more information.

Oracle LogMiner

Describes restrictions for Oracle LogMiner on Autonomous Database.

Archived Log Retention Maximum is 48 Hours

Autonomous Database archived log files are kept for up to 48 hours. LogMiner can only access up to 48 hours of archived log files.

If you attempt to mine log files older than 48 hours, LogMiner reports `ORA-1285: "error reading file"`.

Choose a Character Set for Autonomous Database

The Autonomous Database default database character set is Unicode AL32UTF8 and the default national character set is AL16UTF16. When you provision a database, depending on the workload type, you can select a database character set and a national character set.

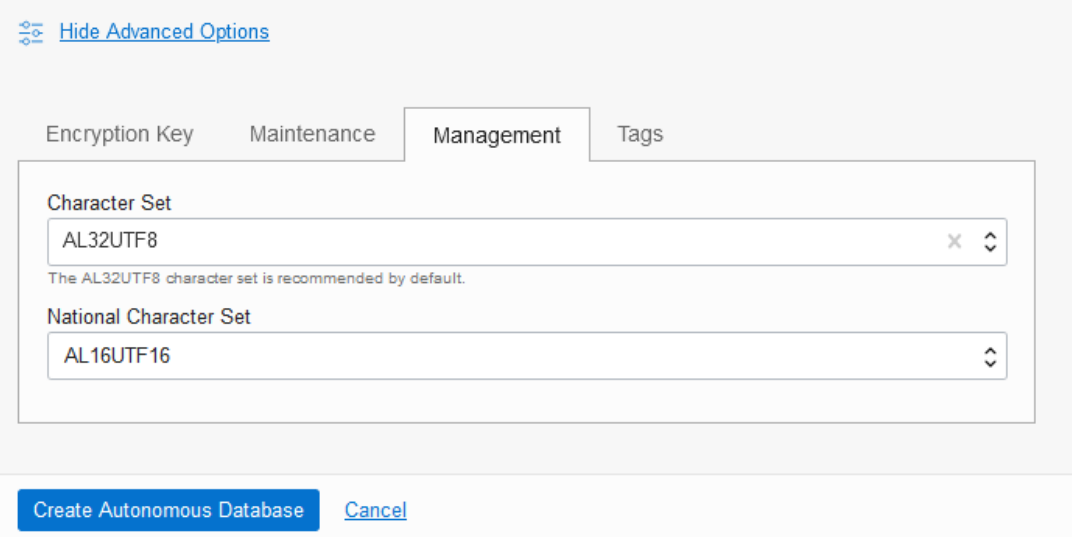
> **Note:**
>
> Oracle recommends using the default Unicode database character set (AL32UTF8) for its universality and compatibility with contemporary and future technologies and language requirements.

If you are using an on-premises database with a non-Unicode character set, migrating to the default Unicode character set can be a convoluted process requiring complex data analysis. Thus, Autonomous Database lets you choose a character set when you provision an Autonomous Database instance.

1. Create an Autonomous Database following the provisioning steps.

   See Provision an Autonomous Database Instance for more information.

2. On the **Create Autonomous Database** page expand **Show Advanced Options**.
3. In the advanced options area, select the **Management** tab.



   - In the **Character Set** field, use the selector to choose a character set.
   - In the **National Character Set** field, use the selector to choose a national character set.

   To make selecting a value easier, when you type in the text area this filters the list. For example, if you type JA you see only the options containing JA, including: JA16EUC, JA16EUCTILDE, JA16SJIS, JA16JA16SJISTILDE, and JA16VMS.

4. Click **Create Autonomous Database** to provision the Autonomous Database.

See the following for more information:

- Choosing an Oracle Database Character Set

- Character Set Migration
- Support Note 788156.1
- Notes for Character Set Selection
  Provides notes and limitations for selecting a character set and a national character set on Autonomous Database.

Notes for Character Set Selection

Provides notes and limitations for selecting a character set and a national character set on Autonomous Database.

- When you provision an Autonomous Database instance you can only select a character set for **Data Warehouse**, **Transaction Processing**, or **APEX** workload types.

  The **JSON Database** workload type only supports the default character set and you cannot select a different character set.
- Always Free Autonomous Database does not support character set selection.
- You cannot select a different character set when you clone an instance. A cloned Autonomous Database instance has the same character set as the source database.
- You cannot change the character set of an existing Autonomous Database instance.
- APEX developer and administration pages are supported in English only when the database character set is different from the default (AL32UTF8).

  Translation of user applications into other languages or setting an application's primary language to a value other than English is not supported. If you use APEX in a language other than English, you may experience issues such as illegible, garbage text. You can, however, use your APEX applications to process non-English user data, as long as the languages of the data are supported by the selected database character set.

  If you need full globalization support in APEX, migrate your applications and data to AL32UTF8, the universal Unicode character set, which is the default and recommended character set for Autonomous Database.

  When you choose a character set other than the default Unicode character set (AL32UTF8) on Autonomous Database, the APEX language selector only shows languages that are supported in that database character set. For example, if you choose the database character set WE8ISO8859P1 (ISO 8859-1 West European), the language selector does not show Japanese or Korean.
- While you are provisioning an instance the **Character Set** selector on the **Management** tab lists the supported character set names. If you want to see a list of supported database character sets before you provision a database, refer to the following:
  - See Table A-4 in Recommended Database Character Sets
  - See Table A-6 in Other Character Sets

Query the Alert Log

You can query the database alert log using `V$DIAG_ALERT_EXT` to search for log data.

`V$DIAG_ALERT_EXT` shows the contents of the XML-based alert log in the Automatic Diagnostic Repository (ADR). For example, if you want to search for `ORA` errors or client connection tracing, you can query `V$DIAG_ALERT_EXT`. Also you can query `V$DIAG_ALERT_EXT` if you want to pull the data from these logs to another analytic solution for tracing `ORA` errors, such as `ORA-600` or `ORA-7445`.

Database Features Unavailable in Autonomous Database

Lists the Oracle Database features that are not available in Autonomous Database. Additionally, database features designed for administration are not available.

List of Unavailable Oracle Features

- Oracle Real Application Security Administration Console (RASADM)
- Oracle Industry Data Models
- Oracle Database Lifecycle Management Pack
- Oracle Data Masking and Subsetting Pack

> **Note:**
>
> Oracle Data Safe, available with Autonomous Database provides Data Masking. See Use Oracle Data Safe with Autonomous Database for more information.

- Oracle Cloud Management Pack for Oracle Database
- Oracle Multimedia: Not available in Autonomous Database and deprecated in Oracle Database 18c.
- Oracle Sharding
- Oracle XStream
- Custom locale objects, including: language, territory, character set, and collation (linguistic sort) are not supported. Custom locale definitions created using Oracle Locale Builder cannot be deployed on Autonomous Database. See Customizing Locale Data for more information.
- Original export/import methods are not supported in Autonomous Database. See Migrate Oracle Databases to Autonomous Database for information on recommended migration methods.

Logical Partition Change Tracking and Materialized Views

Describes information about the Logical Partition Change Tracking (LPCT) metadata framework and Query Rewrite with Logical Partition Change Tracking in Autonomous Database.

- About Logical Partition Change Tracking
- Using Logical Partition Change Tracking
  Logical Partition Change Tracking (LPCT) logically partitions a table using a specified key column and method.
- Example: Logical Partition Change Tracking
  Shows the steps to use Logical Partition Change Tracking (LPCT) using a Materialized View that contains joins and aggregates.

About Logical Partition Change Tracking

Logical Partition Change Tracking (LPCT) tracks the staleness of materialized views.

Logical Partition Change Tracking enables you to create logical partitions on base tables. It evaluates the staleness of the base tables for individual logical partitions without using a materialized view log or requiring any of the tables used in the materialized view to be partitioned.

When one or more dependent base tables of a materialized view are updated, a materialized view becomes `STALE` and cannot be used for query rewrite under the default enforced integrity mode.

Logical Partition Change Tracking (LPCT) provides the capability to leverage the user-supplied logical partitioning information of base tables of a materialized view for a more fine-grained, partition-level tracking of stale data for both refresh and rewrite purposes. While classical Partitioning Change Tracking (PCT) relies on the physical partitioning of tables, LPCT has no dependency on tables being physically partitioned; LPCT can be used with both partitioned and nonpartitioned tables.

Logical Partition Change Tracking mechanism makes use of the `FRESH` subsets (partitions) of materialized views despite other subsets being `STALE`. Faster response times can be achieved for user queries because pre-computed results in materialized views are used more often. Other than increasing the usability of materialized views, PCT and LPCT also allows incremental refreshing of the materialized views without the need of materialized view logs; refresh can be both `ON DEMAND` or `ON COMMIT`.

Similar to Partitioning Change Tracking (PCT), Logical Partition Change Tracking (LPCT) is associated with a base table and can accurately identify the rows in a Materialized View affected by data changes on the base table, according to the logical partition boundaries defined.

See Advanced Materialized Views Advanced Materialized Views for more information.

Using Logical Partition Change Tracking

Logical Partition Change Tracking (LPCT) logically partitions a table using a specified key column and method.

Logical Partition Change Tracking creation syntax is analogous to physical partitions. Unlike physical partitions, which must be created as part of table creation, LPCT can be freely specified independent of the table creation and its shape, allowing more flexibility to address your requirements. LPCT creation is metadata only.

- Creating Logical Partitions - BNF
  Describes the syntax to create BNF logical partitions.
- Choosing the Logical Partition Key Column
  The logical partitioning key is specified to define the boundaries of each logical partition.
- Freshness of Materialized Views Using Logical Partition Change Tracking
  The Logical Partition Change Tracking (LPCT) staleness tracking mechanism automatically records and consolidates the change statistics internally based on the specified logical partition key and partitioning method during each data change.
- Rewrite with Materialized Views Using Logical Partition Change Tracking
  Using Logical Partition Change Tracking (LPCT), Oracle knows that a materialized view is `STALE` with respect to some logical partitions of the base table, but `FRESH` with respect to other portions.
- Refresh of Materialized Views Using Logical Partition Change Tracking
  Logical Partition Change Tracking (LPCT) refresh can be implemented using the finer-grained data staleness to incrementally refresh `STALE` subsets of a materialized view, eliminating costly complete refresh or log-based fast refresh.

- Logical Partition Change Tracking – Data Dictionary Views
  Describes the data dictionary views to find information about logical partitions.

Creating Logical Partitions - BNF

Describes the syntax to create BNF logical partitions.

