Oracle Base Database Service

Oracle Base Database Service

Release Latest Cloud Release

April 2026

Oracle Base Database Service Oracle Base Database Service, Release Latest Cloud Release

Copyright © 2022, 2026, Oracle and/or its affiliates.

Primary Author: suresh.m.mohan@oracle.com

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.

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.

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.

Contents

1 Overview

What's New in Oracle Base Database Service 1
About Oracle Base Database Service 10
- Supported Database Editions 10
- About Easy Create 11
- About Oracle Base Database Service for Developers 11
- Supported Database Versions 12
- About Oracle AI Database 26ai 12
- Upgrade the DB System 12
- Update the Operating System of a DB System 12
- Upgrade the Database in a DB System 13
- Update the Database in a DB System 13
- Oracle Database Preview Version Availability 13
- Oracle Database Preview Version Restrictions 13
- Per-Second Billing for DB Systems 14
- Customer Managed Keys for Databases 14
- Backup and Recovery 14
- Move Databases to Oracle Cloud Using Zero Downtime Migration 14
- Support Requirements 15

About DB Systems 15
- Available Shapes and How It Determines the Resources Allocated 16
- Available Database Versions 19
- How various configurations affect the usable storage 19
- Service Limits 21
- Fast Provisioning Option 21
- Storage Scaling Considerations While Using Fast Provisioning 22
- Fault Domain Considerations for Multi-Node RAC DB Systems 22
- Reboot a DB System Node for Planned Maintenance 22
- Security Hardening Tool for DB systems 23
- Boot Volume Backups 23

2 Configure the Network

VCN and Subnets 1

Copyright © 2022, 2026, Oracle and/or its affiliates.

Overview 1
IPv4/IPv6 Dual Stack Network Support 2
Private Subnet with Service Gateway 2
Public Subnet with Internet Gateway 4
Requirements for IP Address Space 6
VCN Creation Wizard 7
DNS for the DB System 7
Use the Internet and VCN Resolver With Your DB System 8
Hostname Restrictions for Using the Internet and VCN Resolver 9
Custom DNS Resolver 9
Use a Custom DNS Resolver With Your DB System 10
Hostname Restrictions When Using a Custom DNS Resolver 10
DNS: Between On-Premises Network and VCN 11
Set Up DNS for a DB System 11
Security Rules for the DB System 12
General Rules Required for Basic Connectivity 12
General Ingress Rule 1: Allows SSH Traffic From Anywhere 12
General Ingress Rule 2: Allows Path MTU Discovery Fragmentation Messages 13
General Ingress Rule 3: Allows Connectivity Error Messages Within the VCN 13
General Egress Rule 1: Allows All Egress Traffic 13
Custom Security Rules 14
Custom Ingress Rule 1: Allows ONS and FAN Traffic From Within the VCN 14
Custom Ingress Rule 2: Allows SQL*NET Traffic From Within the VCN 14
Custom Egress Rule 1: Allows Outbound SSH Access 15
Custom Egress Rule 2: Allows Access To Oracle Services Network 15
Ways to Implement the Security Rules 15
Use Network Security Groups 16
Use Security Lists 16
Manage Network Security Groups for a DB System 17
Update the Security List for the DB System 17

3 Create

Overview of Creating a DB System 1
Prerequisites 1
Default Options for the Database 2
Use a Backup to Create the Database 2
Custom IP Addresses for the DB Systems 3
Easy Create a DB System 3
Easy Create configuration 3
Procedure 3
Basic details 4

Copyright © 2022, 2026, Oracle and/or its affiliates.

Administrator credentials 4
Network information 4
Default parameters 5
Create 5
Create a DB System 5
General Information 6
Prerequisites 6
Procedure 6
Basic details 6
Shape 6
Storage 8
DB system configuration 9
Network information 10
Diagnostics collection 11
Advanced options 12
Database basic details 13
Administrator credentials 14
Database backups 15
Advanced options 18
Create 19
Create a DB System from a Backup 19
General Information 19
Procedure 20
Select a Backup 20
Basic details 21
Shape 21
Storage 23
DB system configuration 24
Network information 25
Diagnostics collection 26
Advanced options 27
Database basic details 28
Administrator credentials 29
Create 29

4 Update

Upgrade a DB System 1
Prerequisites 1
About Upgrading a DB System 1
Roll Forward a Failed Upgrade 3
Roll Back a Failed Upgrade 3

Copyright © 2022, 2026, Oracle and/or its affiliates.

After your Upgrade is Complete 3
Upgrade the Operating System of a DB System 3
Upgrade the Grid Infrastructure of a DB System 4
View the Upgrade History of a DB System 4
Roll Back a Failed Upgrade 5
Update a DB System 5
Prerequisites 5
Currently Available Updates 6
About Updating a DB System 7
Update the Operating System of a DB System 7
Update the Grid Infrastructure of a DB System 8
View the Update History of a DB System 8
Update the DB System Resources Using dbcli 8
Check for Installed and Available Updates 9
Update Server Components 10
Update Database Home Components 11
Check for Available Operating System Updates for DB System Nodes 13
Update the Operating System of a DB System Node 14
Upgrade a Database 16
About Upgrading Databases 16
Managing Guaranteed Restore Points 20
Upgrade a Database 20
View the Database Upgrade History 21
Roll Back a Failed Database Upgrade 21
Convert a Non-Container Database To an Oracle Database 19c PDB 21
Update a Database 23
About Updating a Database 23
Apply Interim Updates Using a Database Software Image 24
Update a Database 24
View the Database Update History 25
Apply an Interim Update 25

5 Manage

Pluggable Databases 1
About Pluggable Databases 1
Create 1
Restore 1
Relocate 2
Clone 2
Refreshable Clone 3
Convert Refreshable Clone to Regular Pluggable Database 3

Copyright © 2022, 2026, Oracle and/or its affiliates.

SQL Worksheets 3
Backup 4
Open Modes 4
Limitations of Pluggable Database Management 4
Create a Pluggable Database 4
Restore a Pluggable Database 6
In-Place Restore 6
Out-of-Place Restore 7
Relocate a Pluggable Database 7
Clone a Pluggable Database 9
Refresh a Pluggable Database 12
Convert a Refreshable Clone to a Regular Pluggable Database 13
Stop a Pluggable Database 13
Start a Pluggable Database 14
Delete a Pluggable Database 14
Get Connection Strings for a Pluggable Database 15
SQL Worksheets 15
Create a Connection 16
Launch SQL Worksheet 17
Use SQL Worksheet 17
DB Systems 17
List the DB Systems 17
Procedure 17
Filter List Results 18
Change Sort Order 18
Actions 18
Check the Status of a DB System 18
Scale the DB System 19
Change the Shape of a DB System 20
Clone a DB System 24
General Information 24
Limitations 24
Procedure 25
Basic details 25
Network information 25
Diagnostics collection 27
Advanced options 27
Database basic details 28
Administrator credentials 29
Advanced options 29
Clone 29
Manage Licenses on a DB System 29

Copyright © 2022, 2026, Oracle and/or its affiliates.

Move a DB System to Another Compartment 30
Manage Tags for the DB System 30
Terminate a DB System 31
Stop a DB System 31
Start a DB System 32
Reboot a DB System 32
View Work Requests for the DB System 33
Oracle True Cache 33
Migrate to a Dual Stack Network 34
Connect 35
Overview of Connecting to a DB System 35
Prerequisites 35
Database Services and Connection Strings 35
Create an Application Service 36
Database Connection Strings 37
Get the Connection Strings 37
Derive the Connection String 37
Connect to a Database by Using SQLNet 38
Connect to a Database with a Public IP by Using SSH Tunneling 39
Connect to a Database By Using SSH and the Bequeath Protocol 40
Connect From a UNIX-style System 40
Connect From a Windows System 40
Access a Database After You Connect 41
Troubleshoot Connection Issues 42
Manage Serial Console Connection to the DB System 42
Create a Serial Console Connection to the DB System 43
Delete a Serial Console Connection to your DB System 43
Monitor 44
Monitor Base Database Service 44
Available Metrics for Base Database Service Resources 45
Metrics for the DB System in the oci_database_cluster Namespace 45
Metrics for the Database in the oci_database Namespace 46
View Metrics for Base Database Service Resources 50
General Information 51
Prerequisites for Viewing Metrics 51
View Metrics for a DB System 52
View Metrics for a Database 53
View Metrics for a Database in a Compartment 53
View Metrics for a Pluggable Database 54
Monitor Using Database Management Service 54
Manage Database Management for Base Database Service Resources 55
Enable Database Management for a Database 56

Copyright © 2022, 2026, Oracle and/or its affiliates.

Edit Database Management for a Database 58
Disable Database Management for a Database 60
Enable Database Management for a Pluggable Database 61
Edit Database Management for a Pluggable Database 63
Disable Database Management for a Pluggable Database 65
View Performance Hub Metrics for Base Database Service Resources 66
View Performance Hub Metrics for a Database 66
View Performance Hub Metrics for a Pluggable Database 67
Monitor Using Oracle Enterprise Manager 67
Monitor a Database with Enterprise Manager Express 68
Enable the EM Express Console and Determine its Port Number 68
Set the Required Permissions On a 2-node RAC DB System 69
Connect to the EM Express Console 70
Monitor a Database with Enterprise Manager Database Control 71
Determine the Port For the Enterprise Manager Database Control Console 72
Connect to the Enterprise Manager Database Control Console 72
Enable the Console For a Version 11.2.0.4 Database On a Multi-node DB System 73
Events 75
Manage Diagnostics Collection for the DB System 75
General Information 75
Procedure 76
Database Service Events 76
Receive Notifications about Database Service Events 77
Database Service Event Types 77
Database Service Events 77
Temporarily Restrict Automatic Diagnostic Collections for Specific Events 82
Manage Oracle Trace File Analyzer 86
Manage Database Service Agent 87
Incident Logs and Trace Files 88
Event Types for Base Database Service 90
Prerequisites 90
Database Event Types 90
Database Information Event Details 92
Database Critical Event Details 92
DB System Event Types 94
DB System Information Event Details 95
DB System Critical Event Details 95
DB Node Event Types 95
DB Node Information Event Details 96
DB Node Critical Event Details 97
Oracle Database Home Event Types 98
Pluggable Database Event Types 99

Copyright © 2022, 2026, Oracle and/or its affiliates.

Data Guard Association Event Types 100
Remediation for Database Service Events 102
HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE 102
AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN 103
AVAILABILITY.DB_GUEST.CRS_INSTANCE.EVICTION 104
AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN 104
AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN 106
AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN 107
HEALTH.DB_CLUSTER.CDB.CORRUPTION 108
HEALTH.DB_CLUSTER.CDB.ARCHIVER_HANG 110
HEALTH.DB_CLUSTER.CDB.DATABASE_HANG 112
HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE 113
HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE 114
Back Up and Recovery 115
Back Up and Recovery in Base Database Service 115
Prerequisites 117
Managed Backup Features 118
Automatic Incremental and Archived Redo Log Backups 118
Backup Retention 119
Long-Term Retention Backup with Recovery Service 119
Restore Options 120
Protection Policies 120
Protected Databases 121
Real-time Data Protection 121
Backup Deletion Options After Database Termination 121
Backup Scheduling 121
On-Demand Full Backups 122
Standalone Backups 122
Cancel a Running Full or Incremental Backup 122
Backup and Restore from a Standby Database in a Data Guard Association 123
Cross Region Restore 123
Audit and Trace File Retention for Databases Using Automatic Backups 125
Use the API 126
Configure Recovery Service 126
Prerequisite Configuration Checklist for Recovery Service 126
Assign Policies to Allow Access to Recovery Service and Related Resources 127
About Using a Private Subnet for Recovery Service 129
Review Subnet Size Requirements and Security Rules for Recovery Service Subnet 130
Review Networking Service Permissions to Configure a Subnet 131
Create a Recovery Service Subnet in the Database VCN 131
Register Recovery Service Subnet 133
Ensure that the Recovery Service Subnet Can Communicate with Oracle Services 135

Copyright © 2022, 2026, Oracle and/or its affiliates.

Ensure that Your Database Has TDE Fully Configured 135
Turn Off Any Manual Operational Backups 135
Back Up a Database Using the Console 136
Navigate to the List of Standalone Backups for Your Current Compartment 136
Configure Automatic Backups for a Database 136
Configure Automatic Backups for a Standby Database 140
Create an On-Demand Backup of a Database 143
Change the Retention Period of a Long-Term Retention Backup 144
View Details of a Protected Database 145
View Status of a Backup 145
Cancel a Backup 145
Delete Full Backups of a Database 146
Delete Standalone Backups of a Database 146
Back Up a Database to Object Storage Using RMAN 147
Install the Backup Module On the DB System 148
Configure RMAN 149
Back Up the Database 150
Recover a Database Using the Console 151
Recover a Database from Object Storage Using RMAN Backup 152
Set Up Storage on the DB system 153
Perform the Database Restore and Recovery 154
Recover a Database from the OCI Classic Object Store 157
Set Up Storage on the DB System 157
Choose an ORACLE_HOME 159
Copy the Source Database Wallets 159
Install the Oracle Database Backup Module 160
Set Environment Variables 160
Allocate an RMAN SBT Channel 161
Ensure Decryption is Turned On 161
Restore Spfile 161
Set the Database Parameters 161
Restore the Control File 162
Restore the Database 163
Reset the Logs 164
Prepare to Register the Database 164
Register the Database on the DB System 165
Update tnsnames.ora 166
Roll Back Patches on a Version 11.2 Database 166
Post Restore Checklist 167
Oracle Data Guard 168
Use Oracle Data Guard on a DB System 168
General Information 168

Copyright © 2022, 2026, Oracle and/or its affiliates.

Availability Domain and Fault Domain Considerations for Oracle Data Guard 171
Enable Oracle Data Guard on a DB System 171
General Information 171
Create the first standby 172
Create multiple standby databases 179
Perform Database Switchover and Failover 186
Database Switchover 186
Database Failover 186
Edit the Oracle Data Guard Association 187
Reinstate a Database 187
Terminate an Oracle Data Guard Association on a DB System 188
Use Oracle Data Guard with the Database CLI 189
Prerequisites 189
Create a Primary DB System 190
Create a Standby DB System 191
Prepare the Primary DB System 192
Prepare the Standby Database 195
Configure Data Guard 202
Configure Observer (Optional) 204
Encryption Keys 205
Manage Encryption Keys 205
General Information 205
Rotate Encryption Key 206
Assign Key Version 206
Change Key Management 206
Clone, Remote Clone, and Relocate Pluggable Database 207
Rotate Encryption Key for a Database 207
Assign a New Key Version for a Database 208
Change Key Management Type for a Database 208
Rotate Encryption Key for a Pluggable Database 209
Assign a New Key Version for a Pluggable Database 210

6 Secure

Security Guide for Base Database Service 1
Security Overview 1
Security Features 3
User Security 4
Security Settings 5
Security Processes 6
Network Security 7
User Responsibilities for Security Settings 10

Copyright © 2022, 2026, Oracle and/or its affiliates.

Enable Additional Security Capabilities 11
Use Identity and Access Management Authentication with Base Database Service 13
About IAM Authentication 13
Prerequisites 14
Change External Identity Providers 16
Create IAM Groups and Policies for IAM Users 18
Add IAM Users 19
Add IAM Roles 20
Create IAM Database Password for IAM Users 22
Connect to Database with IAM Authentication 22
Configure a Client Connection for SQL*Plus that Uses an IAM Database Password 24
Configure Client Connection for SQL*Plus that Uses an IAM Token 25
Use Instance Principal to Access Database with IAM Authentication 26
Configure Proxy Authentication 27
Use Database Link with IAM Authenticated Users 28
Disable IAM Authentication 29
Use Azure Active Directory Authentication with Base Database Service 29
About Integrating Azure AD with Base Database Service 29
Prerequisites 29
Configure Base Database Service for Integration with Azure AD 31
Map Oracle Database Schemas and Roles 32
Configure Client Connections to Azure ADs 32
Testing the Accessibility of the Azure Endpoint 32
Trace Files Used for Troubleshooting Connections 32
Add SSH Keys to a DB System 32
Open Ports on the DB System 33
Manage Administrator and TDE Wallet Passwords 34
Enable FIPS, SE Linux, and STIG on the DB System Components 34
Security Technical Implementation Guide (STIG) Tool for the DB System 36
Manage Database Security with Oracle Data Safe 38
Security Zone Integration 39
Zero Trust Packet Routing 39
Manage ZPR 40
Manage ZPR Policies 40
Manage Security Attributes 41
Manage Security Attributes for the DB System 41
General Information 41
Add a Security Attribute 42
Edit a Security Attribute 42
Remove a Security Attribute 42
Policy Details for Base Database Service 43
Resource-Types 43

Copyright © 2022, 2026, Oracle and/or its affiliates.