Following is the syntax to create BNF logical partitions:

```
CREATE LOGICAL PARTITION TRACKING ON table_name
    PARTITION BY RANGE (partition_key)
    INTERVAL (interval_clause)
    (partition_specification);
```

- Only `RANGE` and `INTERVAL` logical partitioning methods are supported.
- Only a single logical partition key column is supported.
- The partition key column can be of these data types:
  - `NUMBER`
  - `DATE`
  - `CHAR`
  - `VARCHAR`
  - `VARCHAR2`
  - `TIMESTAMP`
  - `TIMESTAMP WITH TIME ZONE`

Choosing the Logical Partition Key Column

The logical partitioning key is specified to define the boundaries of each logical partition.

The logical partition key is not physical, this means that table rows belonging to a key range are not segregated into a separate physical partition. The table can be non-partitioned or partitioned on a key that is different from the logical partition key. The logical partition key can be chosen freely, and partition boundaries can be made flexible.

To choose a Logical Partition Change Tracking (LPCT) key column, you can consider a clustered column, that is, a column where data is close to sorted by column value, that are frequently referenced in the query filter predicates. For a clustered column, less logical partitions are likely to be affected during data loads, this means that less `STALE` logical partitions need to be refreshed and more `FRESH` logical partitions are ready to be used for rewrites. If a table is already partitioned, it is recommended to create a LPCT using a different column other than the partition key column. LPCT offers similar benefits as Partitioning Change Tracking (PCT), and the combined benefits are not maximized if data tracking is done on the same column.

Freshness of Materialized Views Using Logical Partition Change Tracking

The Logical Partition Change Tracking (LPCT) staleness tracking mechanism automatically records and consolidates the change statistics internally based on the specified logical partition key and partitioning method during each data change.

Adjacent change data is grouped into a “logical” partition. Unlike with Partitioning Change Tracking (PCT), which is tied to physical partition bounds the LPCT scheme offers flexibility in managing and grouping the data changes resulting from DMLs applied to the base table.

During conventional DMLs and Direct-loads, LPCT adopts the same algorithm that PCT uses to track staleness. During Query rewrites, LPCT adopts the same algorithm that PCT uses to calculate rewrite containment.

When a table is logically partitioned using key ranges, a materialized view defined on the table is eligible to use LPCT for staleness tracking, refresh and query rewrite, provided that the materialized view contains the logical partition key.

> **Note:**
>
> All types of Materialized Views are supported for LPCT.

Rewrite with Materialized Views Using Logical Partition Change Tracking

Using Logical Partition Change Tracking (LPCT), Oracle knows that a materialized view is `STALE` with respect to some logical partitions of the base table, but `FRESH` with respect to other portions.

Having the finer-grained data staleness information of the base tables, the associated materialized view would be used more frequently due to LPCT rewrite.

Oracle transparently identifies and makes use of the `FRESH` subset of materialized views for query rewrite to answer complicated queries of base tables when `QUERY_REWRITE_INTEGRITY = ENFORCED |TRUSTED`.

If the materialized view rows are partially `FRESH` with respect to those logical partitions, a partial rewrite might take place to answer the query partially using materialized view, that is, `FRESH` logical partitions, and partially using the base table, that is, the `STALE` logical partitions.

Refresh of Materialized Views Using Logical Partition Change Tracking

Logical Partition Change Tracking (LPCT) refresh can be implemented using the finer-grained data staleness to incrementally refresh `STALE` subsets of a materialized view, eliminating costly complete refresh or log-based fast refresh.

If LPCT refresh is specified, the `STALE` logical partitions are identified and targeted refresh operations will be performed to those logical partitions only.

To invoke refresh using logical partition change tracking you specify `'L'` or `'l'` (“logical”) as refresh method.

For example: `execute DBMS_MVIEW.REFRESH(<materialized_view_name>,'L');`

If `REFRESH FORCE` is specified, a `FAST` refresh is chosen and performed if possible, or else it performs a `COMPLETE` refresh. During materialized view `FORCE` refresh, LPCT refresh has the same priority as Partitioning Change Tracking (PCT) refresh.

Logical Partition Change Tracking – Data Dictionary Views

Describes the data dictionary views to find information about logical partitions.

Query the following data dictionary views to retrieve information about logical partitions.

- `ALL_MVIEW_DETAIL_LOGICAL_PARTITION`: This view displays the freshness information of the materialized views, with respect to an LPCT detail logical partition, accessible to the current user. See ALL_MVIEW_DETAIL_PARTITION for more information.
- `DBA_MVIEW_DETAIL_ LOGICAL_PARTITION`: displays freshness information for all materialized views in the database, with respect to a LPCT detail logical partition. See DBA_MVIEW_DETAIL_PARTITION for more information.
- `USER_MVIEW_DETAIL_ LOGICAL_PARTITION`: displays freshness information for all materialized views, with respect to a LPCT detail logical partition, owned by the current user. See USER_MVIEW_DETAIL_PARTITION for more information.

Example: Logical Partition Change Tracking

Shows the steps to use Logical Partition Change Tracking (LPCT) using a Materialized View that contains joins and aggregates.

1. Create base tables with logical change partitions.
   a. Create `MYSALES` table.

```
CREATE TABLE mysales ( time_id DATE, prod_id NUMBER, cust_id NUMBER,
channel_id CHAR(1), promo_id NUMBER(6), quantity_sold
      NUMBER(3), amount_sold NUMBER(10,2))
      PARTITION BY LIST (prod_id)
      (PARTITION p1 VALUES (1,2,3),
      PARTITION p2 VALUES (4,5,6),
      PARTITION p3 VALUES (7,8,9),
      PARTITION p4 VALUES(DEFAULT));
```

   This creates the `MYSALES` table.

   b. Insert records into the `MYSALES` table.

```
INSERT INTO mysales (time_id, prod_id, cust_id, amount_sold) VALUES
(TO_DATE('2007-06-05','yyyy-mm-dd'), 1, 2088, 189.98);
INSERT INTO mysales (time_id, prod_id, cust_id, amount_sold) VALUES
(TO_DATE('2007-07-05','yyyy-mm-dd'), 2, 1354, 12.99);
INSERT INTO mysales (time_id, prod_id, cust_id, amount_sold) VALUES
(TO_DATE('2007-09-05','yyyy-mm-dd'), 5, 2088, 189.98);
INSERT INTO mysales (time_id, prod_id, cust_id, amount_sold) VALUES
(TO_DATE('2007-10-05','yyyy-mm-dd'), 18, 2088, 42);
COMMIT;
```

   This populates the `MYSALES` table.

c. Create logical partition tracking for the MYSALES table.

```
CREATE LOGICAL PARTITION TRACKING ON mysales
  PARTITION BY RANGE (time_id)
  INTERVAL(NUMTOYMINTERVAL(2, 'YEAR'))
  (
   PARTITION p0 VALUES LESS THAN (TO_DATE('7-15-2005', 'MM-DD-YYYY')),
   PARTITION p1 VALUES LESS THAN (TO_DATE('7-15-2007', 'MM-DD-YYYY'))
  );
```

This creates a logical partition tracking for the MYSALES table using the TIME_ID as key.

d. Create MYCUSTOMERS table.

```
CREATE TABLE mycustomers (cust_id NUMBER, age NUMBER, gender CHAR(1),
address VARCHAR(100));
```

This creates the MYCUSTOMERS table.

e. Insert records into the MYCUSTOMERS table.

```
INSERT INTO mycustomers(cust_id, age, gender) VALUES (2088, 35, 'F');
INSERT INTO mycustomers(cust_id, age, gender) VALUES (1234, 54, 'M');
INSERT INTO mycustomers(cust_id, age, gender) VALUES (1354, 17, 'F');
INSERT INTO mycustomers(cust_id, age, gender) VALUES (6666, 15, 'F');
COMMIT;
```

This populates the MYCUSTOMERS table.

f. Create logical partition tracking for the MYCUSTOMERS table.

```
CREATE LOGICAL PARTITION TRACKING ON mycustomers
 PARTITION BY RANGE (age) INTERVAL (20.5)
 (PARTITION m0 values less than (20));
```

This creates a logical partition tracking for the MYSALES table using the AGE as key.

2. Create a materialized view on top of tables with logical partition change tracking.

a. Create materialized view on MYSALES and MYCUSTOMERS tables.

```
CREATE MATERIALIZED VIEW sales_age_time
 REFRESH FAST
 ENABLE QUERY REWRITE
 AS SELECT SUM(s.amount_sold) amount_total, c.age, s.time_id
 FROM mysales s, mycustomers c
 WHERE s.cust_id = c.cust_id
 GROUP BY c.age, s.time_id;
```

This creates the SALES_AGE_TIME materialized view.

b. Query the DBA_MVIEW_DETAIL_LOGICAL_PARTITION data dictionary view.

```
SELECT mview_name, DETAILOBJ_NAME, DETAIL_LOGICAL_PARTITION_NAME
LPARTNAME,
```

```
DETAIL_LOGICAL_PARTITION_NUMBER LPART#, FRESHNESS
FROM DBA_MVIEW_DETAIL_LOGICAL_PARTITION
WHERE mview_name = 'SALES_AGE_TIME'
ORDER BY 1,2,3;
```

It shows the following output.

```
MVIEW_NAME        DETAILOBJ_NAME     LPARTNAME      LPART#
FRESHNESS
---------------   ---------------    ----------     ----------
---------
SALES_AGE_TIME    MYCUSTOMERS        M0             0           FRESH
SALES_AGE_TIME    MYSALES            P0             0           FRESH
SALES_AGE_TIME    MYSALES            P1             1           FRESH
```

c. Use `EXPLAIN_MVIEW` to assess the logical partition related refresh and rewrite capabilities.

```
EXECUTE DBMS_MVIEW.EXPLAIN_MVIEW ('sales_age_time');

SELECT CAPABILITY_NAME, RELATED_TEXT, POSSIBLE
FROM MV_CAPABILITIES_TABLE
WHERE MVNAME = 'SALES_AGE_TIME' AND CAPABILITY_NAME LIKE '%LPT%'
ORDER BY 1, 2;
```

It shows the following output.

```
CAPABILITY_NAME              RELATED_TEXT                POSSIBLE
--------------------------   -------------------------
---------------
LPT                                                       Y
LPT_TABLE                    MYCUSTOMERS                  Y
LPT_TABLE                    MYSALES                      Y
LPT_TABLE_REWRITE            MYCUSTOMERS                  Y
LPT_TABLE_REWRITE            MYSALES                      Y
REWRITE_LPT                                               Y
REFRESH_FAST_LPT                                          Y
```

3. Observe the impact of DMLs on your materialized view.

a. Introduce a new logical partition on the `MYSALES` table.

```
INSERT INTO mysales (time_id, prod_id, cust_id, amount_sold) VALUES
      (TO_DATE('2019-02-05','yyyy-mm-dd'), 99, 2108, 33);
```