Supported Variables 43
Details for Verb + Resource-Type Combinations 43
Permissions Required for Each API Operation 47
Transparent Data Encryption 49

7 Reference

Oracle Database CLI Reference 1
Operational Notes 1
Syntax 1
CLI Update Command 1
Agent Commands 2
Autologcleanpolicy Commands 3
Backup Commands 4
Backupconfig Commands 7
Component Command 13
Database Commands 14
Dbhome Commands 20
Dbstorage Commands 24
Dgconfig Commands 25
Job Commands 26
Latestpatch Command 28
Logcleanjob Commands 29
Netsecurity Commands 30
Objectstoreswift Commands 32
Pendingjob Command 35
Rmanbackupreport Commands 36
Schedule Commands 38
Scheduledexecution Command 39
Server Command 40
System Command 42
TDE Commands 42
Compartment Quotas 45
General Information 45
Available Quotas 46
Manage Time Zone 47
Time Zone Options 47
View the Current Time Zone 48
Change the Time Zone of the DB System 49
Change the Time Zone of the Host on DB Systems that Use Grid Infrastructure 51
Network Time Protocol 56
Software Images 57

Copyright © 2022, 2026, Oracle and/or its affiliates.

List the Software Images 58
Create a Grid Infrastructure Software Image 58
Create a Database Software Image 59
Create a Database Software Image from an Existing Database 59
View Details of a Software Image 60
Delete a Software Image 60
Move a Software Image to a Different Compartment 60
Provision a Database Using a Database Software Image 61
Update a Database Using a Database Software Image 61
Update Grid Infrastructure using a Grid Infrastructure Software Image 61
Oracle Data Guard Associations 61
Verify the Updates Applied to an Oracle Home 61
Using the API 62
Policy Details for Database Software Images 62
Tags for Base Database Service Resources 63
Importance of Tagging 64
Adding Tags 64
Oracle Standard Tags 65
List of Compliance Regulations 66
Oracle Application Name Tags 67
67
Common Database Export Mechanisms 68
Transport the Data 70
Security and Compliance 70
Performance Best Practices 70
Service-Specific Notes 71
References and Resources 71
Known Issues 72
Existing pluggable databases in new database 72
PDB in existing Data Guard configuration 72
Billing issue when changing license type 73
Backing up to Object Storage using dbcli fails due to certificate change 73
Backing up to Object Storage using RMAN fails due to certificate change 75
Breaking changes in Database service SDKs 76
Unable to use Managed Backups in your DB system 76
Managed Automatic Backups fail on the VM.Standard1.1 shape due to a process crash 78
Oracle Data Pump operations return "ORA-00439: feature not enabled" 79
Unable to connect to the EM Express console from your 1-node DB system 79
OCI Console information not synced for Data Guard enabled databases when using dbaascli 80
Grid Infrastructure does not start after offlining and onlining a disk 80
DB system storage scaling failure 84

Copyright © 2022, 2026, Oracle and/or its affiliates.

DB system shape scaling failure 84
Troubleshoot 84
84
Troubleshoot Backup Failures 93
Identify the Cause of Failure 94
Database Service Agent Issues 95
Oracle Clusterware Issues 96
Object Store Connectivity Issues 97
Host Issues 98
Database Issues 98
TDE Wallet Issues 99
Other Causes of Backup Failures 103
Get Additional Help 105
Troubleshoot Update Failures 105
Identify the Cause of Failure 106
Database Service Agent Issues 107
Object Store Connectivity Issues 108
Host Issues 109
Oracle Clusterware Issues 109
Database Issues 110
Get Additional Help 111
Troubleshoot Shape Change Failures 111
Use the OCI Console to Troubleshoot 112
Use dbcli to Troubleshoot 112
Get Additional Help 113
Troubleshoot Network Connectivity Failures 113
Identify the Cause of Failure 114
Resolve Network Connectivity Failures 115
Get Additional Help 115

Copyright © 2022, 2026, Oracle and/or its affiliates.

1

Overview

What's New in Oracle Base Database Service

Oracle is constantly adding new capabilities to the Base Database Service. This article provides a brief overview of the new features and enhancements made to the Base Database Service. It is organized by the date a particular feature or capability became available.

Table 1-1 Whats New in Oracle Base Database Service

Release Date	Feature	Description
September 02, 2026	Easy Create DB systems	You can now use Easy Create to provision a DB system with minimal configuration. Easy Create applies predefined compute, storage, licensing, and database settings. You can modify supported settings after provisioning. For more information, see *Easy Create a DB system*.
June 29, 2026	New VM.BaseDB.x86 flexible shape	Introducing VM.BaseDB.x86, a new generation-agnostic x86 shape for Base Database DB systems. This shape lets you create and manage DB systems with flexible ECPU allocation, including the ability to update ECPU allocations after provisioning. Supported DB systems using VM.Standard.x86 can also be migrated to VM.BaseDB.x86. For more information, see *About DB Systems*.
June 23, 2026	Change in Backup Destination for New Tenancies in Select OCI Regions (Effective June 23, 2026)	Starting June 23, 2026, Autonomous Recovery Service is the only available backup destination for automatic backups enabled through the OCI Console for tenancies created on or after June 23, 2026, in the following OCI regions: Netherlands Northwest (Amsterdam) AMS, Colombia Central (Bogota) BOG, Brazil East (Sao Paulo) GRU, Indonesia North (Batam) HSG, India South (Hyderabad) HYD, US East (Ashburn) IAD, South Korea Central (Seoul) ICN, Malaysia West 2 (Kulai) JBP, Saudi Arabia West (Jeddah) JED, Morocco West (Casablanca) LEJ, UK South (London) LHR, Italy Northwest (Milan) LIN, Spain Central (Madrid) MAD, Israel Central (Jerusalem) MTZ, Italy North (Turin) NRQ, Spain Central (Madrid 3) ORF, Chile Central (Santiago) SCL, Chile West (Valparaiso) VAP, Brazil Southeast (Vinhedo) VCP, Canada Southeast (Montreal) YUL, and Canada Southeast (Toronto) YYZ. For more information, see *Create a DB System*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 1-1 (Cont.) Whats New in Oracle Base Database Service

Release Date	Feature	Description
May 18, 2026	Multiple Standby Databases using Oracle Data Guard Now Available in All OCI Regions	Multiple Standby Databases using Oracle Data Guard feature is now available in all OCI regions, including UAE Central (Abu Dhabi) AUH, UAE East (Dubai) DXB, Saudi Arabia West (Jeddah) JED, Israel Central (Jerusalem) MTZ, and Saudi Arabia Central (Riyadh) RUH. For more information, see *Use Oracle Data Guard on a DB System*.
May 15, 2026	Change in Backup Destination for New Tenancies in Select OCI Regions (Effective May 15, 2026)	Starting May 15, 2026, Autonomous Recovery Service is the only available backup destination for automatic backups enabled through the OCI Console for tenancies created on or after May 15, 2026, in the following OCI regions: Sweden Central (Stockholm) ARN, UAE Central (Abu Dhabi) AUH, India West (Mumbai) BOM, France Central (Paris) CDG, UK West (Newport) CWL, UAE East (Dubai) DXB, South Africa Central (Johannesburg) JNB, Japan Central (Osaka) KIX, Australia Southeast (Melbourne) MEL, France South (Marseille) MRS, Mexico Northeast (Monterrey) MTY, US Midwest (Chicago) ORD, Mexico Central (Queretaro) QRO, Saudi Arabia Central (Riyadh) RUH, US West (San Jose) SJC, Singapore (Singapore) SIN, Australia East (Sydney) SYD, Singapore West (Singapore) XSP, South Korea North (Chuncheon) YNY, and Switzerland North (Zurich) ZRH. For more information, see *Create a DB System*.
April 02, 2026	Multiple Standby Databases using Oracle Data Guard	This enhancement provides the ability to use Oracle Data Guard to create and manage multiple local and remote standby databases linked to a primary database, providing flexibility for both data protection and disaster recovery. This feature is available in all regions except UAE Central (Abu Dhabi) AUH, UAE East (Dubai) DXB, Saudi Arabia West (Jeddah) JED, Israel Central (Jerusalem) MTZ, and Saudi Arabia Central (Riyadh) RUH. For more information, see *Use Oracle Data Guard on a DB System*.
March 28, 2026	Update Operating System	You can now update the operating system of a DB system using the Console or APIs. For more information, see *Update a DB System*.
September 10, 2025	Oracle Base Database Service for Developers	Oracle Base Database Service for Developers allows developers to build applications on the Oracle Cloud without paying any Oracle Database license fees. For more information, see *About Oracle Base Database Service*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 1-1 (Cont.) Whats New in Oracle Base Database Service

Release Date	Feature	Description
August 27, 2025	Generation-agnostic X86 flexible shape	The generation-agnostic X86 shape is the latest generation-agnostic and flexible shape that enables you to customize the number of ECPUs allocated to an instance. You can now create and manage DB systems on the X86 shape. For more information, see *About DB Systems*.
August 06, 2025	Change in Backup Destination for New Tenancies in Select OCI Regions (Effective August 06, 2025)	Starting August 06, 2025, the Autonomous Recovery Service is now the exclusive backup destination during automatic backup configuration for newly created tenancies in the following OCI regions: Frankfurt (FRA), Phoenix (PHX), and Tokyo (NRT). For more information, see *Create a DB System*.
June 25, 2025	New experience in Oracle Cloud Console	We've launched a major update to the Oracle Cloud Console. The new experience includes intuitive designs across the Console to enhance usability and consistency, with significant updates to resource list pages, resource details pages, and resource creation workflows. For more information, see New Console Experience.
May 29, 2025	Cross-region Data Guard support for OCI Virtual Vault	You can now set up a cross-region Data Guard between two databases. It enables you to have the encryption keys used for the primary and standby databases available in the primary and standby regions, respectively, so that they provide protection against a single point of failure for the OCI Vault key. For more information, see *Use Oracle Data Guard on a DB System*.
February 20, 2025	IPv4/IPv6 dual stack network support	You can now provision DB systems with IPv4/IPv6 dual stack networking. This enables applications to scale seamlessly by leveraging the expanded IPv6 address space, addressing IPv4 exhaustion concerns, and ensuring readiness for future growth without network limitations. For more information, see *VCN and Subnets*.
January 23, 2025	Long-term retention backup (LTR) for Autonomous Recovery Service	Long-term retention backup (LTR) allows you to store full backups for periods up to ten years for compliance, regulatory, or other business needs with complete LTR lifecycle management and immutability. For more information, see *Back Up and Recovery in Base Database Service*.
December 11, 2024	Upgrade Database and Grid Infrastructure software version to 26ai	You can now upgrade the Oracle Database and the Oracle Grid Infrastructure software version to 26ai using the Console or APIs. For more information, see *Upgrade a Database* and *Upgrade a DB System*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 1-1 (Cont.) Whats New in Oracle Base Database Service

Release Date	Feature	Description
November 25, 2024	Migration from Intel X7 or AMD E4 to AMD E5 Shape	You can now migrate your DB system from Intel-based X7 shapes or AMD-based E4 flexible shapes to AMD-based E5 flexible shapes by using the *change shape* operation. For more information, see *Change the Shape of a DB System*.
November 22, 2024	Unified Auditing	You can now enable unified auditing in Base Database Service. It provides selective and effective auditing inside the Oracle database using policies and conditions. For more information, see *Create a DB System*.
October 01, 2024	Zero Trust Packet Routing (ZPR)	You can now protect your resources with additional security using Oracle Cloud Infrastructure Zero Trust Packet Routing (ZPR). ZPR protects sensitive data from unauthorized access through intent-based security policies. For more information, see *Zero Trust Packet Routing*.
August 28, 2024	Assign new key versions to container databases and pluggable databases	You can now assign new key versions and rotate keys for container databases and pluggable databases independently. For more information, see *Manage Encryption Keys*.
June 26, 2024	Custom Software Images	The ability to create a custom software image for Oracle Database and Grid Infrastructure with all the required updates bundled together and certified in the customer environment will allow developers and database administrators to build an approved and reusable *gold image*. For more information, see *Software Images*.
June 11, 2024	AMD E5 flex shape	You can now use the AMD E5 flexible shape for your DB system. Flexible shapes let you customize the number of OCPUs allocated to an instance. For more information, see *About DB Systems*.
May 08, 2024	Backup and restore from a standby database in a Data Guard association for Autonomous Recovery Service	You can now backup and restore from a standby database in a Data Guard association for the Autonomous Recovery Service. For more information, see *Back Up and Recovery in Base Database Service*.
May 02, 2024	Oracle AI Database 26ai is available on Base Database Service	Oracle AI Database 26ai is a regular production release available on Base Database Service. With this release, you can perform all the lifecycle operations on the 26ai database. For more information, see *About Oracle Base Database Service*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 1-1 (Cont.) Whats New in Oracle Base Database Service

Release Date	Feature	Description
April 03, 2024	Restore a backup to create a database across regions	You can now use an existing backup and restore it to create a database (out-of-place restore) across availability domains within the same region or in a different region using a backup created with Object Storage or Autonomous Recovery Service. You can also restore a backup taken on either a database that was configured using host-based wallets (local wallets) or OCI Vault. For more information, see *Back Up and Recovery in Base Database Service*.
February 05, 2024	OCI Vault integration for cross-region Data Guard	You can now set up a cross-region Data Guard between two databases if their keys are residing in a virtual private vault and are managed by the OCI Vault service. It enables you to have the encryption keys used for the primary and standby databases available in the primary and standby regions, respectively, so that they provide protection against a single point of failure for the OCI Vault key. For more information, see *Use Oracle Data Guard on a DB System*
December 13, 2023	Use SQL Worksheet for Pluggable Databases	The SQL Worksheet provides a web-based SQL workspace where you can work with SQL statements directly in the browser. You can now launch and use the SQL Worksheet for your Pluggable Databases from the Console. For more information, see *SQL Worksheets*.
December 06, 2023	Upgrade operating system to Oracle Linux 8	You can now upgrade the operating system of your DB system to Oracle Linux 8 (OL8) using the Console or APIs. For more information, see *Upgrade a DB System*.
December 06, 2023	Backup and restore from a standby database in a Data Guard association	You can now backup and restore from a standby database in a Data Guard association. For more information, see *Back Up and Recovery in Base Database Service*.
October 11, 2023	Enhancements to Pluggable Database (PDB) management	You can now restore, relocate, and refresh Pluggable Databases. For more information, see *About Pluggable Databases*.
August 21, 2023	Autonomous Recovery Service as the Default Backup Destination	Autonomous Recovery Service is now the default backup destination for automatic backups in all regions and includes default limits without having to request them. For more information, see *Back Up a Database Using the Console*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 1-1 (Cont.) Whats New in Oracle Base Database Service