This introduces a new partition (partition #6) on the `MYSALES` table.

b. Introduce a new logical partition on the `MYSALES` table.

```
INSERT INTO mysales (time_id, prod_id, cust_id, amount_sold) VALUES
(TO_DATE('2019-02-05','yyyy-mm-dd'), 99, 2108, 33);
```

This introduces a new partition (partition #6) on the `MYSALES` table.

c. Introduce a new logical partition on the `MYCUSTOMERS` table.

```
INSERT INTO mycustomers(cust_id, age, gender) VALUES (1399, 80, 'F');
```

This introduces a new partition (partition #3) on the `MYCUSTOMERS` table.

d. Introduce a new logical partition on the `MYSALES` table.

```
INSERT INTO mysales (time_id, prod_id, cust_id, amount_sold) VALUES
(TO_DATE('2019-02-09','yyyy-mm-dd'), 99, 1997, 79.9);
```

This introduces a new partition (partition #7) on the `MYSALES` table.

e. Introduce a new logical partition on the `MYSALES` table.

```
INSERT INTO mysales (time_id, prod_id, cust_id, amount_sold) VALUES
(TO_DATE('2010-02-09','yyyy-mm-dd'), 110, 1997, 108.98);
COMMIT;
```

This introduces a new partition (partition #2) on the `MYSALES` table.

f. Query the `DBA_MVIEW_DETAIL_LOGICAL_PARTITION` data dictionary view.

```
SELECT mview_name, DETAILOBJ_NAME, DETAIL_LOGICAL_PARTITION_NAME
LPARTNAME,
 DETAIL_LOGICAL_PARTITION_NUMBER LPART#, FRESHNESS
 FROM DBA_MVIEW_DETAIL_LOGICAL_PARTITION
 WHERE mview_name = 'SALES_AGE_TIME'
 ORDER BY 1,2,3;
```

Now, it shows the following output.

```
MVIEW_NAME        DETAILOBJ_NAME     LPARTNAME      LPART#
FRESHNESS
---------------   ---------------    ----------     ----------
---------
SALES_AGE_TIME    MYCUSTOMERS        M0             0             FRESH
SALES_AGE_TIME    MYCUSTOMERS        SYS_88904P3    3             STALE
SALES_AGE_TIME    MYSALES            P1             0             FRESH
SALES_AGE_TIME    MYSALES            P1             1             FRESH
SALES_AGE_TIME    MYSALES            SYS_88899P3    2             STALE
SALES_AGE_TIME    MYSALES            SYS_88899P7    6             STALE
```

g. Perform LPCT rewrite on subset of lpart #1 on the `MYSALES` and lpart #0 on the `MYCUSTOMERS` tables.

```
DELETE FROM rewrite_table;
DECLARE
 stmt varchar2(2000) := q'#select sum(s.amount_sold) amount_total,
 c.age, s.time_id
 FROM mysales s, mycustomers c
```

```
 WHERE s.cust_id = c.cust_id
 AND s.time_id < TO_DATE ('07-07-2007', 'MM-DD-YYYY')
 AND c.age < 18
 GROUP BY c.age, s.time_id#';
BEGIN
 dbms_mview.explain_rewrite (stmt,'sales_age_time');
END;
/
SELECT mv_name, sequence, pass, message FROM rewrite_table;
```

h. Query `REWRITE_TABLE` to verify the rewrites.

```
SELECT mv_name, sequence, pass, message FROM rewrite_table;
```

i. Execute the following query.

```
SELECT SUM(s.amount_sold) amount_total,
 c.age, s.time_id
 FROM mysales s, mycustomers c
 WHERE s.cust_id = c.cust_id
 AND s.time_id < TO_DATE ('07-07-2007', 'MM-DD-YYYY')
 AND c.age < 18
 GROUP BY c.age, s.time_id;
```

j. View the explain plan for the above query to verify the rewrites.

```
SELECT * FROM TABLE(dbms_xplan.display_cursor);


PLAN_TABLE_OUTPUT
_____________________________________________________
SQL_ID  ampuzk8tbp6df, child number 0
-------------------------------------
SELECT SUM(s.amount_sold) amount_total,
c.age, s.time_id
FROM mysales s, mycustomers c
WHERE s.cust_id = c.cust_id
AND s.time_id < TO_DATE ('07-07-2007', 'MM-DD-YYYY')
AND c.age < 18
GROUP BY c.age, s.time_id;

Plan hash
        value: 3902795718

-------------------------------------------------------------------------
-----------------------
        | Id  | Operation                    | Name           | Rows  |
Bytes | Cost
        (%CPU)| Time

|-------------------------------------------------------------------------
-----------------------
        |   0 | SELECT STATEMENT             |                |
          |       |     2 (100)|
        ||*  1 |  MAT_VIEW
```

```
        REWRITE ACCESS FULL| SALES_AGE_TIME |     1 |    35 |
          2   (0)| 00:00:01

|-------------------------------------------------------------------------
------------------------
        Predicate Information (identified by operation id):
        ---------------------------------------------------
        1 - filter(("SALES_AGE_TIME"."TIME_ID"<TO_DATE('2007-07-07
00:00:00', 'syyyy-mm-dd hh24:mi:ss') AND "SALES_AGE_TIME"."AGE"<18))
Note
-----
   - dynamic statistics used: dynamic sampling (level=2)
26 rows selected.
```

4. Leverage LPCT for incremental refresh.
   a. Execute the following code to perform LPCT refresh.

```
EXECUTE DBMS_MVIEW.REFRESH('SALES_AGE_TIME', 'L');
```

   b. Verify the refresh using the following query.

```
SELECT mview_name, DETAILOBJ_NAME, DETAIL_LOGICAL_PARTITION_NAME
LPARTNAME, DETAIL_LOGICAL_PARTITION_NUMBER LPART#, FRESHNESS
 FROM DBA_MVIEW_DETAIL_LOGICAL_PARTITION
 WHERE mview_name = 'SALES_AGE_TIME'
 ORDER BY 1,2,3;
```

Migrating MySQL and Third-Party Databases to Autonomous Database

Migration is the process of copying the schema objects and data from a source MySQL or a third-party (non-Oracle) database, such as Amazon Redshift, Microsoft SQL Server, Sybase Adaptive Server, IBM DB2 (UDB), or Teradata, to Autonomous Database.

Using Oracle SQL Developer Migration Wizard you can perform migration in an efficient and largely automated way.

- Migrating Amazon Redshift to Autonomous Database
  The SQL Developer Amazon Redshift Migration Assistant, available with SQL Developer 18.3 and later versions provides a framework for easy migration of Amazon Redshift environments on a per-schema basis.
- Migrating MySQL to Autonomous Database
  Migration is the process of copying the schema objects and data from a source MySQL to Autonomous Database. Using Oracle SQL Developer Migration Wizard you can perform migration in an efficient and largely automated way.
- Migrating Microsoft SQL Server or Sybase Adaptive Server to Autonomous Database
  Migration is the process of copying the schema objects and data from a source Microsoft SQL Server or Sybase Adaptive Server to Autonomous Database. Using Oracle SQL Developer Migration Wizard you can perform migration in an efficient and largely automated way.

- Migrating IBM DB2 (UDB) Database to Autonomous Database
  Migration is the process of copying the schema objects and data from a source IBM DB2 (UDB) database to Autonomous Database. Using Oracle SQL Developer Migration Wizard you can perform migration in an efficient and largely automated way.
- Migrating Teradata Database to Autonomous Database
  Migration is the process of copying the schema objects and data from a source Teradata database to Autonomous Database. Using Oracle SQL Developer Migration Wizard you can perform migration in an efficient and largely automated way.

Migrating Amazon Redshift to Autonomous Database

The SQL Developer Amazon Redshift Migration Assistant, available with SQL Developer 18.3 and later versions provides a framework for easy migration of Amazon Redshift environments on a per-schema basis.

This section describes the steps and the workflow for both an online migration of Amazon Redshift and for the generation of scripts for a scheduled, manual migration that you run at a later time.

- Autonomous Database Redshift Migration Overview
  Using SQL Developer you can migrate database files from Amazon Redshift to Autonomous Database.
- Connect to Amazon Redshift
  Using SQL Developer you can migrate database Schemas and Tables from Amazon Redshift to Autonomous Database. To extract metadata and data from Amazon Redshift for a migration, you need to connect to Amazon Redshift with SQL Developer.
- Connect to Autonomous Database
  Using SQL Developer you create a connection to Autonomous Database
- Start the Cloud Migration Wizard
  Invoke the Cloud Migration Wizard from the Tools menu of SQL Developer to initiate your Amazon Redshift migration to Autonomous Database.
- Review and Finish the Amazon Redshift Migration
  The summary shows a summary of the information that you have specified.
- Use Generated Amazon Redshift Migration Scripts
  When you choose to generate migration scripts a new subdirectory is created in the local directory specified in the migration Wizard. You can run these scripts in real time or use them for programmatic processing.
- Perform Post Migration Tasks
  After successful migration of your Redshift environment you should consider the following post-migration tasks:

Autonomous Database Redshift Migration Overview

Using SQL Developer you can migrate database files from Amazon Redshift to Autonomous Database.

- **Capture:** Captures Metadata schemas and tables from source database and stores in Migration Repository.
- **Convert:** Redshift Datatypes are mapped to Oracle Datatypes. Redshift Object names are converted to Oracle names based on Oracle Naming Convention. The Column Defaults that use Redshift functions are replaced by their Oracle equivalents.
- **Generate:** Generate schemas and DDLs based on the converted metadata.

- **Deploy:** Deploy the generated schemas and DDLs.
- **Copy Data:** Unload data from Redshift tables to Amazon Storage S3 then copy data from Amazon Storage to Autonomous Database tables(in schemas) that were Deployed earlier.

Connect to Amazon Redshift

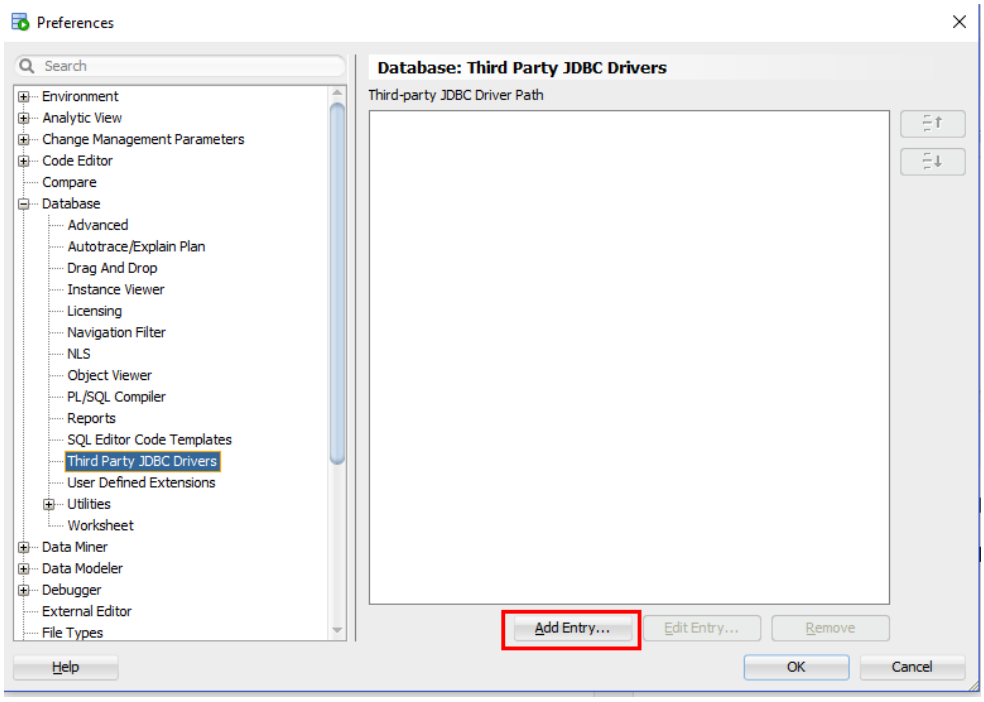
Using SQL Developer you can migrate database Schemas and Tables from Amazon Redshift to Autonomous Database. To extract metadata and data from Amazon Redshift for a migration, you need to connect to Amazon Redshift with SQL Developer.

Download Amazon Redshift JDBC Driver and Add the Third Party Driver

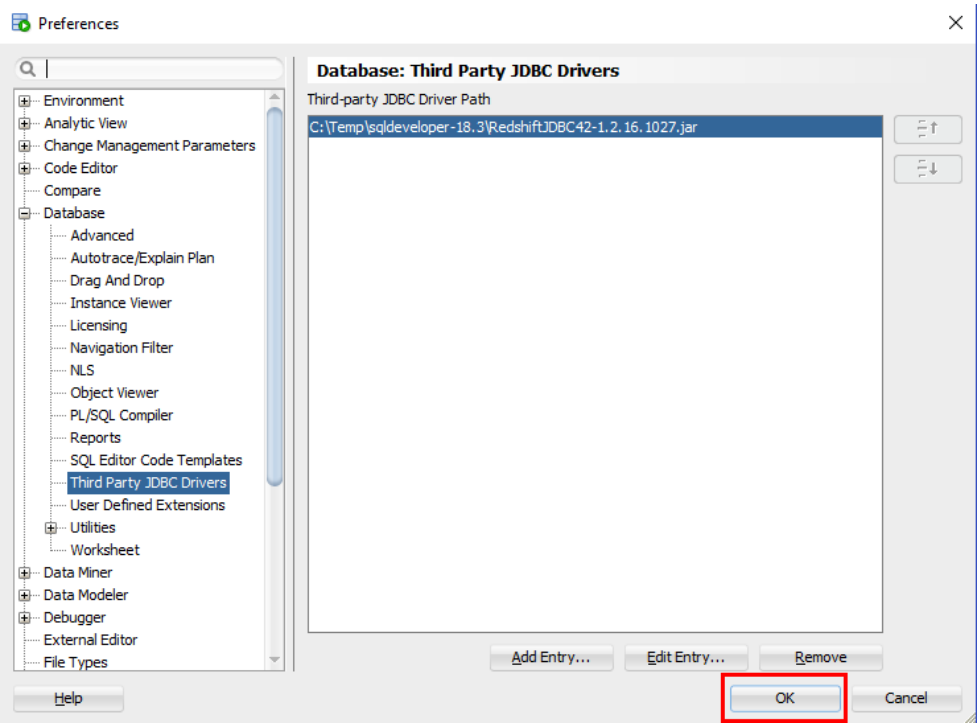
1. Download an Amazon Redshift JDBC driver to access Amazon Redshift. Consult the Amazon Redshift documentation for the location of the most recent JDBC driver. For more information, see Configure a JDBC Connection.

> **Note:**
>
> Use the Amazon Redshift JDBC Driver JDBC 4.2–compatible driver.

2. Store the Amazon Redshift JDBC driver in a local directory where SQL Developer can access the Amazon Redshift JDBC driver.
3. Add the Amazon Redshift JDBC driver as third party to SQL Developer before making a connection. Within SQL Developer, go to **Tools > Preferences > Database > Third Party JDBC Drivers** (for Mac, this is **Oracle SQL Developer > Preferences Database > Third Party JDBC Drivers**).
4. Click **Add Entry** and select the path to the Amazon Redshift JDBC Driver that you download.



5. Click **OK** to add the Amazon Redshift JDBC driver that you download.

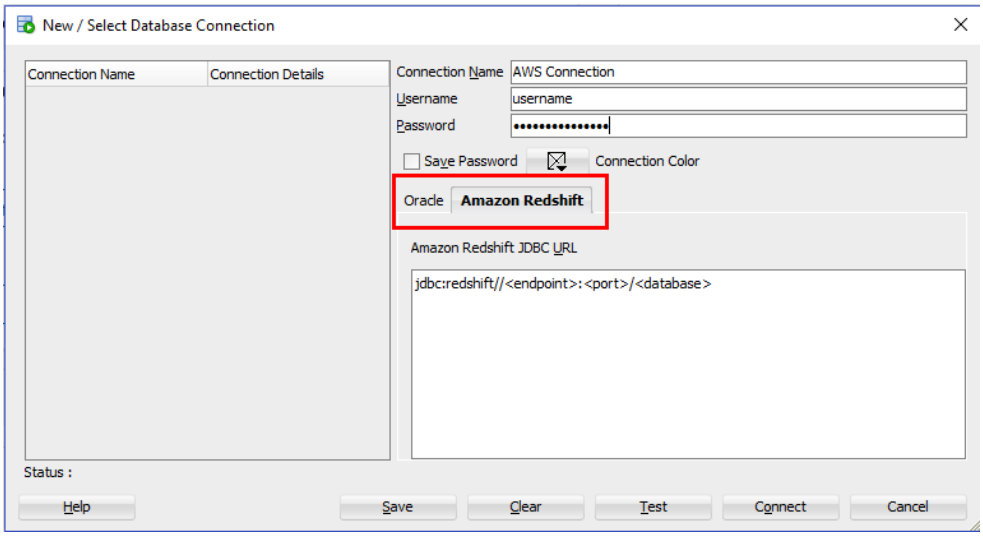


Add Connection to Amazon Redshift Database

Connect to the Amazon Redshift database.

1. In the Connections panel, right-click **Connections** and select **New Connection...**.
2. Select the **Amazon Redshift** tab and enter the connection information for Amazon Redshift.

   If you are planning to migrate multiple schemas it is recommended to connect with the Master username to your Amazon Redshift system.



- For more details for configuring a JDBC Connection and obtaining the Amazon Redshift JDBC URL, see AWS: Configure a JDBC Connection.
- For more details for configuring security options for the connection (in case of "Amazon [500150] connection error"), see AWS: Configure Security options for Connection (in case of "Amazon [500150] connection error").
- If you deployed your Amazon Redshift environment within a Virtual Private Cloud (VPC) you have to ensure that your cluster is accessible from the Internet. See http://docs.aws.amazon.com/redshift/latest/gsg/rs-gsg-authorize-cluster-access.html for details of how to enable public Internet access.
- If your Amazon Redshift client connection to the database appears to hang or times out when running long queries, see http://docs.aws.amazon.com/redshift/latest/mgmt/connecting-firewall-guidance.html for possible solutions to address this issue.

Test the connection before you save it.

Additional Information for Amazon Authentication and Access Control

- AWS: Security
- AWS: Managing Cluster Security Groups Using the Console

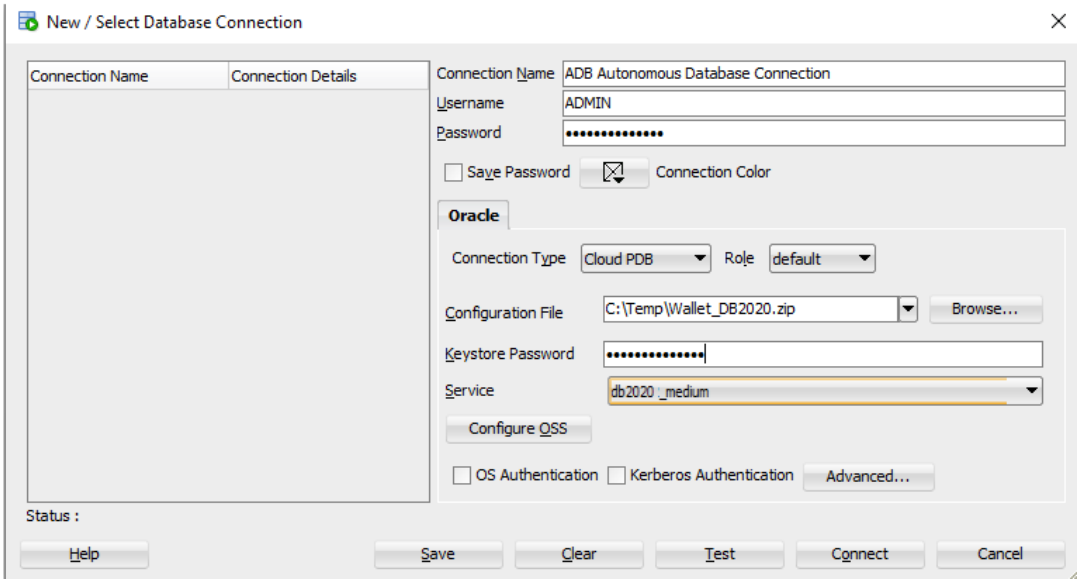
Connect to Autonomous Database

Using SQL Developer you create a connection to Autonomous Database

Obtain the client credentials wallet zip file. For more information, see Download Client Credentials (Wallets).

1. In the Connections panel, right-click **Connections** and select **New Connection...**.
2. Select the **Oracle** tab and enter the connection information for Autonomous Database.
3. For the AWS Redshift Migration connection, select the **_low** connection to your database.

   For more information, see See Database Service Names for Autonomous Database for more information..
4. Add a connection to Autonomous Database.