Release Date	Feature	Description
July 26, 2023	Restore a backup to create a database across availability domains within the same region	You can now use an existing backup and restore it to create a new database (out-of-place restore), either within the same availability domain or in a different availability domain within the same region, using a backup created with Object Storage or Autonomous Recovery Service. For more information, see *Create a DB System from a Backup*.
June 28, 2023	Ampere A1 flex shape	You can now use Arm-based Ampere A1 flexible shape for your DB system. Flexible shapes let you customize the number of OCPUs allocated to an instance. For more information, see *About DB Systems*.
May 17, 2023	Enhanced controls to configure automatic full (L0) and incremental (L1) backups	You can now use these enhanced controls for database backups. With these enhanced controls, you can: • Specify whether you want the initial L0 backup to start immediately or according to the L0 schedule. • Choose a time window for the future full backups to start. • Choose a time window for the incremental backups to start, which can be different from the time window for the L0 backups. The time windows will remain the same, two-hour scheduling windows and the default six-hour window. For more information, see: • *Back Up and Recovery in Base Database Service* • *Back Up a Database Using the Console*
April 26, 2023	Cancel a running full or incremental backup	You can now cancel an ongoing backup, allowing you to free up system resources. You will no longer have to call the Oracle Cloud Operations team to have this backup job canceled. For more information, see: • *Back Up and Recovery in Base Database Service* • *Back Up a Database Using the Console*
March 29, 2023	Health and performance metrics for DB Systems and databases in the Console	You can now monitor the health, capacity, and performance of your DB systems and databases with metrics, alarms, and notifications. You can use Console, Monitoring APIs, or Database Management APIs to view metrics. For more information, see: • *View Metrics for Base Database Service Resources* • *Available Metrics for Base Database Service Resources*

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 1-1 (Cont.) Whats New in Oracle Base Database Service

Release Date	Feature	Description
March 22, 2023	Configure Oracle Database Autonomous Recovery Service as a backup destination	Oracle Database Autonomous Recovery Service provides an optimized policy-driven automatic backup and recovery system for the Base Database Service. It also offers a real-time data protection feature that enables protected databases with zero data loss recovery in the event of a database failure. Since real-time data protection is an extra cost option, you can choose to enable or disable it. For more information, see *Back Up and Recovery in Base Database Service*.
January 20, 2023	Service gateway rules to support upcoming features	To support upcoming features, OCI cloud automation needs to communicate with various OCI services such as Identity, Object Storage Service (OSS) etc. either using an internet gateway or a service gateway. For more information, see *VCN and Subnets*.
January 11, 2023	Intel X9 flex shape	You can now use Intel X9 flexible shape for your DB system. Flexible shapes let you customize the number of OCPUs allocated to an instance. For more information, see *About DB Systems*.
January 11, 2023	Performance Metrics for Pluggable Databases	You can now monitor the health and performance of your PDBs by using Database Management metrics and Performance Hub. For more information, see: • *View Metrics for Base Database Service Resources* • *View Performance Hub Metrics for Base Database Service Resources*
December 21, 2022	Microsoft Azure Active Directory integration with Base Database Service	You can now configure the database in the Base Database Service to use Microsoft Azure Active Database authentication and authorization to allow Azure AD users to access the database with Azure AD credentials. For more information, see *Use Azure Active Directory Authentication with Base Database Service*.
December 21, 2022	Documentation added on troubleshooting network connectivity failures	New articles have been added to the Base Database Service documentation to help you in troubleshooting network connectivity failures. For more information, see *Troubleshoot Network Connectivity Failures*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 1-1 (Cont.) Whats New in Oracle Base Database Service

Release Date	Feature	Description
December 21, 2022	Documentation added on managing Database Management	New articles have been added to the Base Database Service documentation to help you in managing Database Management. For more information, see: • *Manage Database Management for Base Database Service Resources* • *View Metrics for Base Database Service Resources* • *View Performance Hub Metrics for Base Database Service Resources*
December 1, 2022	Migration from Intel X7 to AMD E4 Shape	You can now migrate your DB system from Intel-based X7 shapes to AMD-based flexible shape E4 by using the **Change shape** operation. For more information, see *Change the Shape of a DB System*.
December 1, 2022	New documentation added on security	New article have been added to the Base Database Service documentation to help you with security features. For more information, see *Security Guide for Base Database Service*.
November 16, 2022	Documentation added on troubleshooting failures	New articles have been added to the Base Database Service documentation to help you in troubleshooting failures. For more information, see: • *Troubleshoot Backup Failures* • *Troubleshoot Update Failures* • *Troubleshoot Shape Change Failures*
September 27, 2022	Identity and Access Management (IAM) integration with Base Database Service	With the latest Release Update, you can now configure the database in a DB system to use OCI Identity and Access Management (IAM) authentication and authorization to allow IAM users to access the database with IAM credentials. As of this release, IAM authentication and authorization: • will be available to enable on newly provisioned databases in a DB system and for existing virtual machines that have been patched to set the `WALLET_ROOT` system parameter. • can not be used with databases configured with Data Guard. For more information, see *Use Identity and Access Management Authentication with Base Database Service*.
September 20, 2022	Oracle Standard Tags for Base Database Service Resources	Using the OCI tagging system, Base Database Service resources may now be tagged according to your organizational scheme, allowing you to group resources, manage costs, and gain usage insights. For more information, see *Tags for Base Database Service Resources*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 1-1 (Cont.) Whats New in Oracle Base Database Service

Release Date	Feature	Description
August 30, 2022	Database Service Events and Log Collection	The diagnostics collection and notifications feature enables Oracle Cloud Operations and you to identify, investigate, track, and resolve guest VM issues quickly and effectively. Database Service Events feature enables you to get notified about health issues with your Oracle Databases or other components on the DB system. For more information, see: • *Manage Diagnostics Collection for the DB System* • *Database Service Events* • *Incident Logs and Trace Files* • *Manage Log and Diagnostic Files* • *Remediation for Database Service Events*
July 27, 2022	TCPS integration	You can now use TCPS to connect when enabling database management support for the Base Database Service. For more information, see Enable Database Management for Oracle Cloud Databases.
July 19, 2022	Database Service Events name change	The event names for Base Database Service have changed. For more information, see *Event Types for Base Database Service*.
July 14, 2022	Data Guard enhancements	The Data Guard has now been enhanced with the following functionalities: • Select your desired time zone for the standby database, • Change the database image for the standby database, • Have a different fault domain for primary and standby databases, For more information, see *Enable Oracle Data Guard on a DB System*.
July 07, 2022	Database upgrade with Data Guard	You can now upgrade the database that has a Data Guard association. The upgrade options are available in the Console for Data Guard associations created using the Console. For more information, see *Upgrade a Database*.
May 18, 2022	AMD E4 flex shape	Flexible shapes let you customize the number of OCPUs allocated to an instance. When you create a DB system using a flexible shape, you select the number of OCPUs that you need for the workloads that run on the instance. This flexibility lets you build DB systems that match your workload, enabling you to optimize performance and minimize cost. For more information, see *About DB Systems*.
April 30, 2022	Image based patching	You can now update a database using the database software images using the Console. For more information, see *Update a Database*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 1-1 (Cont.) Whats New in Oracle Base Database Service

Release Date	Feature	Description
April 21, 2022	Upgrade operating system and Oracle Grid Infrastructure	You can now upgrade the operating system and the Oracle Grid Infrastructure using the Console or APIs. For more information, see *Upgrade a DB System*.

About Oracle Base Database Service

Oracle Base Database Service enables you to maintain absolute control over your data while using the combined capabilities of Oracle Database and Oracle Cloud Infrastructure.

Oracle Base Database Service offers database systems (DB systems) on virtual machines (VMs). They are available as single-node DB systems and multi-node RAC DB systems on Oracle Cloud Infrastructure (OCI). You can manage these DB systems by using the OCI Console, the OCI API, the OCI CLI, the Database CLI (DBCLI), Enterprise Manager, or SQL Developer.

> **Note**
>
> Oracle Base Database Service has announced the retirement of VM.Standard.x86 shape, the Enterprise Edition - Extreme Performance (EE-EP) software edition, and OCPU-metered service SKUs, effective July 31, 2026. Existing deployments will continue to run, but new provisioning and contract renewals will be affected.
>
> Plan migration to VM.BaseDB.x86 using the ECPU metric or to Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS) for EE-EP workloads, as applicable. For information about affected SKUs, customer impact, and migration guidance, see My Oracle Support Article KB921246.

Supported Database Editions

Oracle Base Database Service supports the following Oracle Database editions:

- Standard Edition
- Enterprise Edition
- Enterprise Edition - High Performance
- Enterprise Edition - Extreme Performance
- Enterprise Edition - Developer

> **Note**
>
> Oracle Enterprise Edition - Extreme Performance is required for multi-node RAC DB systems.

Copyright © 2022, 2026, Oracle and/or its affiliates.

About Easy Create

Easy Create provides a simplified workflow for provisioning a DB system with minimal configuration. The workflow requires only the database name, an Oracle AI Database software edition, administrator credentials, and network information.

Easy Create applies predefined compute, storage, licensing, and database settings and displays them for review before provisioning. Certain settings remain configurable after the DB system is provisioned. The standard creation workflow remains available for DB systems that require advanced configuration during provisioning.

For more information, see *Easy Create a DB system*.

About Oracle Base Database Service for Developers

Oracle Base Database Service for Developers allows developers to build applications on Oracle Base Database Service in the Oracle Cloud without paying any Oracle Database license fees. The infrastructure resources, such as storage, memory, network bandwidth, etc., will incur usual usage costs.

A DB system with Oracle Base Database Service for Developers is provisioned in the same manner as the other generally available DB systems by selecting Enterprise Edition Developer as the database edition. There are no additional limitations on the number of OCPUs, memory, database storage, or PDBs that can be deployed per DB system. For detailed steps to provision a DB system, see *Create a DB System* .

Oracle Base Database Service for Developers is available for Oracle Database versions 26ai and 19c starting with the 23.8.0.0, 19.27.0.0, and later release updates (RU) only.

Oracle Base Database Service for Developers software cannot be used for production workloads.

Oracle Base Database Service for Developers Restrictions

The following restrictions apply to the DB systems using Oracle Database Enterprise Edition Developer and will be enforced by the cloud automation and the database instance:

- Only a single instance can be provisioned. Two-instance RAC deployments are not supported.
- Supported only on Ampere A1 shape.
- Supported only on Logical Volume Manager.
- Supported only on single-node DB systems.
- Oracle Database software edition cannot be changed after provisioning.
- Data Guard association is not supported.
- Cross-region backup and restore is not supported.
- Bring Your Own License (BYOL) is not supported.
- Only the backups created using Enterprise Edition Developer can be used to create a DB system from a backup using Enterprise Edition Developer.
- Additionally, the following restrictions apply to the PDBs and will be automatically enforced by the database:
  - Maximum of 2 threads per PDB.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Maximum of 8 GB memory per PDB.
- Maximum of 20 GB database size per PDB.
- Maximum of 30 sessions per PDB.

Supported Database Versions

Oracle Base Database Service supports the following Oracle Database versions:

- Oracle AI Database 26ai
- Oracle Database 21c
- Oracle Database 19c

About Oracle AI Database 26ai

Oracle AI Database 26ai is the next long-term support release of Oracle Database. It includes over 300 new features, with a focus on artificial intelligence (AI) and developer productivity. Features such as AI Vector Search enable you to leverage a new generation of AI models to generate and store vectors of documents, images, sound, and so on; index them and quickly look for similarity while leveraging the existing analytical capabilities of Oracle Database. This, combined with the already extensive set of Machine Learning algorithms, enables you to quickly create sophisticated AI-enabled applications.

For more information about the new features, see Oracle Database New Features.

Upgrade the DB System

You can upgrade the DB system that uses earlier versions to the current version. For more information, see *Upgrade a DB System*.

Update the Operating System of a DB System

> **Tip**
>
> Run a precheck to determine whether updates are available for the operating system.

Perform the following steps to update the operating system (OS) of a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, select the **Updates (OS)** tab to view the list of available OS updates for the DB system.
3. Under the **OS update** section, from the **Actions** menu of the status, select the following actions:
   a. **Precheck**: Check for any prerequisites to ensure that the update can be successfully applied.
   b. **Apply OS Update**: Apply the selected update. The latest available OS update is always applied. Oracle recommends that you run the precheck for an update before applying it.

Copyright © 2022, 2026, Oracle and/or its affiliates.

4. To check the status or review details of prechecks or OS updates, go to the **Update History** tab. From the **Actions** menu for the relevant operation, select **View details**.

For 2-node RAC DB systems, you can apply OS updates to individual nodes one at a time, or to both nodes simultaneously by selecting them during the precheck and update process.

> **Warning**
>
> For a single-node DB system, you can expect downtime for the DB system while the OS is updated.

In the list, the **State** displays the status of the operation. While an update is being applied, the State changes to **Updating** and the status of the DB system changes to **Updating**. Lifecycle operations on the DB system and its resources might be temporarily unavailable. When the update completes successfully, the State changes to **Applied** and the status of the DB system changes to **Available**.

Upgrade the Database in a DB System

You can upgrade database instances that use earlier Oracle Database versions to later Oracle Database versions. For more information, see *Upgrade a Database*.

Update the Database in a DB System

You must update the database in your DB system periodically to ensure proper functioning. Oracle recommends updating the DB system before you update the database within that DB system. For more information, see *Update a Database*.

Oracle Database Preview Version Availability

OCI periodically offers preview software versions of Oracle Database for testing purposes. You can provision a DB system using preview version software to test applications before the general availability of the software in the Database service. When you provision a DB system with preview version software, the DB system remains available to you until you decide to terminate it.

Preview version DB systems are provisioned in the same manner as generally available DB systems. If available, preview version software is displayed as one of the choices in the database version selector.

Oracle Database Preview Version Restrictions

Preview version software cannot be used for production databases. The following restrictions apply to preview version software:

- Preview version software is not available for DB systems using RAC.
- Uses Logical Volume Manager (LVM) storage management software only. Automatic Storage Management (ASM) is not available.
- Patching and database version upgrades are not available.
- You cannot upgrade the preview version software to its generally available release.
- You cannot create a new DB system from a backup of a database that uses preview version software.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Standalone backups cannot be created.
- Data Guard is not available.
- Preview version software DB systems cannot be created from backups.
- In-place restores are supported.

Per-Second Billing for DB Systems

OCI uses per-second billing. This means that OCPU and storage usage are billed by the second, with a minimum usage period of one minute for the DB systems.

Customer Managed Keys for Databases

OCI provides customer-managed keys that enable you to encrypt your data using encryption keys that you control using the OCI Vault service. The Vault service provides you with centralized key management capabilities that are highly available and durable. For more information, see *Manage Encryption Keys*.

Backup and Recovery

OCI enables you to create and store automatic daily backups and on-demand full backups. You can store backups in your DB system's local storage or in OCI Object Storage.

You can recover a database from the following sources:

- Object Storage using RMAN backup.
- OCI Classic Object Store.

You can create a DB system from the following sources:

- Daily automatic backups or on-demand full backups.
- The last archived redo log backup. Requires that you have automatic backups enabled. This backup combines data from the most recent daily automatic backup and data from archived redo logs, and represents the most current backup available.
- Daily automatic backup data used to create a point-in-time copy of the source database based on a specified time stamp.
- Standalone Backups.

For more information, see *Back Up and Recovery in Base Database Service*.

Move Databases to Oracle Cloud Using Zero Downtime Migration

Oracle now offers two new solutions for migrating Oracle Database workloads from on-premises and OCI Classic to a variety of Oracle Database Cloud Services: **Oracle Cloud Infrastructure Database Migration Service** and **Zero Downtime Migration**.

Zero Downtime Migration (ZDM) is an installable tool that provides you with a simplified and automated migration experience, providing zero to negligible downtime for the production system.

OCI Database Migration Service is based on the ZDM tool, and as a managed OCI service, it provides you with a user interface to move Oracle Databases to Oracle Cloud.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Support Requirements

Oracle Database Support is required to run a database service on Base Database Service. Customers running the license included option have support through their database cloud subscription. Customers bringing their own licenses to Base Database Service must have a valid support contract covering the database licenses they are running in the database cloud service. Customers are responsible for applying updates to their databases in a timely manner to protect from product issues and security vulnerabilities.