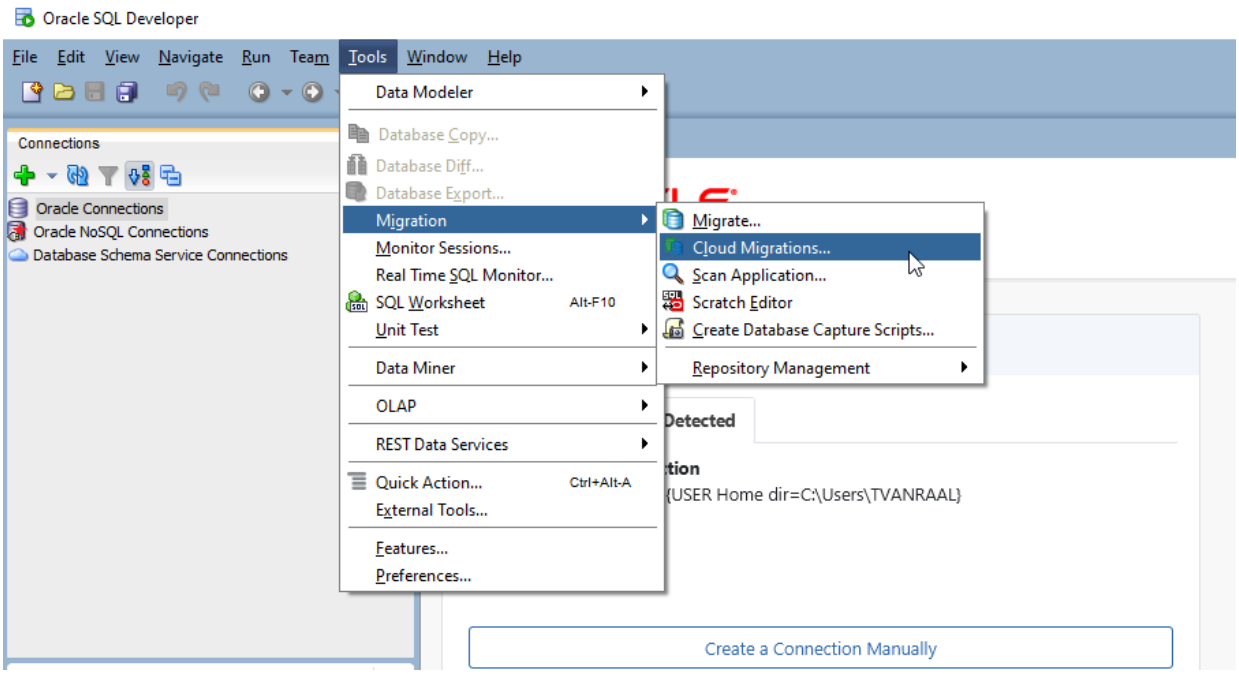
See Connect Oracle SQL Developer with a Wallet (mTLS) for more information.

Test the connection before you save it.

Start the Cloud Migration Wizard

Invoke the Cloud Migration Wizard from the Tools menu of SQL Developer to initiate your Amazon Redshift migration to Autonomous Database.

The Cloud Migration wizard starts when you click Cloud Migrations from **Migration** in the **Tools** menu. The Cloud Migrations wizard enables you to migrate schemas, objects (tables), and data from an Amazon Redshift database to Autonomous Database.



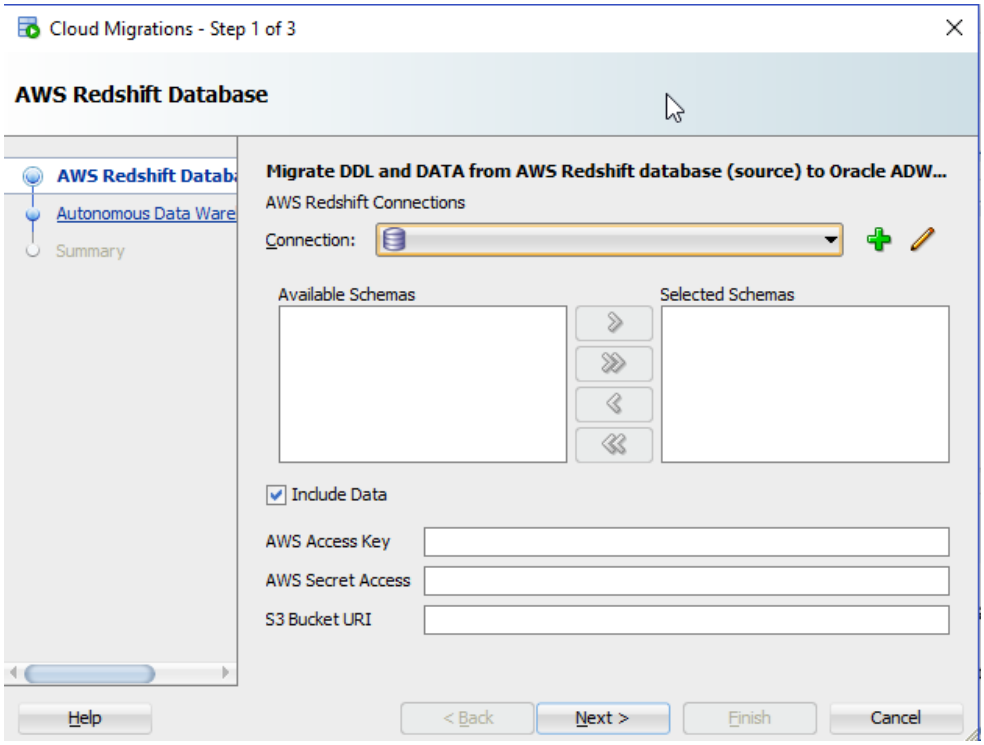
The Cloud Migration Wizard is an easy set of steps. The Cloud Migration Wizard guides you to:

- Identify the schemas in your Amazon Redshift database that you want to migrate.
- Identify the target Autonomous Database.
- Define whether you want to migrate the metadata (DDL), the data, or both.
- Choose to migrate your system online, to generate scripts, or both.

Identify the Amazon Redshift Database

Identify the schemas in the Amazon Redshift database to migrate. All objects, mainly tables, in the schema will be migrated. Migration to Autonomous Database is on a per-schema basis. Schemas cannot be renamed as part of the migration.

1. In the AWS Redshift Migration, specify the Connection.



- **Connection:** Name of the Redshift database connection.
- **Available Schemas:** Schemas available for the specific connection.
- **Selected Schemas:** Click the Add icon to select the schemas you want to migrate from **Available Schemas**.
- **Include Data:** DDL and DATA migrates the selected schemas and data.

When you migrate data, you have to provide the AWS access key, AWS Secret Access Key, and an existing S3 bucket URI where the Redshift data will be unloaded and staged. The security credentials require certain privileges to store data in S3. It is recommended to create new, separate access keys for the migration. The same access key is used later on to load the data into the Autonomous Database using secure REST requests.

- **AWS Access Key:** For more information on access keys, see AWS Identity and Access Management.

- **AWS Secret Access:** For more information on access keys, see AWS Identity and Access Management.
- **S3 Bucket URI:** For information on common S3 ServiceException errors, see S3ServiceException Errors.

  For more information an S3 buckets, see Creating and Configuring an S3 Bucket.

Amazon S3 Bucket URI Format

For the source files that reside in Amazon S3, see the following for a description of the URI format for accessing your files: Accessing a Bucket For example the following refers to the file folder 'folder_name' in the adwc bucket in the us-west-2 region.

```
https://s3-us-west-2.amazonaws.com/adwc/folder_name
```

S3 Bucket Configuration Example 1

If you provide the following S3 Bucket URI :

```
https://s3-us-west-2.amazonaws.com/my_bucket
```

The wizard verifies the entire path including *my_bucket*. An attempt is made to write a test file, if it is not accessible there is a prompt. In case, `my_bucket` does not exist, there is an error reported:

```
Validation Failed
```

Then the code generation creates the following path for the `DBMS_CLOUD.COPY_DATA` function:

```
file_uri_list => "https://s3-us-west-2.amazonaws.com/my_bucket/
oracle_schema_name/oracle_table_name/*.gz"
```

The migration assistant creates these folders: *oracle_schema_name*/*oracle_table_name* inside the bucket: *my_bucket*.

S3 Bucket Example 2

If you provide the following S3 Bucket URI :

```
https://s3-us-west-2.amazonaws.com/my_bucket/another_folder
```

The wizard verifies the entire path including *my_bucket*. An attempt is made to write a test file, if it is not accessible there is a prompt. In case, `my_bucket` does not exist, there is an error reported:

```
Validation Failed
```

In this case the *another_folder* does not have to exist. The migration creates the *another_folder* bucket inside *my_bucket*.

Then the code generation creates the following path for the `DBMS_CLOUD.COPY_DATA` function:

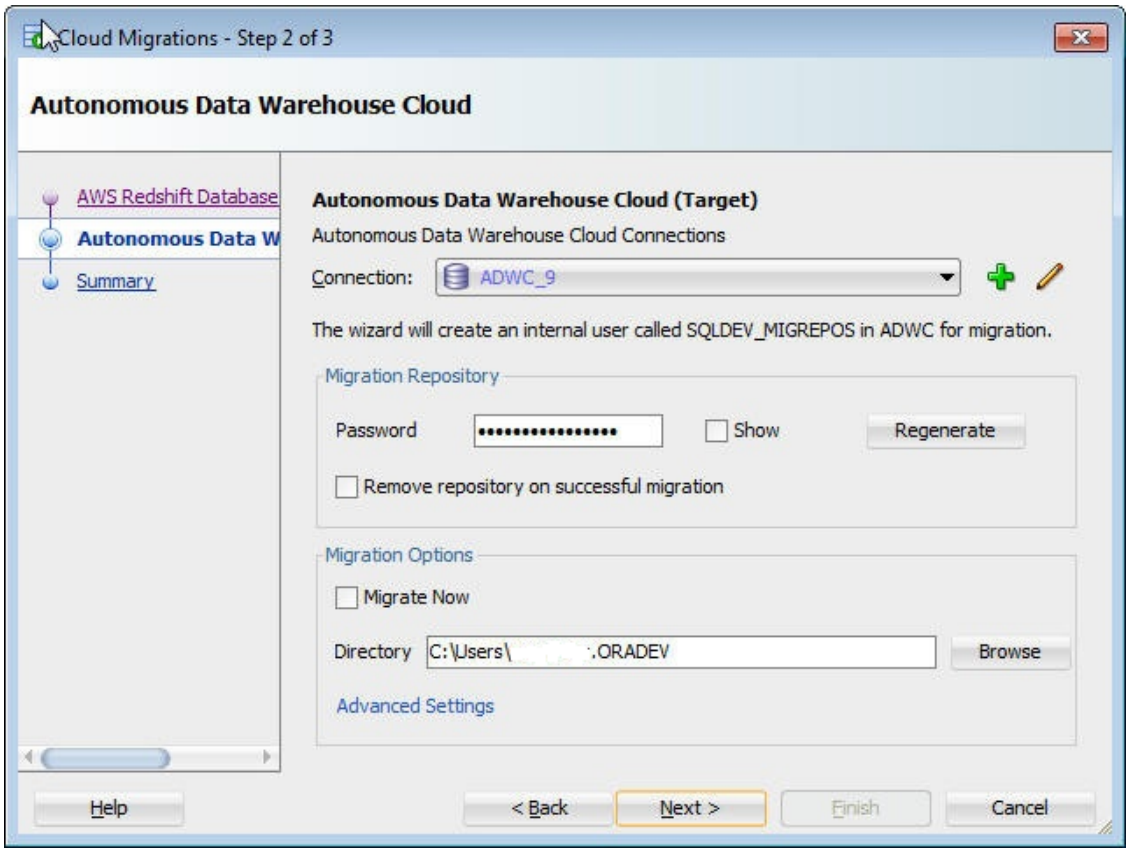
```
file_uri_list => 'https://s3-us-west-2.amazonaws.com/my_bucket/another_folder/
oracle_schema_name/oracle_table_name/*.gz
```

Step 2 of 3: Autonomous Data Warehouse Cloud

First create a connection for your target Autonomous Database See Connect to Autonomous Database. The user for this connection must have the administrative privileges; the connection is used throughout the migration to create schemas and objects.

> **Note:**
>
> Use the ADMIN user or a user with admin role.

The Amazon Redshift Migration Assistant allows you to do an online migration right away, to generate all scripts necessary for a migration, or both. If you chose to store the scripts in a local directory you have to specify the local directory (the directory must be writable by the user).



- **Connection:** Name of the Autonomous Data Warehouse Cloud connection. Create a connection for the Autonomous Database if required. The user must have administrative privileges since this connection is used throughout the migration to create schemas and objects.

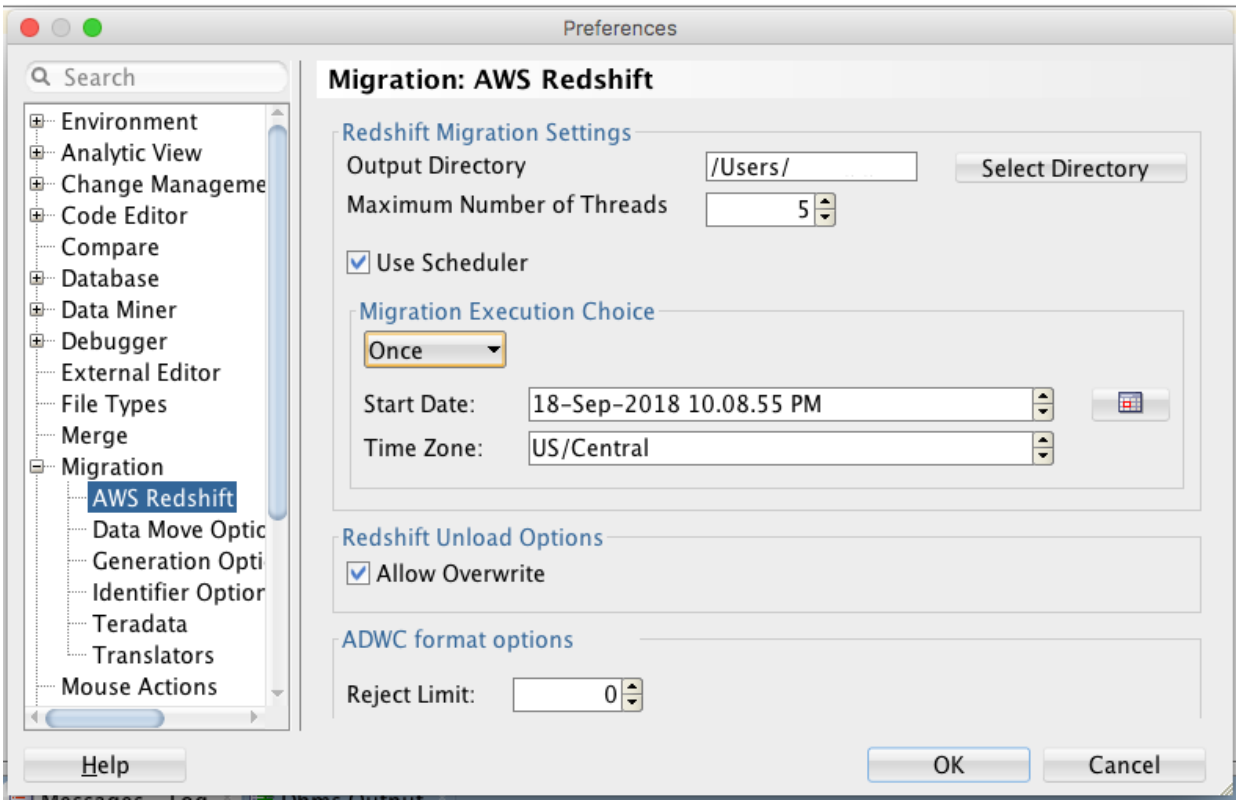
- **Migration Repository Password:** Password for the migration repository that is installed in the Autonomous Database as part of the schema migration. Either use the pre-filled password or enter a new password.
- **Remove repository on successful migration:** Select this option to remove the repository after the migration is completed. The repository is not required after migration.
- **Migrate Now:** Select this option to perform an online migration immediately.

> **Note:**
>
> – If **Include Data** from Step 1 and **Migrate Now** are both unselected, you are opting for just generation of all required SQL Scripts for manual migration.
>
> – If **Include Data** from Step 1 is unchecked and **Migrate Now** is selected, then all selected schemas and their tables will be deployed in Autonomous Database but data will not be loaded into tables.
>
> – If **Include Data** from Step 1 and **Migrate Now** are both selected, then all selected schemas and their tables will be deployed in Autonomous Database and data will be loaded into tables.

- **Directory:** Specify the director to store the generated scripts necessary for the migration; this saves the scripts in a local directory.

Advanced Settings (Optional)

The default settings should work unless you want to control the format options when Unloading to S3 storage or Copying from S3 storage to Autonomous Database. For more information on Format Options, see DBMS_CLOUD Package Format Options. To use advanced options, click **Advanced Settings**.



Output Directory: Enter the path or click Select Directory to select the directory or folder for the migration.

Maximum Number of Threads: Enter the number of parallel threads to enable when loading data to tables in Autonomous Database.

Use Scheduler: Select this option to enable the scheduler for migration. You can schedule jobs for data load migration operations from the AWS S3 bucket to Autonomous Database. You have the option to run the scheduled jobs immediately or at a future date and time. To monitor the data load scheduled jobs, use the Scheduler node in the Connections navigator.

Migration Execution Choice:

- **Immediate** runs the scheduler as soon as the Redshift migration is triggered.
- **Once** runs the scheduler on a future date. You specify the **Start Date** and **Time Zone**. By default, the Start Date displays the current date and time of the local system. To change the start date, use the calendar icon to double-click and select the date or use the spinner to highlight the date and then click the field to set it.

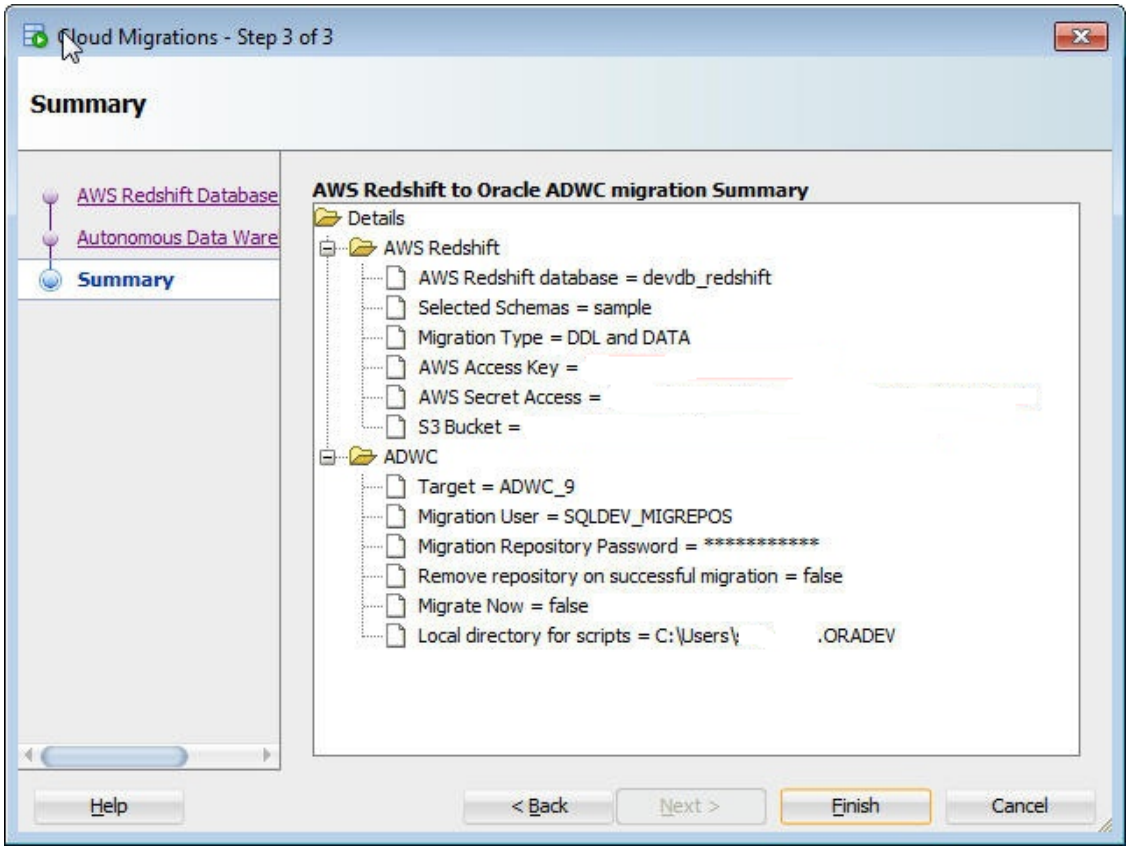
Redshift Unload Options: Allow Overwrite: If this option is enabled, the unload process will overwrite existing files, including the manifest file (lists the data files that are created by the unload process). By default, unload fails if there are files that can be overwritten.