Oracle Applications Unlimited (specifically Oracle E-Business Suite, PeopleSoft, JD Edwards, Siebel and Hyperion applications) require ongoing updates to ensure continued interoperability. As such, Oracle requires customers to have an active Oracle support contract for any such programs running on Oracle Base Database Service. Note that third-party support does not satisfy this requirement. This requirement shall not be interpreted as superseding anything in the Oracle Technical Support Policies, including, but not limited to, the Matching Service Level requirement.

About DB Systems

Oracle Cloud Infrastructure (OCI) offers DB systems on virtual machines.

There are two types of database systems (DB systems) on virtual machines:

- Single-node DB system: A 1-node DB system consists of one virtual machine.
- Multi-node RAC DB system: A 2-node DB system consists of two virtual machines.

If you must provision a DB system for development or testing purposes, a special fast-provisioning single-node DB system is available.

When you create a DB system, you select the Oracle Database edition and version that applies to the database on that DB system. You cannot change the selected edition. Depending on your selected Oracle Database edition and version, your DB system can support multiple pluggable databases (PDB). See the following Oracle Database licensing topic for information about the maximum number of pluggable and container databases (CDB) available for your selected Oracle Database version.

- Oracle Database 19c: Permitted Features, Options, and Management Packs by Oracle Database Offering

A DB system can have only a single Database Home, which in turn can have only a single database. A DB system database uses OCI block storage instead of local storage. You specify a storage size when you create the DB system, and you can scale up the storage as required at any time. To change the number of CPU cores on an existing DB system, you must change the shape of that DB system. For more information, see *Change the Shape of a DB System*.

> **Note**
>
> The shape change operation takes place in a rolling fashion for multi-node RAC DB systems, enabling you to change the shape with no database downtime.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Available Shapes and How It Determines the Resources Allocated

When you create a DB system, you select a *shape*, which determines the resources allocated to the DB system. After you create the DB system, you can change its shape to adapt to new processing capacity requirements. The following shapes are available:

> **Note**
>
> Oracle Base Database Service has announced the retirement of VM.Standard.x86 shape, the Enterprise Edition - Extreme Performance (EE-EP) software edition, and OCPU-metered service SKUs, effective July 31, 2026. Existing deployments will continue to run, but new provisioning and contract renewals will be affected.
>
> Plan migration to VM.BaseDB.x86 using the ECPU metric or to Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS) for EE-EP workloads, as applicable. For information about affected SKUs, customer impact, and migration guidance, see My Oracle Support Article KB921246.

Flexible Shapes

Flexible shapes let you customize the compute resources - ECPUs or OCPUs allocated to an instance. When you create an instance using a flexible shape, you select the compute resource that you require for the workloads that run on the instance. This flexibility lets you build instances that match your workload, enabling you to optimize performance and minimize cost. The amount of memory allowed is based on the number of ECPUs or OCPUs selected, and the ratio of memory to ECPUs or OCPUs depends on the shape.

Flexible shapes are available with Ampere, AMD, and Intel processors. The following table shows the available shapes.

Table 1-2 Flexible Shapes

Shape	Compute	Memory
VM.BaseDB.x86	Minimum is **4 ECPUs** and maximum is **256 ECPUs**. Should be in increments of 4.	**8 GB** for every **4 ECPUs**. Minimum is **8 GB** and maximum is **512 GB** total memory.
VM.Standard.x86	Minimum is **4 ECPUs** and maximum is **252 ECPUs**. Should be in increments of 4.	**8 GB** for every **4 ECPUs**. Minimum is **8 GB** and maximum is **504 GB** total memory.
Ampere VM.Standard.A1.Flex	Minimum is **1 OCPU** and maximum is **57 OCPUs**.	**8 GB** per OCPU. Minimum is **8 GB** and maximum is **456 GB** total memory.
AMD VM.Standard.E5.Flex	Minimum is **1 OCPU** and maximum is **64 OCPUs**.	**16 GB** per OCPU. Minimum is **16 GB** and maximum is **1024 GB** total memory.
AMD VM.Standard.E4.Flex	Minimum is **1 OCPU** and maximum is **64 OCPUs**.	**16 GB** per OCPU. Minimum is **16 GB** and maximum is **1024 GB** total memory.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 1-2 (Cont.) Flexible Shapes

Shape	Compute	Memory
Intel X9 VM.Standard3.Flex	Minimum is **1 OCPU** and maximum is **32 OCPUs**.	**16 GB** per OCPU. Minimum is **16 GB** and maximum is **512 GB** total memory.

> **Note**
>
> - Arm-based Ampere A1 shape is available for Oracle Database versions 26ai and 19c starting with the 23.7.0.0, 19.19.0.0, and later release updates (RU) only.
> - AMD E5 shape is available for Oracle Database versions 26ai and 19c starting with the 23.4.0.24.05, 19.21.0.0, and later release updates (RU) only.
> - AMD E4 shape is available for Oracle Database versions 26ai, 21c, and 19c with the 23.4.0.24.05, 21.6.0.0, 19.15.0.0, and later release updates (RU) only.
> - Intel X9 shape is available for Oracle Database versions 26ai, 21c, and 19c with the 23.4.0.24.05, 21.8.0.0, 19.17.0.0, and later release updates (RU) only.
> - Multi-node RAC DB systems require a minimum of two OCPUs per node.

Generation-Agnostic X86 Shape

The **VM.BaseDB.x86** shape and **VM.Standard.x86** shape are generation-agnostic, flexible x86 shapes that let you customize the number of ECPUs allocated to a DB system. These shapes provide a consistent experience across supported x86 infrastructure generations while allowing you to scale compute resources based on workload requirements.

The x86 shapes have the following characteristics:

- Use **ECPU-based billing**, enabling more granular, usage-based pricing.
- VM.BaseDB.x86 supports Oracle AI Database 26ai and Oracle Database 19c.
- VM.Standard.x86 supports Oracle AI Database 26ai only.
- Support the following Oracle Database software editions:
  - Standard Edition
  - Enterprise Edition
  - Enterprise Edition - High Performance
- Support **Logical Volume Manager** for storage management.
- Support **single-node DB systems** only.

> **Note**
>
> The x86 shapes do not support **fault domain**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Arm-based Ampere A1 Shape

Arm-based Ampere A1 shapes are flexible and enable you to customize the number of OCPUs allocated to an instance. The following are some additional details about Ampere A1 shapes:

- Ampere A1 shape is only supported on Logical Volume Manager.
- Ampere A1 shape is only supported on single-node DB systems.
- Oracle Database Standard Edition is not supported on Ampere A1 shape-based DB systems.
- A database software image cannot be used for creating a database on Ampere A1 shape-based DB systems.
- Ampere A1 shape-based DB system provisioning and restoration are not supported if the backup destination for the database is the Autonomous Recovery Service.
- Ampere A1 shape is not supported for databases that use OCI vault encryption.
- The shape of Ampere A1 shape-based DB systems cannot be changed to Intel or AMD shape-based DB systems, and vice versa.
- A backup of an Ampere A1 shape-based database cannot be restored on Intel or AMD shape-based DB systems, and vice versa.
- Ampere A1 shape-based DB systems do not support Data Guard associations with Intel or AMD shape-based DB systems.

Standard Shapes

Standard shapes are available with Intel processors.

The following table shows the available shapes in the X7 series.

Table 1-3 VM Available Shapes X7 Series

Shape	CPU Cores	Memory
VM.Standard2.1	1	15 GB
VM.Standard2.2	2	30 GB
VM.Standard2.4	4	60 GB
VM.Standard2.8	8	120 GB
VM.Standard2.16	16	240 GB
VM.Standard2.24	24	320 GB

Note

- Intel X7 shapes are available for Oracle Database versions 26ai, 21c, and 19c only.
- VM.Standard2.1 shape cannot be used for multi-node RAC DB system.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Available Database Versions

OCI supports the creation of DB systems using older database versions. For each shape, the latest version and the two prior versions of the release are available at provisioning with the following specifications.

- Arm-based Ampere A1 shape is available for Oracle Database versions 26ai and 19c starting with the 23.7.0.0, 19.19.0.0, and later release updates (RU) only.
- AMD E5 shape is available for Oracle Database versions 26ai and 19c starting with the 23.4.0.24.05, 19.21.0.0, and later release updates (RU) only.
- AMD E4 shape is available for Oracle Database versions 26ai, 21c, and 19c with the 23.4.0.24.05, 21.6.0.0, 19.15.0.0, and later release updates (RU) only.
- Intel X9 shape is available for Oracle Database versions 26ai, 21c, and 19c with the 23.4.0.24.05, 21.8.0.0, 19.17.0.0, and later release updates (RU) only.
- Intel X7 shapes are available for Oracle Database versions 26ai, 21c, and 19c only.
- Migration from Intel X9 shape to AMD E4 and E5 shapes is not supported.
- Migration to AMD E4 shape is supported for instances using the base image with 21.6.0.0, 19.15.0.0, and later release updates only. For instances created before those release updates, updating and migrating them is not possible as the base image itself does not support migration.
- Migration to Intel X9 shape is supported for instances using the base image with 21.8.0.0, 19.17.0.0, and later release updates only. For instances created before those release updates, updating and migrating them is not possible as the base image itself does not support migration.

If you must create a DB system with an older database version, see Critical Patch Updates for information about known security issues with your chosen database version. You must also analyze and patch known security issues for the operating system included with the older database version. For information about security best practices for databases in OCI, see Securing Databases.

How various configurations affect the usable storage

The DB systems use OCI block storage. The following table shows details of the available storage options. Total storage includes available storage plus recovery logs.

General Information

- You can scale your data storage and recovery storage separately. Oracle recommends keeping recovery storage at 20% of total storage or higher.
- For multi-node RAC DB systems, storage capacity is shared between the nodes.
- The recovery area storage is determined based on the storage selected. However, you can change the recovery area storage independently after provisioning.
- The total storage is the sum of available data storage, recovery area storage, and 200 GB reserved for the DB system software.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Available data storage for flexible shapes

Table 1-4 Available data storage for flexible shapes

Available data storage (GB)	Recovery area storage (GB)	Total storage (GB)
256	256	712
512	256	968
1024	512	1736
2048	512	2760
4096	1024	5320
8192	2048	10440
12288	4096	16584
16384	4096	20680
24576	8192	32968
32768	8192	41160
40960	10240	51400
49152	12288	61640
57344	14336	71880
65536	16384	82120
73728	18432	92360
81920	20480	102600

Available data storage for standard shapes

Table 1-5 Available data storage for standard shapes

Available data storage (GB)	Recovery area storage (GB)	Total storage (GB)
256	256	712
512	256	968
1024	256	1480
2048	408	2656
4096	820	5116
6144	1228	7572
8192	1640	10032
10240	2048	12488
12288	2456	14944
14336	2868	17404
16384	3276	19860
18432	3688	22320
20480	4096	24776
22528	4504	27232
24576	4916	29692
26624	5324	32148
28672	5736	34608
30720	6144	37064
32768	6552	39520

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 1-5 (Cont.) Available data storage for standard shapes

Available data storage (GB)	Recovery area storage (GB)	Total storage (GB)
34816	6964	41980
36864	7372	44436
38912	7784	46896
40960	8192	49352

Service Limits

The following limits apply to the Base Database resources.

Table 1-6 Service limits

Resources	Oracle Universal Credits	Pay As You Go or Trial
Total VM DB Block Storage	150TB	2TB
VM.Standard1 -Total OCPUs	300 cores	2 cores
VM.Standard2 -Total OCPUs	300 cores for US West (Phoenix), 300 cores for US East (Ashburn), 50 cores for Germany Central (Frankfurt), 50 cores for UK South (London)	2 cores

> **Note**
>
> Total VM DB Block Storage includes block storage for all VM.Standard1 and VM.Standard2 virtual machine databases.

Fast Provisioning Option

For single-node DB systems, OCI provides a "fast provisioning" option that enables you to create a DB system using Logical Volume Manager (LVM) as your storage management software. The standard way ("standard provisioning") is to provision with Automatic Storage Management (ASM).

The following details apply to the fast provisioning option:

- When using the fast provisioning option, the number and size of the block volumes specified during provisioning determines the maximum total storage available through scaling.
- Multi-node RAC DB systems require ASM and cannot be created using the fast provisioning option.
- You can clone DB systems that have been created using the fast provisioning option.
- You cannot use a custom database software image when provisioning a DB system with LVM.

For more information, see:

- Logical Volume Manager

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Oracle Automatic Storage Management
- Oracle Database Software Images

Storage Scaling Considerations While Using Fast Provisioning

> **Note**
>
> This topic applies only to single-node DB systems.

When you provision a DB system using the fast provisioning option, the **Available storage (GB)** value you specify during provisioning determines the maximum total storage available through scaling. The following table details the maximum storage value available through scaling for each setting offered in the provisioning workflow:

Table 1-7 Storage Scaling Considerations While Using Fast Provisioning

Initial storage specified during provisioning (GB)	Maximum storage available through scaling (GB)
256	2560
512	2560
1024	5120
2048	10240
4096	20480
8192	40960

Fault Domain Considerations for Multi-Node RAC DB Systems

When you provision a multi-node RAC DB systems, the system assigns each node to a different fault domain by default. Using the **Advanced options** link in the provisioning dialog, you can select the fault domain(s) to be used for your multi-node RAC DB systems and the system will assign the nodes to your selected fault domains. Oracle recommends that you place each node of a multi-node RAC DB system in a different fault domain.

For more information on fault domains, see Regions and Availability Domains.

Reboot a DB System Node for Planned Maintenance

The DB system nodes use underlying physical hosts that periodically must undergo maintenance. When such maintenance is required, OCI schedules a reboot of your DB system node and notifies you of the upcoming reboot. The reboot enables your DB system node to be migrated to a new physical host that is not in need of maintenance. (Stopping and starting the node will also result in the migration to a new physical host.) The only effect to your DB system node is the reboot itself. The planned maintenance of the original physical hardware takes place after your node has been migrated to its new host and has no effect on your DB system.

If your DB system node is scheduled for a maintenance reboot, you can proactively reboot your node (by stopping and starting it) using the Console or the API. This lets you control how and when your node experiences downtime. If you choose not to reboot before the scheduled time, then OCI will reboot and migrate your node at the scheduled time.

Copyright © 2022, 2026, Oracle and/or its affiliates.

To identify the DB system nodes that you can proactively reboot, navigate to your system's **DB System Details** page in the Console and check the **Node maintenance reboot** field. If the instance has a maintenance reboot scheduled and can be proactively rebooted, this field displays the date and start time for the reboot. When the **Maintenance reboot** field does not display a date, your DB system has no scheduled node maintenance events.

To check for scheduled maintenance events using the API, use the `GetDbNode` operation to check the `timeMaintenanceWindowEnd` field of the `DbNode` resource. This field specifies when the system will begin the next scheduled node reboot.

To locate nodes that have scheduled maintenance reboots, you can use the Search Service with a predefined query to find all DB systems that have a scheduled maintenance reboot.

For instructions about using the Console to reboot a node, see *Reboot a DB System*.

Security Hardening Tool for DB systems

The DB systems provisioned using Oracle Linux 7 include a Python script, referred to as the Security Technical Implementation Guide (STIG) tool, that you can use to perform security hardening for your DB system.

For more information, see:

- *Security Technical Implementation Guide (STIG) Tool for the DB System*.
- *Enable FIPS, SE Linux, and STIG on the DB System Components*.

Boot Volume Backups

Oracle maintains a weekly boot volume backup of your DB system so that the system can be easily restored in the event of a serious error or system failure. Boot volume backups are currently not accessible to users (there is no Console, API, or CLI access to a DB system boot volume backup), and Oracle bears the cost of keeping and maintaining the backup. In the event of a system failure, contact My Oracle Support to request that Oracle perform a restore of your DB system from the boot volume backup.

Copyright © 2022, 2026, Oracle and/or its affiliates.

2

Configure the Network

VCN and Subnets

This article describes how to manage virtual cloud networks (VCNs) and the subnets in them.

Overview

A VCN is a software-defined network that you set up in the Oracle Cloud Infrastructure (OCI) data centers in a particular region. A subnet is a subdivision of a VCN.

Before you set up a DB system, you must set up a VCN and other Networking service components.

DB systems support subnets configured with either IPv4-only or IPv4/IPv6 dual stack addresses.

To launch a DB system, you must have:

- A VCN in the region where you want the DB system.
- At least one subnet in the VCN (either a public subnet or a private subnet).
- Connectivity to the Oracle Services Network.
- Custom route table with appropriate rules.
- Security rules.

> **Note**
>
> Oracle recommends using the internet gateway for public subnets and service gateway for private subnets along with the appropriate security list and routing table rules.

You could use availability domain specific subnets or regional subnets which span all availability domains in the region.

> **Note**
>
> - Oracle recommends using regional subnets, which span all availability domains in the region.
> - Certain details of the VCN and subnet configuration depend on your choice for DNS resolution within the VCN.

For more information on:

- VCN and subnets, see Overview of VCNs and Subnets.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Networking, see Networking Overview.
- DNS, see *DNS for the DB System*.

IPv4/IPv6 Dual Stack Network Support

You can now provision DB systems with IPv4/IPv6 dual stack networking. This enables applications to scale seamlessly by leveraging the expanded IPv6 address space, addressing IPv4 exhaustion concerns, and ensuring readiness for future growth without network limitations.

Base Database Service supports GUA, BYOIP, and ULA IPv6 prefixes. While provisioning a DB system, a subnet should have only one IPv6 prefix. Base Database Service does not support subnets with multiple IPv6 prefixes. For more information, see IPv6 Addresses.

Virtual IPs (VIPs) and Single Client Access Names (SCANs) will be configured for IPv4 and IPv6 networks. The Oracle Clusterware private interconnect continues to use IPv4 only. For more information, see Requirements for IP Address Space.

The following are supported for this release:

- Provision a new DB system on the subnet configured with dual stack addresses.
- Clone a DB system from an IPv4-only network configuration to a dual stack network configuration and vice versa.
- Data Guard association can exist between DB systems in IPv4-only networks and dual stack networks. Oracle recommends having open traffic for both IPv4 and IPv6 for Data Guard peers.
- A DB system in an IPv4-only network can be migrated to a dual stack network using Oracle Data Guard. For detailed steps, see *Migrate to a Dual Stack Network*.
- Dual stack networking for single-node DB systems with Grid Infrastructure is available for Oracle Database versions 26ai and 19c starting with the 23.8.0.25.04, 19.27.0.0, and later release updates (RU) only.
- For a DB system in an IPv6 network, the Grid Infrastructure can be upgraded to version 23.8.0.25.04 or later only.

Private Subnet with Service Gateway

You can connect to the Oracle Services Network using the service gateway for private subnets. The subnet is private and cannot be reached from the internet. Oracle recommends this option for a production system. The following image provides the architecture for private subnet with service gateway.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Figure 2-1 Architecture for Private Subnet with Service Gateway

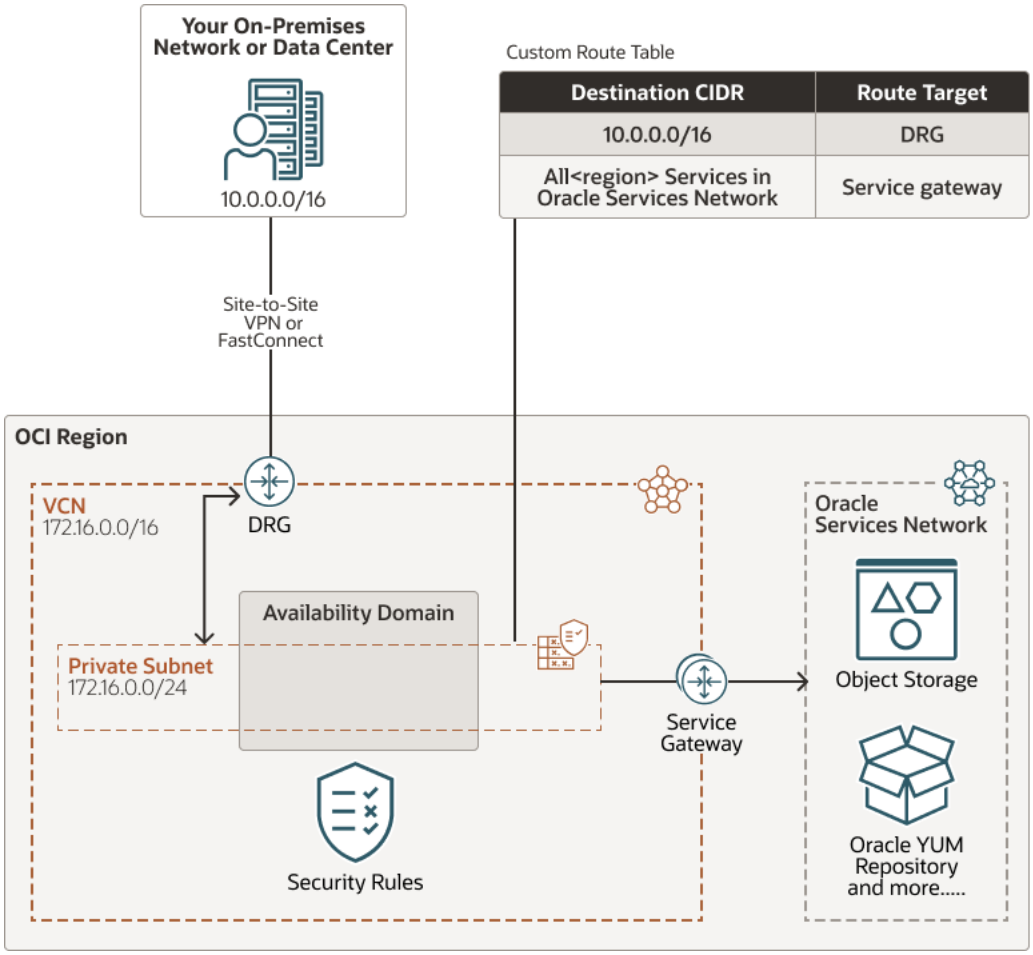

Perform the following steps to setup private subnet with service gateway.

- Private subnet.
- Gateways for the VCN:
  - Dynamic Routing Gateway (DRG), with a FastConnect or Site-to-Site VPN to your on-premises network.
  - Service gateway to reach Oracle Services Network for database provisioning, backups and patching, and to reach Oracle YUM repos for OS updates.
- Route table: A custom route table for the subnet, with these rules:
  - A route for the on-premises network's CIDR, and **`target = DRG`**.
  - A rule for the CIDR label called **All *<region>* Services in Oracle Services Network**, and **`target = the service gateway`**.
- Security rules to enable the desired traffic to and from the DB system nodes.
  - The following rule enables the DB system to communicate with the Oracle services (for public subnet with internet gateway), or with the Oracle Services Network, which includes all the Oracle services (for private subnet with service gateway). It is redundant with the general egress rule for basic connectivity (and in the default

Copyright © 2022, 2026, Oracle and/or its affiliates.

security list). It is optional but recommended in case the general rule (or default security list) is inadvertently changed.

* **Stateless**: No (all rules must be stateful)
* **Destination Type**: Service
* **Destination Service**:
  * When using public subnet (with internet gateway), use the CIDR **`0.0.0.0/0`**
  * When using private subnet (with service gateway), use the CIDR label called **All *<region>* Services in Oracle Services Network**
* **IP Protocol**: TCP
* **Source Port Range**: All
* **Destination Port Range**: 443 (HTTPS)
* **Description**: An optional description of the rule.

For more information on:

- private subnet, see Private Subnet.
- service gateway, see Access to Oracle Services: Service Gateway.
- Dynamic Routing Gateway, see Dynamic Routing Gateway (DRG).
- FastConnect, see FastConnect.
- Site-to-Site VPN, see Site-to-Site VPN.
- route table, see Route Table.
- security rules, see Security Rules and *Security Rules for the DB System*.
- networking, see Networking Overview.

Public Subnet with Internet Gateway

You can connect to the Oracle Services Network using the internet gateway for public subnets. You can use this setup in production if you want to use an internet gateway with the VCN, or if you have services that run only on a public network and need access to the database. This option can be useful when doing a proof-of-concept or development work. The following image provides the architecture for public subnet with internet gateway.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Figure 2-2 Architecture for Public Subnet with Internet Gateway

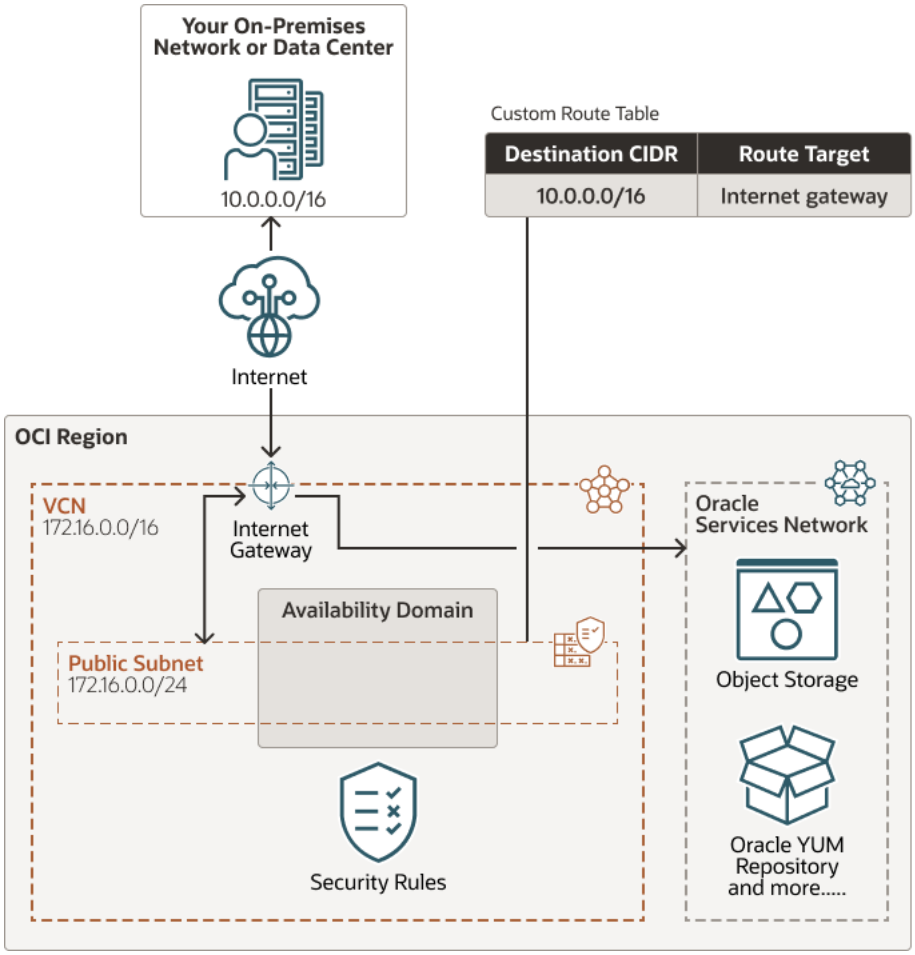

Perform the following steps to setup public subnet with internet gateway.

- Public subnet.
- Internet gateway.
- Route table: A custom route table for the subnet, with a rule for CIDR **`0.0.0.0/0`**, and **`target = internet gateway`**.
- Security rules to enable the desired traffic to and from the DB system nodes.

For more information on:

- public subnet, see Public Subnet.
- internet gateway, see Internet Gateway.
- route table, see Route Table.
- security rules, see Security Rules and *Security Rules for the DB System*.

> **Note**
>
> See this known issue for information about configuring route rules with *service gateway* as the target on route tables associated with public subnets.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Requirements for IP Address Space

If you are setting up DB systems (and thus VCNs) in more than one region, make sure the IP address space of the VCNs does not overlap.

The subnet you create for a DB system cannot overlap with **`192.168.16.16/28`**, which is used by the Oracle Clusterware private interconnect on the database instance.

> **Warning**
>
> Any change in the VCN may impact RAC DB functionality. Oracle recommends that you assess the changes required on the CRS side before making any changes in the VCN.
>
> For more information, see How to Modify Public Network Information including VIP in Oracle Clusterware (Doc ID 276434.1).

> **Note**
>
> While the subnet itself uses **`192.168.16.0/24`** as default, the actual addresses the Oracle Clusterware private interconnect uses is **`192.168.16.16/28`**. Effectively, even though you cannot deploy a DB system using **`192.168.16.0/24`**, the private interconnect will be able to communicate with hosts using these addresses if you modify the private interconnect subnet to **`192.168.16.16/28`**.

The following table lists the *minimum* required subnet size.

> **Note**
>
> The Networking service reserves three IP addresses in each subnet. Allocating a larger space for the subnet than the minimum required (for example, at least /25 instead of /28) can reduce the relative impact of those reserved addresses on the subnet's available space.
>
> For more information, see IP Addresses Reserved for Use by Oracle.

DB System Type	# Required IP Addresses	Minimum Subnet Size
1-node virtual machine	1 + 3 reserved in subnet = **4**	/30 (4 IP addresses)
2-node RAC virtual machine	(2 addresses * 2 nodes) + 3 for SCANs + 3 reserved in subnet = **10**	/28 (16 IP addresses)

Copyright © 2022, 2026, Oracle and/or its affiliates.

VCN Creation Wizard

> **Note**
>
> Oracle recommends not to use this VCN creation wizard for production.

The Networking section of the Console includes a wizard that creates a VCN along with related resources. It can be useful if you just want to try launching an instance. However, the wizard automatically creates a public subnet and an internet gateway. You may not want this for your production network, so Oracle recommends you create the VCN and other resources individually yourself instead of using the wizard.

For more information on the wizard, see Virtual Networking Quickstart.

DNS for the DB System

You can use DNS and hostname resolution for the DB system.

Oracle recommends using a private DNS resolver to enable the use of hostnames when on-premises hosts and VCN resources communicate with each other.

The following table shows which choices are supported with each type of DB system, and the endpoints that need to be resolved for the DB system to function.

DB System Type	Supported DNS Choices	Endpoints to Be Resolved
Single-node virtual machine	• Recommended: Default (Internet and VCN resolver) • Custom DNS resolver of your choice	• Object Storage endpoints (includes both the Object Storage endpoints and Swift endpoints) • Oracle YUM repo endpoints
Multi-node RAC virtual machine	• Default (Internet and VCN resolver)	• Object Storage endpoints (includes both the Object Storage endpoints and Swift endpoints) • Oracle YUM repo endpoints • Single Client Access Names (SCANs)

The following sections give more details about the DNS choices.

Default (Internet and VCN Resolver)

See the preceding table for the types of DB systems that support the Internet and VCN resolver.

Oracle recommends using the Internet and VCN resolver for DNS. It's the default, built-in DNS functionality that comes with each VCN. It enables hosts in a VCN to resolve these items:

- Hostnames of other hosts in the same VCN.
- Hostnames that are publicly published on the Internet.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For a DB system, the Internet and VCN resolver handles resolution of all necessary endpoints: Object Storage endpoints (includes both the Object Storage endpoints and Swift endpoints), YUM repos, and SCANs (SCANs are used only with multi-node RAC DB systems).

By default, each VCN is configured to use the Internet and VCN resolver. If you plan to use a custom DNS resolver, you must configure the VCN in a different way.