ADWC format options: Reject Limit: Enter the number of rows to reject when loading data to tables in Autonomous Database. The migration operation will error out after the specified number of rows are rejected. The default is 0.

Review and Finish the Amazon Redshift Migration

The summary shows a summary of the information that you have specified.

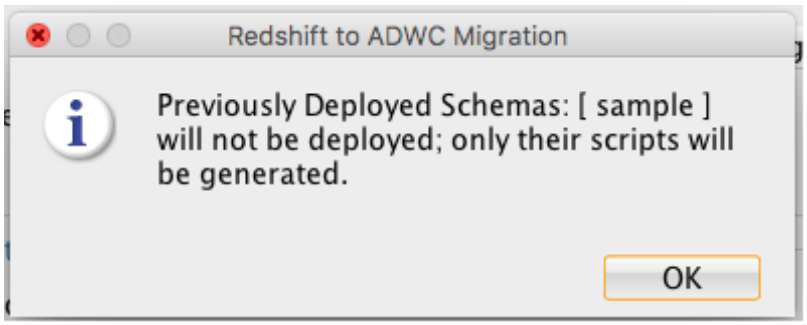
To change any information, press **Back** as needed.



If you have chosen an immediate migration, then the dialog of the migration wizard stays open until the migration is finished. If you select generate scripts, the migration process generates the necessary scripts in the specified local directory, and does not run the scripts.

To perform the migration, click **Finish**

If the selected schema name in AWS Redshift already exists in Autonomous Database, the migration process excludes deploying these selected schemas and displays a dialog:



Summary: What The Migration Assistant Creates

- Creates a new Autonomous Database user using the *schema_name* from Redshift.
- Creates a new bucket on S3 based on the *schema name*.
- Creates sub-folders on S3 for each table.

Use Generated Amazon Redshift Migration Scripts

When you choose to generate migration scripts a new subdirectory is created in the local directory specified in the migration Wizard. You can run these scripts in real time or use them for programmatic processing.

The directory contains the following scripts:

- `redshift_s3unload.sql`
- `adwc_ddl.sql`
- `adwc_dataload.sql`
- `adwc_dataload_scheduler.sql`

These scripts contain all necessary commands to migrate your Amazon Redshift system to Autonomous Database. You can run these scripts in real time or use them for programmatic processing.

Unload Your Amazon Redshift Data into S3

The first step of a successful migration is to unload your Amazon Redshift data into Amazon S3, which acts as a staging area. Script **redshift_s3unload.sql** has all the Amazon Redshift unload commands to unload the data using the access credentials and the S3 bucket that were specified in the Migration Wizard workflow.

Connect to your Amazon Redshift environment to run this script.

Create Your Data Warehouse Objects

To prepare your Autonomous Database create your empty data warehouse schema prior to loading data. The Amazon Redshift Migration Assistant converted all Amazon Redshift schema structures into Oracle structures in script **adwc_ddl.sql**.

The script must be executed while you are connected to your Autonomous Database as privileged user; for example, ADMIN.

By default, the schema created for the migration has the same name as the schema in Amazon Redshift. You must change the password to the valid password for the specified user either in the script or after the script runs. If you want to change the schema name then change the schema name and all references to the name.

Load Your Amazon Redshift Data into Your Oracle Autonomous Database

The script **adwc_dataload.sql** contains all the load commands necessary to load your unloaded Amazon Redshift data straight from S3 into your Autonomous Database.

Execute the script while connected to your Autonomous Database as a privileged user; for example ADMIN.

If you want to change the target schema name when you create your data warehouse objects then you must adjust the target schema names in this script accordingly.

Use of JOB SCHEDULER

SQL Developer provides a graphical interface for using the DBMS_SCHEDULER PL/SQL package to work with Oracle Scheduler objects. To use the SQL Developer scheduling features, please refer 'Scheduling Jobs Using SQL Developer' topic of SQL Developer User Guide and **Oracle Database Administrator's Guide** to understand the concepts and essential tasks for job scheduling.

The Scheduler node for a connection appears in the Connections navigator and in the DBA navigator. Use ADWC 'admin' user to navigate which displays Scheduler objects owned by the 'admin' monitoring status of data load jobs.

Under ADWC 'admin' Connection → Scheduler → Jobs, you will see AWS Redshift to ADWC data load jobs are created with name <schema_name>_<table_name>.

To see the status of completion of each data load, please expand each scheduled job and check the status.

Also for more detailed information about data load operation see table `MD_REPORT` in `SQLDEV_MIGREPOS` schema that stores information about table columns: and

```
 OPERATION_ID, LOGFILE_TABLE, BADFILE_TABLE, SOURCE_SCHEMA_NAME,
TARGET_SCHEMA_NAME, SOURCE_TABLE_NAME,
```

and

```
 TARGET_TABLE_NAME, SOURCE_TABLE_ROWS, TARGET_TABLE_ROWS_LOADED, ERROR
MESSAGE,
```

and

```
STATUS (COMPLETED or FAILED)
```

Redshift Migration Log and Report Files

After Redshift Migration, you will find three files:

- `MigrationResults.log` : Log file of Redshift migration
- `readme.txt` : file explains how to use the Generated Amazon Redshift Migration Scripts.
- `redshift_migration_reportxxx.txt` : Contains information about Migration, here is sample:

```
_____________________________________________________________
OPERATION ID : 8566
LOGFILE TABLE : COPY$8566_LOG
BADFILE TABLE : COPY$8566_BAD
SOURCE SCHEMA : sample
TARGET SCHEMA : SAMPLE
SOURCE TABLE : listing
TARGET TABLE : LISTING
SOURCE TABLE ROWS : 192497
TABLE ROWS LOADED : 192497
ERROR MESSAGE : null
STATUS : COMPLETED
```

```
START TIME : 2018-09-27 17:25:18.662075
END TIME : 2018-09-27 17:25:25.012695
_____________________________________________________________
```

Perform Post Migration Tasks

After successful migration of your Redshift environment you should consider the following post-migration tasks:

- Drop schema `SQLDEV_MIGREPOS`
- Drop the Amazon S3 bucket used for staging
- Harden the Amazon account used for accessing S3
- Drop the database credential used for data loading from S3
- Harden your accounts in your Autonomous Database

1. Drop schema `SQLDEV_MIGREPOS`

   As part of the schema migration the Migration Assistant installs a minimal migration repository in the target Autonomous Database. After the migration this account is no longer needed and can be dropped or alternatively locked.

2. Drop the Amazon S3 Bucket Used for Staging

   Unless you desire to use the unloaded Redshift data otherwise you can drop the bucket containing the unloaded data.

3. Harden the Amazon Account Used for Accessing S3

   You should inactivate the security access key used for S3 access unless needed for other purposes.

4. Drop the database credential used for data loading from S3

   The Amazon security credentials to access S3 are stored encrypted as database credential `REDSHIFT_DWCS_CREDS` in your Autonomous Database in the privileged user schema that was used for the migration. Oracle recommends you drop this credential after successful migration unless needed for other purposes. For more information, see DROP_CREDENTIAL Procedure.

5. Harden your Accounts in Your Autonomous Database

   For the new schema created as part of the migration with the Migration Assistant, ensure to change the passwords of these accounts or lock and expire them if they’re solely used for data storage.

Migrating MySQL to Autonomous Database

Migration is the process of copying the schema objects and data from a source MySQL to Autonomous Database. Using Oracle SQL Developer Migration Wizard you can perform migration in an efficient and largely automated way.

See SQL Developer: Migrating Third-Party Databases and Before Migrating From MySQL for more information.

Migrating Microsoft SQL Server or Sybase Adaptive Server to Autonomous Database

Migration is the process of copying the schema objects and data from a source Microsoft SQL Server or Sybase Adaptive Server to Autonomous Database. Using Oracle SQL Developer Migration Wizard you can perform migration in an efficient and largely automated way.

See SQL Developer: Migrating Third-Party Databases and Before Migrating From Microsoft SQL Server or Sybase Adaptive Server for more information.

Migrating IBM DB2 (UDB) Database to Autonomous Database

Migration is the process of copying the schema objects and data from a source IBM DB2 (UDB) database to Autonomous Database. Using Oracle SQL Developer Migration Wizard you can perform migration in an efficient and largely automated way.

See SQL Developer: Migrating Third-Party Databases and Before Migrating From IBM DB2 for more information.

Migrating Teradata Database to Autonomous Database

Migration is the process of copying the schema objects and data from a source Teradata database to Autonomous Database. Using Oracle SQL Developer Migration Wizard you can perform migration in an efficient and largely automated way.

See SQL Developer: Migrating Third-Party Databases and Before Migrating From Teradata for more information.

SODA Collection Metadata

It describes SODA Collection metadata on the database.

- SODA Collection Metadata on Autonomous Database
  Describes default and customized collection metadata on Autonomous Database.

SODA Collection Metadata on Autonomous Database

Describes default and customized collection metadata on Autonomous Database.

- SODA Default Collection Metadata on Autonomous Database
  Describes the default collection metadata on Autonomous Database, that is the metadata for a collection that is added when custom metadata is not supplied.
- SODA Customized Collection Metadata on Autonomous Database
  Describes SODA collection custom metadata on Autonomous Database.

SODA Default Collection Metadata on Autonomous Database

Describes the default collection metadata on Autonomous Database, that is the metadata for a collection that is added when custom metadata is not supplied.

Each SODA implementation provides a way to create a default collection when you supply a collection name. For example, in SODA for Java you use the `createCollection` method and supply just a collection name parameter:

```
db.admin().createCollection("myCol");
```

This creates a collection with default collection metadata. When you create a default collection on your database, the collection metadata includes the following information (regardless of which SODA implementation you use to create the default collection):

```
{
   "keyColumn" :
   {
      "name" : "ID",
      "sqlType" : "VARCHAR2",
      "maxLength" : 255,
      "assignmentMethod" : "UUID"
   },

   "contentColumn" :
   {
      "name" : "JSON_DOCUMENT",
      "sqlType" : "BLOB",
      "jsonFormat" : "OSON"
   },
   "versionColumn" :
   {
     "name" : "VERSION",
     "method" : "UUID"
   },

   "lastModifiedColumn" :
   {
     "name" : "LAST_MODIFIED"
   },

   "creationTimeColumn" :
   {
      "name" : "CREATED_ON"
   },

   "readOnly" : false
}
```

> **Note:**
>
> Using Always Free Autonomous Database with Oracle Database 23ai, the default metadata changes as follows.

```
{
   "keyColumn" :
   {
      "name" : "ID",
      "sqlType" : "VARCHAR2",
      "maxLength" : 255,
      "assignmentMethod" : "UUID"
   },

   "contentColumn" :
   {
      "name" : "JSON_DOCUMENT",
      "sqlType" : "JSON",
   },
   "versionColumn" :
   {
     "name" : "VERSION",
     "method" : "UUID"
   },

   "lastModifiedColumn" :
   {
     "name" : "LAST_MODIFIED"
   },

   "creationTimeColumn" :
   {
      "name" : "CREATED_ON"
   },

   "readOnly" : false
}
```

SODA Customized Collection Metadata on Autonomous Database

Describes SODA collection custom metadata on Autonomous Database.