For more information, see:

- Private DNS resolvers
- Use private DNS in interconnected VCNs and On-premises
- Database Connection Strings
- DNS in Your Virtual Cloud Network
- Use the Internet and VCN Resolver With Your DB System
- Use a Custom DNS Resolver With Your DB System

Use the Internet and VCN Resolver With Your DB System

As part of the overall network setup, perform these tasks:

1. Create the VCN with the required DNS settings:
   - When creating the VCN, select the check box for **Use DNS hostnames in this VCN**.
   - Specify a DNS label for the VCN.
   - Notice that you cannot change these VCN DNS settings after you create the VCN.
2. Create each subnet with the required DNS settings:
   - When creating a subnet in the VCN, select the check box for **Use DNS hostnames in this subnet**.
   - Specify a DNS label for the subnet.
   - Notice that you cannot change these subnet DNS settings after you create the subnet.
3. Use the default set of DHCP options that come with the VCN:
   - When creating each subnet, configure it to use the VCN's default set of DHCP options.
   - By default, the default set of DHCP options is configured to use the Internet and VCN resolver.
4. Create the DB system with a hostname prefix:
   - Later, when creating the DB system, specify a value in the **Hostname prefix** field.
   - Notice that the DB system's **Host domain name** value is automatically assigned based on the VCN and subnet DNS labels.

The resulting DB system has a fully qualified domain name (FQDN) based on the hostname prefix, VCN label, and subnet label you specify.

For more information, see:

- Overview of VCNs and Subnets
- DHCP Options

Copyright © 2022, 2026, Oracle and/or its affiliates.

Hostname Restrictions for Using the Internet and VCN Resolver

When you create the VCN, subnet, and DB system, you must carefully set the following identifiers, which are related to DNS in the VCN:

- VCN DNS label.
- Subnet DNS label.
- Hostname prefix for the DB system.

These values make up the node's FQDN:

```
<hostname_prefix><RAC_node_#>.<subnet_DNS_label>.<VCN_DNS_label>.oraclevcn.com
```

For multi-node RAC DB systems, a node number is automatically appended after the hostname prefix.

For example:

- Node 1: `dbsys1.ad1.acmevcniad.oraclevcn.com`
- Node 2: `dbsys2.ad1.acmevcniad.oraclevcn.com`

Requirement for the DB System's Hostname Prefix:

- Recommended maximum: 16 characters.
- Must start with an alphabetical character.
- Cannot be the string *localhost*.

Requirements for the VCN and Subnet DNS Labels:

- Recommended maximum: 15 characters.
- No hyphens or underscores.
- Recommended: Include the region name in the VCN's name, and include the availability domain name in the subnet's name.
- The FQDN has a maximum total limit of 63 characters, so set the VCN and subnet DNS labels short enough to meet that requirement. Here is a safe general rule:

  ```
  <16_chars_max>#.<15_chars_max>.<15_chars_max>.oraclevcn.com
  ```

> **Note**
>
> The recommended maximums are not enforced when you create the VCN and subnets. However, the DB system deployment fails if the FQDN has more than 63 characters.

Custom DNS Resolver

A custom DNS resolver is a DNS server that you set up in your on-premises network and maintain yourself. It must resolve the endpoints required by the DB system.

Copyright © 2022, 2026, Oracle and/or its affiliates.

By default, the VCN is configured to use the Internet and VCN resolver. Therefore, if you instead want to use a custom DNS resolver, you must configure the VCN and DHCP options in a different way.

Use a Custom DNS Resolver With Your DB System

As part of the overall network setup, perform these tasks:

1. Create the VCN with the recommended DNS settings:
   - When creating the VCN, Oracle recommends that you select the check box for **Use DNS hostnames in this VCN** and then specify a DNS label for the VCN.
   - Notice that you cannot change the preceding VCN DNS settings after you create the VCN. They are optional for a custom DNS server, but required if you use the Internet and VCN resolver. Therefore, Oracle recommends that you configure them now in case you later want to use the Internet and VCN resolver.
2. Create each subnet with the recommended DNS settings:
   - When creating a subnet in the VCN, Oracle recommends that you select select the check box for **Use DNS hostnames in this subnet** and then specify a DNS label for the subnet.
   - Notice that you cannot change the preceding subnet DNS settings after you create the subnet. They are optional for a custom DNS server, but required if you use the Internet and VCN resolver. Therefore, Oracle recommends that you configure them now in case you later want to use the Internet and VCN resolver.
3. Edit the default set of DHCP options to use a custom resolver:
   - When creating each subnet, configure it to use the VCN's default set of DHCP options.
   - Edit the default set of DHCP options so that **DNS type** is set to **Custom resolver**. Provide the IP address for at least one DNS server (maximum three). Optionally provide a single search domain (which will automatically be added to the host's `/etc/resolv.conf` file).
4. Create the DB system with required DNS entries:
   - Later, when creating the DB system, specify a **Hostname prefix**.
   - For the **Host domain name**: If you selected the check box for **Use DNS hostnames** in the preceding steps, the **Host domain name** is automatically generated from the VCN and subnet DNS labels. Otherwise, you must provide a value for the **Host domain name**.
   - Notice that when launching the DB system, an IP address is automatically assigned from the VCN's CIDR block and the address is resolved locally based on the host's `/etc/hosts` file. Your custom DNS resolver does not need to resolve the hostname in advance for the DB system launch to succeed.

For more information, see:

- Overview of VCNs and Subnets
- DHCP Options

Hostname Restrictions When Using a Custom DNS Resolver

Requirement for the DB System's Hostname Prefix:

- Recommended maximum: 16 characters.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Must start with an alphabetical character.
- Cannot be the string *localhost*.

Requirements for the VCN and Subnet DNS Labels:

- You can provide a value for the DNS labels only if you select the check box for **Use DNS hostnames** when creating the VCN and subnets. The resulting FQDN for the DB system follows this format:

  ```
  <hostname_prefix>.<subnet_DNS_label>.<VCN_DNS_label>.oraclevcn.com
  ```

- Recommended maximum for each DNS label: 15 characters.
- No hyphens or underscores.
- Recommended: Include the region name in the VCN's name, and include the availability domain name in the subnet's name.
- The FQDN has a maximum total limit of 63 characters, so set the VCN and subnet DNS labels short enough to meet that requirement. Here is a safe general rule:

  ```
  <16_chars_max>.<15_chars_max>.<15_chars_max>.oraclevcn.com
  ```

> **Note**
>
> The recommended maximums are not enforced when you create the VCN and subnets. However, the DB system deployment fails if the FQDN has more than 63 characters.

Requirements for the DB System's Host Domain Name:

- You can provide a value in the **Host domain name** field only if you did not select the check box for **Use DNS hostnames** when creating the VCN and subnets.
- No hyphens or underscores.
- Ensure that the value results in an FQDN that is no longer than 63 characters. Otherwise the DB system deployment will fail.

DNS: Between On-Premises Network and VCN

If you are using the Internet and VCN resolver and want to enable the use of hostnames when on-premises hosts and VCN resources communicate with each other, you can set up an instance in the VCN to be a custom DNS server.

For an example of an implementation of this scenario with the Oracle Terraform provider, see Hybrid DNS Configuration.

Set Up DNS for a DB System

DNS lets you use host names instead of IP addresses to communicate with a DB system.

You can use the Internet and VCN resolver (the DNS capability built into the VCN) as described in DNS in Your Virtual Cloud Network.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Alternatively, you can use your choice of DNS server. You associate the host name and domain name to the public or private IP address of the DB system. You can find the host and domain names and IP addresses for the DB system on the **Database** page in the Console.

To associate the host name to the DB system's public or private IP address, contact your DNS administrator and request a custom DNS record for the DB system's IP address. For example, if your domain is `example.com` and you want to use `clouddb1` as the host name, you would request a DNS record that associates `clouddb1.example.com` to your DB system's IP address.

If you provide the public IP address to your DNS administrator as described above, you should also associate a custom domain name to the DB system's public IP address:

1. Register your domain name through a third-party domain registration vendor, such as register.com.
2. Resolve your domain name to the DB system's public IP address, using the third-party domain registration vendor console. For more information, refer to the third-party domain registration documentation.

Security Rules for the DB System

This article lists the security rules to use with your DB system. Security rules control the types of traffic allowed in and out of the DB system's compute nodes. The rules are pided into two sections.

For more information about security rules, see Security Rules. For more information about different ways to implement these rules, see Ways to Implement the Security Rules.

> **Note**
>
> Your instances running Oracle-provided DB system images also have firewall rules that control access to the instance. Make sure that both the instance's security rules and firewall rules are set correctly. Also see *Open Ports on the DB System*.

General Rules Required for Basic Connectivity

The following sections has several general rules that enable essential connectivity for hosts in the VCN.

If you use security lists to implement your security rules, be aware that the rules that follow are included by default in the default security list. Update or replace the list to meet your particular security needs. The two ICMP rules (general ingress rules 2 and 3) are required for proper functioning of network traffic within the Oracle Cloud Infrastructure environment. Adjust the general ingress rule 1 (the SSH rule) and the general egress rule 1 to allow traffic only to and from hosts that require communication with resources in your VCN.

For more information on default security list, see Security Lists.

General Ingress Rule 1: Allows SSH Traffic From Anywhere

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0 (IPv4), ::/0 (IPv6)

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 22

> **Note**
>
> IPv6 CIDR is required only if you want to use IPv6 address to connect to SSH.

General Ingress Rule 2: Allows Path MTU Discovery Fragmentation Messages

This rule enables hosts in the VCN to receive Path MTU Discovery fragmentation messages. Without access to these messages, hosts in the VCN can have problems communicating with hosts outside the VCN.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0 (IPv4), ::/0 (IPv6)
- **IP Protocol:** ICMP
- **Type:** 3
- **Code:** 4

General Ingress Rule 3: Allows Connectivity Error Messages Within the VCN

This rule enables the hosts in the VCN to receive connectivity error messages from each other.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** Your VCN's CIDR
- **IP Protocol:** ICMP
- **Type:** All
- **Code:** All

General Egress Rule 1: Allows All Egress Traffic

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** 0.0.0.0/0 (IPv4), ::/0 (IPv6)
- **IP Protocol:** All

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> - IPv6 destination CIDR is required only for outgoing communication to IPv6 networks.
> - Destination CIDR can be restricted.

Custom Security Rules

The following rules are necessary for the DB system's functionality.

> **Note**
>
> Custom ingress rules 1 and 2 only cover connections initiated from within the VCN. If you have a client that resides *outside the VCN*, Oracle recommends setting up two *additional* similar rules that instead have the **Source CIDR** set to the public IP address of the client.

Custom Ingress Rule 1: Allows ONS and FAN Traffic From Within the VCN

This rule is recommended and enables the Oracle Notification Services (ONS) to communicate about Fast Application Notification (FAN) events.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** Your IPv4 and IPv6 VCN's CIDR
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 6200
- **Description:** An optional description of the rule.

Custom Ingress Rule 2: Allows SQL*NET Traffic From Within the VCN

This rule is for SQL*NET traffic and is required only if you need to enable client connections to the database.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** Your IPv4 and IPv6 VCN's CIDR
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 1521
- **Description:** An optional description of the rule.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Custom Egress Rule 1: Allows Outbound SSH Access

This rule enables SSH access between nodes in a 2-node DB system. It is redundant with the general egress rule in General Rules Required for Basic Connectivity (and in the default security list). It is optional but recommended in case the general rule (or default security list) is inadvertently changed.

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** 0.0.0.0/0 (IPv4), ::/0 (IPv6)
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 22
- **Description:** An optional description of the rule.

Custom Egress Rule 2: Allows Access To Oracle Services Network

This rule enables the DB system to communicate with the Oracle services (for public subnet with internet gateway), or with the Oracle Services Network, which includes all the Oracle services (for private subnet with service gateway). It is redundant with the general egress rule in General Rules Required for Basic Connectivity (and in the default security list). It is optional but recommended in case the general rule (or default security list) is inadvertently changed. OCI services communicate with IPv4 only.

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** Service
- **Destination Service:**
  - When using IPv4 public subnet (with internet gateway), use the CIDR `0.0.0.0/0`
  - When using IPv4 private subnet (with service gateway), use the CIDR label called **All *<region>* Services in Oracle Services Network**
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 443 (HTTPS)
- **Description:** An optional description of the rule.

For more information about networking, see Networking Overview.

Ways to Implement the Security Rules

The Networking service offers two ways to implement security rules within your VCN:

- Network security groups
- Security Lists

For a comparison of Security Lists and Network Security Groups, see Security Rules.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Use Network Security Groups

If you choose to use network security groups (NSGs), here is the recommended process:

1. Create a network security group for DB systems. Add the following security rules to that NSG:
   - The rules listed in *General Rules Required for Basic Connectivity*.
   - The rules listed in *Custom Security Rules*.
2. When the database administrator creates the DB system, they must choose several networking components (for example, which VCN and subnet to use). They can also choose which NSG or NSGs to use. Make sure they choose the NSG you created.

You could instead create one NSG for the general rules and a separate NSG for the custom rules. Then when the database administrator chooses which NSGs to use for the DB system, make sure they choose both NSGs.

Use Security Lists

If you choose to use security lists, here is the recommended process:

1. Configure the subnet to use the required security rules:
   a. Create a custom security list for the subnet and add the rules listed in *Custom Security Rules*.
   b. Associate the following two security lists with the subnet:
      - VCN's default security list with all its default rules. This automatically comes with the VCN.
      - The new custom security list you created for the subnet
2. Later when the database administrator creates the DB system, they must choose several networking components. When they select the subnet that you have already created and configured, the security rules are automatically enforced for the compute nodes created in the subnet.

> ⚠ **Caution**
>
> **Do not remove the default egress rule from the default security list**. If you do, instead make sure to include the following replacement egress rule in the subnet's custom security list:
>
> - **Stateless:** No (all rules must be stateful)
> - **Destination Type:** CIDR
> - **Destination CIDR:** 0.0.0.0/0
> - **IP Protocol:** All

Copyright © 2022, 2026, Oracle and/or its affiliates.

Manage Network Security Groups for a DB System

Your DB system can use up to five network security groups (NSGs). Note that if you choose a subnet with a security list, the security rules for the DB system will be a union of the rules in the security list and the NSGs.

Procedure

Perform the following steps to edit the network security group.

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **DB system information** tab, under the **Network** section, select the **Edit** link in **Network security groups**.
3. In the **Edit network security groups** panel, select **+ Another network security group** to add an NSG to the DB system.
   - To change an assigned NSG, select the drop-down menu displaying the NSG name, then select a different NSG.
   - To remove an NSG from your DB system, select the **X** icon to the right of the displayed NSG name.
4. Select **Save**.

For more information, see Security Lists, Network Security Groups, and *VCN and Subnets*.

Update the Security List for the DB System

You can update the security list for the DB systems using the following steps.

Review the list of ports in *Open Ports on the DB System* and for every port you open in iptables, update the security list used for the DB system, or create a new security list.

> **Note**
>
> The port 1521 for the Oracle default listener is included in iptables, but should also be added to the security list.

Procedure

Perform the following steps to update the security list for the DB system.

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **DB system information** tab, under the **Network** section, note the **Client Subnet** name and select the **VCN** link.
3. On the Virtual Cloud Network details page, on the **Subnets** tab, select the required VCN.
4. On the Subnet details page, on the **Security** tab, select the required security list.
5. On the Security List details page, select the **Security rules** tab.
   - **Ingress Rules** is the first table of security rules on the page. **Egress Rules** is the second table of security rules on the page.

Copyright © 2022, 2026, Oracle and/or its affiliates.

6. To add a rule, select **Add Ingress Rules** (or **Add Egress Rules**).
7. To delete an existing rule, select the checkbox next to the rule and then select **Actions** and then select **Remove**.
8. To edit an existing rule, select the checkbox next to the rule and then select **Actions** and then select **Edit**.

> **Tip**
>
> The source CIDR should be the CIDR block that includes the ports you open for the client connection.

For more information about creating or updating a security list, see Security Lists.

Copyright © 2022, 2026, Oracle and/or its affiliates.

3

Create

Overview of Creating a DB System

This article provides an introduction to the various settings required to create DB systems.

When you create a DB system using the Console, the API, or the CLI, a system is provisioned to support Oracle Database, and an Oracle Database is created based on the options you provide and some default options described later in this article.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be granted security access in a **policy** by an administrator. This access is required whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you get a message that you don’t have permission or are unauthorized, verify with your administrator what type of access you have and which **compartment** to work in.

For administrators: The policy in Let database admins manage Oracle Cloud database systems lets the specified group do everything with databases and related Database resources.

If you're new to policies, see Getting Started with Policies and Common Policies. If you want to dig deeper into writing policies for databases, see Details for the Database Service.

Prerequisites

You'll need the following items to create a DB system:

- The public key, in OpenSSH format, from the key pair that you plan to use for connecting to the DB system via SSH. A sample public key, abbreviated for readability, is shown below.

  ```
  ssh-rsa AAAAB3NzaC1yc2EAAAABJQAA....lo/gKMLVM2xzc1xJr/
  Hc26biw3TXWGEakrK1OQ== rsa-key-20160304
  ```

- A correctly configured virtual cloud network (VCN) to launch the DB system. Its related networking resources (gateways, route tables, security lists, DNS, and so on) must also be configured as necessary.
- Oracle recommends using a service gateway to enable necessary access, if you plan to back up your DB system or use the managed update feature.
- For a multi-node RAC DB system, ensure that port 22 is open for both ingress and egress on the subnet, and that the security rules you create are stateful (the default). Otherwise, the DB system might fail to provision successfully.

For more information, see:

- Managing Key Pairs on Linux Instances
- *VCN and Subnets*

Copyright © 2022, 2026, Oracle and/or its affiliates.

Default Options for the Database

To simplify creating a DB system in the Console, and when using the API, the following default options are used for the database.

- **Console enabled**: False
- **Create container database**: True
- **Create instance only (for standby and migration)**: False
- **Database home ID**: Creates a new database home
- **Database language**: AMERICAN
- **Database sizing template**: `odb2`
- **Database storage**: Automatic Storage Management (ASM). Optionally, for faster provisioning, single-node DB systems can be provisioned using Logical Volume Manager.
- **Database territory**: AMERICA
- **Database unique name**: The user-specified database name and a system-generated suffix, for example, `dbtst_phx1cs`.
- **PDB admin name**: `pdbuser`

For more information, see:

- *Create a DB System*
- Automatic Storage Management

Use a Backup to Create the Database

When creating a new DB system using a backup stored in the Recovery Service or Object Storage as the source of the database, you have the following options:

- **Daily automatic backup.** Requires that you have automatic backups enabled and an available backup to use. If you are creating a database from an automatic backup, you can choose any level 0 weekly backup, or a level 1 incremental backup created after the most recent level 0 backup.
- **On-demand full backup.**
- **Standalone backup.**
- **Last archived redo log backup.** Requires that you have automatic backups enabled. This backup combines data from the most recent daily automatic backup and data from archived redo logs, and represents the most current backup available. The time of the last archived redo log backup is visible on the **Database Details** page in the **Last backup time** field.
- **Point-in-time out of place restore.** Specify a timestamp to create a new copy of the database that included data up to a specified point in time. The timestamp must be earlier or equal to the **Last backup time** time displayed on the **Database Details** page. Note the following limitations when performing a point-in-time out of place restore:
  - The timestamp must be within the recovery window of the database.
  - The timestamp must be available within the database incarnation of the available automatic backups.
  - The timestamp cannot fall within two overlapping database incarnations.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- The create database operation will fail if the database has undergone structural changes since the specified timestamp. Structural changes include operations such as creating or dropping a tablespace.
- The create database operation cannot be started if another point-in-time database copy operation is in progress.

For more information, see *Back Up a Database Using the Console*.

Custom IP Addresses for the DB Systems

When creating a new single-node DB system or cloning an existing DB system, you can optionally define the IP address of the DB system being provisioned. This is useful in development contexts where you create and delete the same DB system over and over, and you need each new iteration of the DB system to use the same IP address.

> **Note**
>
> This facility is not available when creating a multi-node RAC DB system.

Easy Create a DB System

This article provides the details and procedure to easy create a DB system with predefined configuration and minimal required input.

Easy Create configuration

Easy Create provisions a DB system with the following predefined configuration:

- **Operating System**: Oracle Linux 8 (OL8)
- **Database version**: Oracle Database 19c (19.0.0.0)
- **License type**: License included
- **Shape**: VM.Standard.E5.Flex shape with 4 OCPU, 64 GB of memory, 4 Gbps of network bandwidth, and 64K IOPS
  For Enterprise Edition - Developer: VM.Standard.A1.Flex with 8 OCPU, 64 GB of memory, 8 Gbps of network bandwidth, and 128K IOPS
- **Storage**: 256 GB of available storage and 19.2K IOPS
  For Enterprise Edition - Developer: 50 GB of available storage, and 19.2K IOPS
- **Storage management software**: Logical Volume Manager
- **Block volume performance**: Higher performance

Procedure

Perform the following steps to easy create a DB system using the OCI Console:

On the **DB Systems** list page, select **Easy Create DB system**. If you need help finding the list page or the DB system, see *List the DB Systems*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Basic details

- **Database name**: Enter a name for the database. It is also known as the `DB_NAME`. The database name must begin with an alphabetic character and can contain a maximum of eight alphanumeric characters. Special characters are not permitted.
- **Oracle Database software edition**: The database edition supported by the DB system. You cannot change the edition after provisioning. The following editions are supported:
  - Standard Edition
  - Enterprise Edition
  - Enterprise Edition - High Performance
  - Enterprise Edition - Extreme Performance
  - Enterprise Edition - Developer

Administrator credentials

- **Username**: `sys` (This is a read-only field). A database administrator named `sys` will be created with the password you supply.
- **Password**: Enter a password for the administrator. The password must meet the following criteria:
  - A strong password for SYS, SYSTEM, TDE wallet, and PDB administrator.
  - The password must be 9 to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters.
  - The special characters must be _, #, or -.
  - The password must not contain the user name (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reversed order and regardless of casing.
- **Confirm password**: Reenter the password you specified.
- Using a separate TDE wallet password is optional. By default, Easy Create uses the administrator password for the TDE wallet.

  To specify a different TDE wallet password, switch off **Use the administrator password for the TDE wallet** toggle. Enter a password in **Enter TDE wallet password**, and then confirm by reentering it in **Confirm TDE wallet password**.

Network information

If you already have a VCN and subnet, skip creating a VCN.

If you want to create a new VCN with a simplified configuration, select **Create VCN** and provide the following details in the Create default VCN panel.

- **Display name**: Enter a display name for the VCN.
- **Compartment**: Select the compartment in which to create the VCN.
- **CIDR**: Displays the CIDR block for the VCN. The default value is 10.2.0.0/24.
- Select **Create**.

Select the VCN and subnet details.

- Select **Compartment** to select a VCN in a different compartment.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Virtual cloud network**: From the list, select a VCN in which to create the DB system.
- Select **Compartment** to select a subnet in a different compartment.
- **Client subnet**: The subnet to which the DB system attaches. Do not use a subnet that overlaps with 192.168.16.16/28, which is used by the Oracle Clusterware private interconnect on the database instance. Specifying an overlapping subnet causes the private interconnect to malfunction.

  Select a dual stack subnet if you want to configure the DB system with both IPv4 and IPv6 addresses.

Default parameters

Select **View details** to review the predefined configuration of the DB System. The following parameters are provided.

- **DB system name**: The name of the DB system. Easy Create generates the name automatically.
- **Availability domain**: The availability domain that hosts the DB system.
- **Compartment**: The compartment that contains the DB system.
- **License type**: The license type applied to the DB system.
- **Database image**: The database software image version used to create the database.
- **Unified Auditing**: Unified Auditing is not enabled for the DB system.
- **Shape**: The compute shape used for the DB system.
- **Shape details**: Lists the OCPUs, memory, network bandwidth, and IOPS provided by the selected shape.
- **Storage management software**: The software used to manage the DB system storage.
- **Block volume performance**: The performance level configured for the block volumes.
- **Storage details**: Lists the available storage capacity and provisioned IOPS.
- **Character Set**: The database character set. The default value is AL32UTF8.
- **National Character Set**: The national character set. The default value is AL16UTF16.

Create

Select **Create**. The DB system appears in the list with a status of Provisioning. The DB system's icon changes from yellow to green (or red to indicate errors).

After the DB system's icon turns green, with a status of Available, you can select the highlighted DB system name to display details about the DB system. Note the IP addresses. You'll need the private or public IP address, depending on network configuration, to connect to the DB system.

Create a DB System

This article provides the details and procedure to create a DB system.

Copyright © 2022, 2026, Oracle and/or its affiliates.

General Information

The DB systems will be provisioned with Oracle Linux 8 (OL8) for Oracle Database versions 26ai, 21c, and 19c with 23.4.0.24.05, 21.12.0.0, 19.21.0.0, and later release updates (RU).

Prerequisites

- A correctly configured virtual cloud network (VCN) to launch the DB system. Its related networking resources (gateways, route tables, security lists, DNS, and so on) must also be configured as necessary. For more information, see *VCN and Subnets*.
- Oracle recommends using a service gateway to enable necessary access if you plan to back up your DB system or use the managed update feature.
- For a multi-node RAC DB system, ensure that port 22 is open for both ingress and egress on the subnet and that the security rules you create are stateful (the default). Otherwise, the DB system might fail to provision successfully.
- Optionally, a public key, in OpenSSH format, from the key pair that you plan to use for connecting to the DB system via SSH. A sample public key, abbreviated for readability, is shown below. For more information, see Managing Key Pairs on Linux Instances.

  ```
  ssh-rsa AAAAB3NzaC1yc2EAAAABJQAA....lo/gKMLVM2xzc1xJr/
  Hc26biw3TXWGEakrK1OQ== rsa-key-20160304
  ```

Procedure

Perform the following steps to create a DB system using the OCI Console:

On the **DB Systems** list page, select **Create DB system**. If you need help finding the list page or the DB system, see *List the DB Systems*.

Basic details

- **DB system name**: Enter a name for the DB system. Avoid entering confidential information. OCI assigns an unique Oracle Cloud Identifier (OCID) to the DB system, so the name you enter can be nonunique.
- **Compartment**: Select a compartment for your new DB system from the list. By default, the current compartment is selected.
- **Availability domain**: Select an availability domain in which the DB system must reside.
- **Oracle Database software edition**: The database edition supported by the DB system. The database edition cannot be changed later.

Shape

Shape: The shape determines the type of DB system created and the resources allocated to the system. For a complete list of shapes, see Available Shapes and How It Determines the Resources Allocated.

- By default, the **AMD VM.Standard.E4.Flex** shape with **4 OCPU**s is selected.
- The Enterprise Edition Developer is supported only on the Ampere A1 shape.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- If **Enterprise Edition Developer** is selected for the Oracle Database software edition, then by default, the **VM.Standard.A1.Flex** shape with **1 OCPU** is selected.
- The Standard Edition is not supported on the Ampere A1 shape.
- By default, the same shape as the primary is selected for standby.
- To specify a shape other than the default, select **Change shape**.

In the **Change shape** panel, provide the following details to select from available shapes:

Shape series

Select an AMD, Intel, or Ampere processor in the processor group.

- **AMD:** Shapes that use current-generation AMD processors. The AMD shapes are flexible.
- **Intel:** Standard and optimized shapes that use current-generation Intel processors. Both fixed and flexible Intel shapes are available.
- **Ampere**: Shapes that use Arm-based Ampere processors. The Ampere shapes are flexible.

> **Note**
>
> If you select Ampere A1, AMD E4, AMD E5, or Intel X9 flexible shapes, the memory, network bandwidth, and maximum theoretical IOPS scale proportionally.

Configure OCPU

Select the shape you want to use for this instance.

To change the number of OCPUs, perform the following steps:

- From the actions menu of the desired shape, select **Update OCPU count**.
- Select the desired number of OCPUs per node from the list.
- Select **Update**.

You can change the number of OCPUs for Ampere A1, AMD E4, AMD E5, and Intel X9 flexible shapes.

- For the Ampere A1 shape, a minimum of 1 OCPU and a maximum of 57 OCPUs can be selected.
- For AMD E4 and E5 shapes, a minimum of 1 OCPU and a maximum of 64 OCPUs can be selected.
- For the Intel X9 shape, a minimum of 1 OCPU and a maximum of 32 OCPUs can be selected.

The following resources scale proportionately to the number of **OCPUs** you selected.

- **Memory (GB)**: The amount of memory you want to allocate to this instance.
  For Ampere A1, AMD E4, AMD E5, and Intel X9 shapes, the memory will scale proportionally based on the number of OCPUs selected.
  - For the Ampere A1 shape, for each OCPU, 8 GB of memory is allocated. A minimum of 8 GB and a maximum of 456 GB of memory is allocated.
  - For AMD E4 and E5 shapes, for each OCPU, 16 GB of memory is allocated. A minimum of 16 GB and a maximum of 1024 GB of memory is allocated.

Copyright © 2022, 2026, Oracle and/or its affiliates.

  - For the Intel X9 shape, for each OCPU, 16 GB of memory is allocated. A minimum of 16 GB and a maximum of 512 GB of memory is allocated.

- **Network bandwidth (Gbps)**: The amount of network bandwidth you want to allocate to this instance.
  For Ampere A1, AMD E4, AMD E5, and Intel X9 shapes, the bandwidth will scale proportionally based on the number of OCPUs selected. For each OCPU, 1 Gbps of network bandwidth is allocated.
  - For the Ampere A1 shape, a minimum of 1 Gbps and a maximum of 40 Gbps of network bandwidth is allocated.
  - For AMD E4 and E5 shapes, a minimum of 1 Gbps and a maximum of 40 Gbps of network bandwidth is allocated.
  - For the Intel X9 shape, a minimum of 1 Gbps and a maximum of 32 Gbps of network bandwidth is allocated.
- **Theoretical max IOPS**: The amount of input and output per second (IOPS) you want to allocate to this instance. Theoretical max IOPS is also dependent on the storage you select.
  For standard x86 shapes, the theoretical max IOPS will scale proportionally based on the number of ECPUs selected. For each ECPU, 8K theoretical max IOPS is allocated.

  For Ampere A1, AMD E4, AMD E5, and Intel X9 shapes, the theoretical max IOPS will scale proportionally based on the number of OCPUs selected. For each OCPU, 16K theoretical max IOPS is allocated.
  - For the Ampere A1 shape, a minimum of 16K and a maximum of 640K theoretical max IOPS is allocated.
  - For AMD E4 and E5 shapes, a minimum of 16K and a maximum of 640K theoretical max IOPS is allocated.
  - For the Intel X9 shape, a minimum of 16K to a maximum of 512K theoretical max IOPS is allocated.

Storage

- **Storage management software**: The storage determines how much storage to provision and the performance settings. By default, **Logical Volume Manager** with **Higher performance** is selected.
- To specify a storage other than the default, select **Change storage**.

> **Note**
>
> Ampere A1 shape is only supported on Logical Volume Manager. When the Ampere A1 shape is selected, the storage management software type changes to Logical Volume Manager with the Higher Performance option.

In the **Change storage** panel, provide the following details to select from available storage:

Storage management software

For the storage management software type, select one of the following:

- **Oracle Grid Infrastructure** to use Oracle Automatic Storage Management (recommended for production workloads)

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Logical Volume Manager** to quickly provision your DB system using Logical Volume Manager storage management software.

> **Note**
>
> The Available storage (GB) value you specify during provisioning determines the maximum total storage available through scaling. For total storage available for each choice, see Storage Scaling Considerations While Using Fast Provisioning.

Storage performance

For the storage performance type, select one of the following:

- **Balanced** for most workloads that require a good balance between performance and cost savings.
- **Higher performance** for large databases and workloads with high I/O requirements. It is the default performance level.

Provide the following details:

- **Available data storage (GB)**: Select the amount of Block Storage in GB to allocate to the DB system. Available storage can be scaled up as needed after provisioning your DB system.
- **Recovery area storage (GB)**: (Read-only) This field displays the amount of storage available for recovery log data (RECO storage). The recovery area storage is determined based on the storage selected. However, you can change the recovery area storage independently after provisioning. For more information about changing the recovery area storage, see *Scale the DB System*.
- **Expected theoretical max IOPS for data storage**: (Read-only) This field displays the maximum theoretical IOPS that is achievable for the storage you have selected. For Standard x86 shapes, the theoretical max IOPS for data storage will scale proportionally based on the amount of data storage selected. For each GB of data storage, 100 theoretical max IOPS are allocated.