Each SODA implementation provides a way to customize the collection metadata during collection creation. For example, in SODA for Java, you can use the following command:

```
OracleDocument metadata = db.createDocumentFromString("metadata_string");
OracleCollection col = db.admin().createCollection("myCustomColl", metadata);
```

In this example, for *`metadata_string`* you can use the default metadata as the starting point, and customize the following:

- Change `keyColumn.assignmentMethod` to `CLIENT`: Change the value of the `assignmentMethod` under `keyColumn` in the metadata to `CLIENT` (instead of `UUID`).

Valid values for `keyColumn.assignmentMethod` on Autonomous Database:

  - **UUID** (default): Keys are generated by SODA, based on the `UUID`.
  - **CLIENT**: Keys are assigned by the client application.

- Provide a `mediaTypeColumn name` value: A media type column is needed if the collection is to be heterogeneous, that is, it can store documents other than JavaScript Object Notation (JSON). See Media Type Column Name for details.

The following example specifies client-assigned keys and a custom media type column. The `mediaTypeColumn name` is specified with the value *YOUR_MEDIA_TYPE_COLUMN_NAME*. Otherwise, the default settings are used.

```
{
   "keyColumn" :
   {
      "name" : "ID",
      "sqlType" : "VARCHAR2",
      "maxLength" : 255,
      "assignmentMethod" : "CLIENT"
   },

   "contentColumn" :
   {
      "name" : "JSON_DOCUMENT",
      "sqlType" : "BLOB"
   },

   "versionColumn" :
   {
     "name" : "VERSION",
     "method" : "UUID"
   },

   "lastModifiedColumn" :
   {
     "name" : "LAST_MODIFIED"
   },

   "creationTimeColumn" :
   {
      "name" : "CREATED_ON"
   },

   "mediaTypeColumn" :
   {
      "name" : "YOUR_MEDIA_TYPE_COLUMN_NAME"
   },

   "readOnly" : false
}
```

Cloud Support and Identity Information

This lists the procedure to file a service request using Oracle Cloud Support.

- Obtain Tenancy Details
  When you file a service request or when you need Autonomous Database tenancy details for other reasons you can obtain these details for your instance. Tenancy details for an instance are available on the Oracle Cloud Infrastructure Console or you can obtain these details by querying the database.

Obtain Tenancy Details

When you file a service request or when you need Autonomous Database tenancy details for other reasons you can obtain these details for your instance. Tenancy details for an instance are available on the Oracle Cloud Infrastructure Console or you can obtain these details by querying the database.

If you need to file a service request use Oracle Cloud Support or contact your support representative. When you contact support, provide the first five items so that support can identify your instance.

- Database Name
- Region
- Tenancy OCID
- Database OCID
- Compartment OCID
- Outbound IP Address
- Availability Domain

If you are connected to the database you can obtain tenancy details by querying the `CLOUD_IDENTITY` column of the `V$PDBS` view.

For example, running the following provides the cloud identity information for an instance:

```
SELECT cloud_identity FROM v$pdbs;

CLOUD_IDENTITY
--------------------------------------------------------------------
{
  "DATABASE_NAME" : "DBxxxxxxxxxxx",
  "REGION" : "us-phoenix-1",
  "TENANT_OCID" : "OCID1.TENANCY.REGION1..ID1",
  "DATABASE_OCID" : "OCID1.AUTONOMOUSDATABASE.OC1.IAD.ID2",
  "COMPARTMENT_OCID" : "ocid1.tenancy.region1..ID3"
  "PUBLIC_DOMAIN_NAME" : "adb.us-ashburn-1.oraclecloud.com",
  "TENANT_ACCOUNT_NAME" : "test_example",
  "AUTOSCALABLE_STORAGE" : false,
  "BASE_SIZE" : 1099511627776,
  "INFRASTRUCTURE" : "Serverless",
  "SERVICE" : "ADW",
  "APPLICATIONS" :
  [
    "GRAPH",
    "ODI",
    "ORDS",
    "DATABASEACTIONS",
    "OMLMOD",
    "APEX",
```

```
    "OML"
  ],
  "OUTBOUND_IP_ADDRESS" :
  [
       "192.0.2.254"
  ]
  "AVAILABILITY_DOMAIN" : "SoSC:region-1-AD-1"
  "COMPUTE_MODEL" : "ECPU",
  "COMPUTE_COUNT" : 2,
  "COMPUTE_AUTOSCALING" : true,
  "AUTONOMOUS_DATA_GUARD" :
  [
    {
      "STANDBY_TYPE" : "local",
      "AVAILABILITY_DOMAIN" : "SoSC:region-AD-2"
    },
    {
      "STANDBY_TYPE" : "cross-region",
      "AVAILABILITY_DOMAIN" : "SoSC:region-AD-1"
    }
  ]
}
```

If your Autonomous Database instance was created before the tenancy details feature was added and has not been restarted, then this query does not return tenancy details. In this case, as a one-time operation, restart your instance and run the query again. You can restart your instance from the Oracle Cloud Infrastructure console or using the restart API.

Topics

- Database Name
- Region
- Tenancy OCID
- Database OCID
- Compartment OCID
- Outbound IP Address
- Public Domain Name
- Tenant Account Name
- Autoscalable Storage
- Base Size
- Infrastructure
- Service
- Applications
- Availability Domain
- Compute Model
- Compute Count
- Compute Auto Scaling

- Cloud Identity Information for Autonomous Data Guard
- Public IP Address Ranges for Autonomous Database

Database Name

You set the database name when you provision a database or when you rename a database.

The Oracle Cloud Infrastructure Console shows the database name on the Autonomous Database Details page under General Information in the **Database Name** field.

Region

The Oracle Cloud Infrastructure Console shows the region on the console, in the Regions area.

Tenancy OCID

The tenancy details page shows the tenancy OCID.

See Managing the Tenancy for information on accessing the Tenancy Details page.

See Resource Identifiers for information on Oracle Cloud Identifiers.

Database OCID

The Oracle Cloud Infrastructure Console shows the database OCID on the Autonomous Database Details page under General Information in the **OCID** field.

See Resource Identifiers for information on Oracle Cloud Identifiers.

Compartment OCID

See Managing Compartments for more information on compartments.

See Resource Identifiers for information on Oracle Cloud Identifiers.

Outbound IP Address

The `OUTBOUND_IP_ADDRESS` is the outbound IP address of your Autonomous Database instance. You can use the `OUTBOUND_IP_ADDRESS` when you create a database link to another Autonomous Database instance that uses ACLs to restrict access. In this case, you need to allow the specified outbound IP address to connect to the target Autonomous Database and then create the database link.

The `OUTBOUND_IP_ADDRESS` shows the outbound IP address in the following cases:

- When your Autonomous Database instance uses a Public Endpoint.
- When your Autonomous Database instance uses a Private Endpoint and the database property `ROUTE_OUTBOUND_CONNECTIONS` is set to `''` (the default value).

> **Note:**
>
> When your Autonomous Database instance uses a Private Endpoint and you set the database property `ROUTE_OUTBOUND_CONNECTIONS` to `'PRIVATE_ENDPOINT'`, outbound connections go through the private endpoint. See Enhanced Security for Outbound Connections with Private Endpoints for more information.

See Create Database Links from Autonomous Database to a Publicly Accessible Autonomous Database with a Wallet (mTLS) for information on creating database links.

See Configuring Network Access with Access Control Rules (ACLs) for information on configuring access control rules.

Public Domain Name

The `PUBLIC_DOMAIN_NAME` shows the public domain name for your Autonomous Database instance.

Tenant Account Name

The `TENANT_ACCOUNT_NAME` field specifies the tenant account name for your Autonomous Database instance.

See Tenancy Management for more information.

Autoscalable Storage

The `AUTOSCALABLE_STORAGE` field specifies if storage auto scaling is enabled. Possible values are `true` or `false`.

See Storage Auto Scaling for more information.

Base Size

The `BASE_SIZE` show's the Autonomous Database instance base storage size.

See Provision an Autonomous Database Instance for more information.

Infrastructure

The `INFRASTRUCTURE` field shows the deployment type.

See Provision an Autonomous Database Instance for more information.

Service

The `SERVICE` field shows the Autonomous Database instance workload type.

See About Autonomous Database Workload Types and Provision an Autonomous Database Instance for more information.

Applications

The `APPLICATIONS` field shows the tools that are enabled on the Autonomous Database instance.

See Manage Autonomous Database Built-in Tools for more information.

Availability Domain

You can view the Availability Domain of your Autonomous Database instance.

This is useful when you need to deploy your Application Servers when your Autonomous Database instance resides in a multi-Availability Domain region. In order to reduce latency as much as possible you would want to deploy your Application Server and database in the same Availability Domain.

See Regions and Availability Domains for more information on Availability Domains.

Compute Model

The `COMPUTE_MODEL` shows the Autonomous Database instance compute model. Valid values are `ECPU` or `OCPU`.

See Compute Models in Autonomous Database for more information.

Compute Count

The `COMPUTE_COUNT` shows the Autonomous Database instance **ECPU count** (**OCPU count** if your database uses OCPUs).

See Provision an Autonomous Database Instance for more information.

Compute Auto Scaling

The `COMPUTE_AUTOSCALING` field specifies if the Autonomous Database instance has compute auto scaling enabled. Valid values are `true` or `false`.

See Compute Auto Scaling for more information.

Cloud Identity Information for Autonomous Data Guard

The `AUTONOMOUS_DATA_GUARD` field shows the available Autonomous Data Guard disaster recovery standby databases, information for each standby includes the `STANDBY_TYPE`, either `local` or `cross-region`, and the `AVAILABILITY_DOMAIN`.

See Use Standby Databases with Autonomous Data Guard for Disaster Recovery for more information.

Public IP Address Ranges for Autonomous Database

Oracle Cloud Infrastructure provides information about public IP address ranges for Autonomous Databases that are deployed in Oracle Cloud Infrastructure.

See IP Address Ranges for more information.