DB system configuration

- **Total node count**: The number of nodes in the DB system. You can specify either one or two nodes. It also depends on the shape and storage you select.
  - Multi-node RAC DB systems require a minimum of two OCPUs per node and are not available on Logical Volume Manager.
  - Ampere A1 shape and VM.Standard2.1 shape are only available on single-node DB systems.
- **Total storage (GB)**: Read-only. Displays the total amount of storage that will be used by the DB system, including storage required by the DB system software. The size of the backup determines the minimum value for available storage.
- **Cluster name**: Optional. Displays only for multi-node DB systems to enable you to specify the cluster to store the node.
- **Theoretical max IOPS**: Read-only. Displays the maximum IOPS that is supported for the instance. It is the minimum of the network IOPS and storage IOPS you selected in the *Shape* and *storage* sections. Maximum theoretical IOPS is calculated based on a database with an 8K block size.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **IOPS limiting factor**: Read-only. Displays either **Storage** or **Network** based on which the theoretical max IOPS is determined. It helps identify if you need to increase storage or increase the network bandwidth (by increasing the number of OCPUs proportionally) for your shape if more IOPS are required.
- **SSH keys**: Add the public key portion of each key pair you want to use for SSH access. Select one of the following options:
  - **Generate SSH key pair**: Use this option to create a new SSH key pair. Select both **Save private key** and **Save public key** when using this option. The private key is downloaded to your local system and must be stored in a safe location. You cannot download another copy of the private key generated during this operation after completing the operation.
  - **Upload SSH key files**: Select this option to browse or drag and drop your existing public key (.pub) files.
  - **Paste SSH keys**: Select this option to paste in individual public keys. To paste multiple keys, select **+ Another SSH key** and supply a single key for each entry.
- **License type**: The type of license you want to use for the DB system. Your choice affects metering for billing.
  - **License included** means the cost of this OCI Database service resource will include both the Oracle Database software licenses and the service.
  - **Bring Your Own License (BYOL)** means you will use your organization's Oracle Database software licenses for this OCI Database service resource. For more information, see Bring Your Own License.

Network information

- Select **Compartment** to select a VCN in a different compartment.
- **Virtual cloud network**: From the list, select a VCN in which to create the DB system. Select **Change compartment** to select a VCN in a different compartment.
- Select **Compartment** to select a subnet in a different compartment.
- **Client subnet**: The subnet to which the DB system attaches. For both single-node and multi-node RAC DB systems, do not use a subnet that overlaps with 192.168.16.16/28, which is used by the Oracle Clusterware private interconnect on the database instance. Specifying an overlapping subnet causes the private interconnect to malfunction.

  Choose a dual stack subnet if you want to configure the DB system with both IPv4 and IPv6 addresses.
- **Network security groups**: Optionally, you can specify one or more network security groups (NSGs) for your DB system. NSGs function as virtual firewalls, enabling you to apply a set of ingress and egress security rules to your DB system. A maximum of five NSGs can be specified.

  For more information, see Access and Security and *Security Rules for the DB System*.

  > **Note**
  >
  > If you select a subnet with a security list, the security rules for the DB system will be a union of the rules in the security list and the NSGs.

  To use network security groups:

Copyright © 2022, 2026, Oracle and/or its affiliates.

  - Switch on the **Use network security groups to control traffic** toggle. Note that you must have a virtual cloud network selected to be able to assign NSGs to your DB system.
  - Specify the NSG to use with the DB system. You may need to use more than one NSG. If you're not sure, contact your network administrator.
  - To use additional NSGs, select **+ Another network security group**.
- **Hostname prefix**: Enter a hostname prefix for the DB system. The host name must begin with an alphabetic character and can contain only alphanumeric characters and hyphens (-). The maximum number of characters allowed is 16.

  > ⚠ **Caution**
  >
  > The host name must be unique within the subnet. If it is not unique, the DB system will fail to provision.

- **Host domain name**: The domain name for the DB system. If the selected subnet uses the Oracle-provided Internet and VCN Resolver for DNS name resolution, then this field displays the domain name for the subnet, and it can't be changed. Otherwise, you can provide your choice of a domain name. Hyphens (-) are not permitted.
- **Host and domain URL**: Combines the host and domain names to display the fully qualified domain name (FQDN) for the database. The maximum length is 64 characters.
- **Private IP type**: Optionally, for non-RAC DB systems, you can define the IP address of the new DB system. This is useful in development contexts where you create and delete a DB system over and over, and you need each new iteration of the DB system to use the same IP address. If you specify an IP address that is currently in use within the subnet, the provisioning operation will fail with an error message regarding the invalid IP address.

  If a dual stack subnet is selected, then both IPv4 and IPv6 address options are displayed.

  - **IPv4 address**: You can either automatically assign an IPv4 address or enter it manually.
    - Select the **Automatically assign IPv4 addresses from subnet** option to assign an address automatically.
    - Select the **Manually assign IPv4 addresses** option to manually enter a private IP address. The IP address should be within the subnet CIDR range.
  - **IPv6 address**: You can either automatically assign an IPv6 address or enter it manually.
    - Select the **Automatically assign IPv6 addresses from subnet** option to assign an address automatically.
    - Select the **Manually assign IPv6 addresses** option to manually enter an IP address. The IP address should be within the subnet CIDR range.

Diagnostics collection

The diagnostics collection and notifications feature enables Oracle Cloud Operations and you to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to events to get notified about resource state changes. You can enable or disable this feature at anytime.

By default the options are selected for enabling. However, you can select to uncheck the diagnostic collection check boxes if you do not require the diagnostic feature.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Enable diagnostic events**: Enables and allows Oracle to collect and send fault notifications about critical, warning, and information events for you.
- **Enable health monitoring**: This diagnostics collection for Oracle Cloud operations viewing is not available for the Base Database Service.
- **Enable incident logs and trace collection**: Enables and allows Oracle to receive event notifications and collect incident logs and traces for fault diagnosis and issue resolution.

> **Note**
>
> You are opting in with the understanding that the list of events and log files can change in the future. You can opt out of this feature at any time.

Advanced options

Expand **Advanced options** to provide the advanced options for this resource.

Management

Expand **Management** to provide the following details:

- **Fault domain**: The fault domain(s) in which the DB system resides. You can select which fault domain to use for your DB system. For multi-node RAC DB systems, you can specify which two fault domains to use. Oracle recommends that you place each node of a multi-node RAC DB system in a different fault domain. For more information about fault domains, see About Regions and Availability Domains.
- **Time zone**: The default time zone for the DB system is UTC, but you can specify a different time zone. The time zone options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system. For more information, see DB System Time Zone. The following options are available:
  - **UTC:** configures your DB system to use coordinated universal time.
  - **Browser-detected:** The console displays the time zone detected by your browser for this option.
  - **Select another time zone:** To manually specify a time zone, first make a choice using the **Region or country selector** to select a geographic region, then use the **Time zone** selector to select your required time zone.

> **Tip**
>
> If you want to set a time zone other than UTC or the browser-detected time zone, and if you do not see the time zone you want, try selecting "Miscellaneous" in the **Region or country** list.

Security

Expand **Security** to provide security details. Optionally, you can specify one or more security attributes to configure Zero Trust Packet Routing (ZPR) for the DB system.

- Select the **Namespace** in which the required security attribute is available.
- Select the **Key** and **Value** of the required security attribute.
- Select **Add security attribute**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> - Administrators must set up security attribute namespaces and security attributes in a tenancy before users can apply security attributes to the DB systems.
> - A security attribute is effective only with appropriate policies. If a security attribute without any policies is added, all access will be denied by default, even if allowed in the Security List or NSGs.
> - If you use the security attribute, the security rules for the DB system will be a union of the rules in the security attributes along with any rules in the security list and the NSGs.
> - You may need to use more than one security attribute. If you're not sure, contact your network administrator.
> - A maximum of 3 security attributes can be specified for a DB system.

For more information about:

- Add, edit, or remove a security attribute after a DB system is created, see *Manage Security Attributes for the DB System*.
- Defining appropriate ZPR policies, see *Zero Trust Packet Routing*.
- Creating a security attribute, see Creating a Security Attribute.

Tags

Expand **Tags** to provide tag details for this resource.

You can add free-form tags or defined tags to this resource. You must have permission to use the tag namespace for defined tags. For information about using tags to manage your OCI resources, see Resource Tags.

> **Tip**
>
> Tags can be applied later or at any time to a resource.

Database basic details

- **Database name**: Enter a name for the database. It is also known as the `DB_NAME`. The database name must begin with an alphabetic character and can contain a maximum of eight alphanumeric characters. Special characters are not permitted.
- **Database unique name suffix**: *Optional.* The second portion of the database's unique name. The complete database unique name is created by appending the database unique name suffix to the database name you specify.
- **Database unique name**: Read-only. Displays the complete database unique name (`DB_UNIQUE_NAME`). The database unique name is a globally unique name for the database. Primary and standby databases in a Data Guard association can share the same database name but must have different database unique names.
- **Unified Auditing**: Select this to enable the unified auditing framework. Unified auditing provides selective and effective auditing inside the Oracle database using policies and conditions. The unified auditing option needs to be specified during provisioning. Unified

Copyright © 2022, 2026, Oracle and/or its affiliates.

auditing is enabled by default for Oracle AI Database 26ai. For more information, see Introduction to Auditing.

- **Database image**: Determines what Oracle Database version is used for the database. You can mix database versions on the DB system, but not editions. By default, the latest Oracle-published database software image is selected.

  Select **Change database image** to use a different Oracle-published image or a custom database software image that you have created in advance, then select an **Image Type**:

  – **Oracle Database software images**: These images contain generally available versions of Oracle Database software.

  – **Custom database software images**: These images are created by your organization and contain customized configurations of software updates and patches. Use the **Compartment** and **Database version** filters to select the custom database software images from a specific compartment or Oracle Database software major release version.

  > **Note**
  >
  > The custom database software image must be based on an Oracle Database release that meets the following criteria:
  >
  > * The release is currently supported by Oracle Cloud Infrastructure.
  > * The release is supported by the hardware model you are provisioning.

  For more information about database software images, see Oracle Database Software Images.

  After selecting a software image, select the **Select** button.

- **PDB name**: Optional. The name of the pluggable database. The PDB name must begin with an alphabetic character and can contain a maximum of eight alphanumeric characters. The only special character permitted is the underscore (_).

Administrator credentials

- **Username**: sys (This is a read-only field). A database administrator named sys will be created with the password you supply.
- **Password**: Enter a password for the administrator. The password must meet the following criteria:
  – A strong password for SYS, SYSTEM, TDE wallet, and PDB administrator.
  – The password must be 9 to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters.
  – The special characters must be _, #, or -.
  – The password must not contain the user name (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reversed order and regardless of casing.
- **Confirm password**: Reenter the password you specified.
- Using a TDE wallet password is optional. If you are using customer-managed encryption keys stored in a vault in your tenancy, the TDE wallet password is not applicable to your DB system. Use **Show advanced options** at the end of the Database Information section to configure customer-managed keys.

Copyright © 2022, 2026, Oracle and/or its affiliates.

If you are using customer-managed keys, or if you want to specify a different TDE wallet password, switch off the **Use the administrator password for the TDE wallet** toggle. If you are using customer-managed keys, leave the TDE password fields blank. To set the TDE wallet password manually, enter a password in the **Enter TDE wallet password** field, and then confirm by entering it into the **Confirm TDE wallet password** field.

Database backups

Switch on or off the **Enable automatic backups** toggle as per your requirement. If you are enabling automatic backups, you can select to configure the **Recovery Service** or **Object Storage** as the backup destination.

> **Note**
>
> To use Recovery Service as a backup destination, the Oracle Database version must be 19.18 or later.

Your choice to use Recovery Service as the backup destination depends on the available limits in your tenancy and the available capacity in the specific region. The following restrictions apply to all regions when you enable automatic backups and want to use Recovery Service as the backup destination:

- If you have available limits and if there is available capacity in the region, then your choices are Recovery Service (default) and Object Storage.
- If you have exhausted the default limits for the Recovery Service, then you can only use Object Storage. However, you can make an additional limits request and then use Recovery Service.
- If there is no available capacity in the region, then you can use only Object Storage. However, after the required capacity becomes available in the region, you can seamlessly switch from Object Storage to Recovery Storage.
- Ampere A1 shape-based DB systems can only be backed up in the Object Storage.

Effective June 23, 2026, when you enable automatic backups using the OCI Console, the Autonomous Recovery Service will be the only available backup destination under the following conditions:

- The tenancy was created on or after June 23, 2026.
- The database is deployed in one of the following OCI regions:
  Netherlands Northwest (Amsterdam) **AMS**, Colombia Central (Bogota) **BOG**, Brazil East (Sao Paulo) **GRU**, Indonesia North (Batam) **HSG**, India South (Hyderabad) **HYD**, US East (Ashburn) **IAD**, South Korea Central (Seoul) **ICN**, Malaysia West 2 (Kulai) **JBP**, Saudi Arabia West (Jeddah) **JED**, Morocco West (Casablanca) **LEJ**, UK South (London) **LHR**, Italy Northwest (Milan) **LIN**, Spain Central (Madrid) **MAD**, Israel Central (Jerusalem) **MTZ**, Italy North (Turin) **NRQ**, Spain Central (Madrid 3) **ORF**, Chile Central (Santiago) **SCL**, Chile West (Valparaiso) **VAP**, Brazil Southeast (Vinhedo) **VCP**, Canada Southeast (Montreal) **YUL**, and Canada Southeast (Toronto) **YYZ**.
- The Oracle Database version is later than 19.18 or 23.4.

If any of these conditions are not met, OCI Object Storage will be shown as a backup destination.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Effective May 15, 2026, when you enable automatic backups using the OCI Console, Autonomous Recovery Service will be the only available backup destination under the following conditions:

- The tenancy was created on or after May 15, 2026.
- The database is deployed in one of the following OCI regions:
  Sweden Central (Stockholm) **ARN**, UAE Central (Abu Dhabi) **AUH**, India West (Mumbai) **BOM**, France Central (Paris) **CDG**, UK West (Newport) **CWL**, UAE East (Dubai) **DXB**, South Africa Central (Johannesburg) **JNB**, Japan Central (Osaka) **KIX**, Australia Southeast (Melbourne) **MEL**, France South (Marseille) **MRS**, Mexico Northeast (Monterrey) **MTY**, US Midwest (Chicago) **ORD**, Mexico Central (Queretaro) **QRO**, Saudi Arabia Central (Riyadh) **RUH**, US West (San Jose) **SJC**, Singapore (Singapore) **SIN**, Australia East (Sydney) **SYD**, Singapore West (Singapore) **XSP**, South Korea North (Chuncheon) **YNY**, or Switzerland North (Zurich) **ZRH**.
- The Oracle Database version is later than 19.18 or 23.4.

If any of these conditions are not met, OCI Object Storage will be shown as a backup destination.

Effective August 06, 2025, when you enable automatic backups through the OCI Console, Autonomous Recovery Service will be the only available backup destination under the following conditions:

- The tenancy was created on or after August 06, 2025.
- The database is deployed in one of the following OCI regions:
  Germany Central (Frankfurt) **FRA**, US West (Phoenix) **PHX**, or Japan East (Tokyo) **NRT**.
- The Oracle Database version is later than 19.18 or 23.4.

If any of these conditions are not met, OCI Object Storage will be shown as a backup destination.

Recovery Service

If **Recovery Service** is selected as the **Backup destination**, you can configure the following options:

- **Protection policy**: You can select from one of the preset protection policies or a custom policy. The system automatically deletes your backups at the end of your chosen protection policy recovery window.

  The following retention periods are available for **Recovery Service**. The retention periods (in days) are defined in the Recovery Service protection policy.

  - Bronze (14 days)
  - **Silver (35 days)** (default)
  - Gold (65 days)
  - Platinum (95 days)
  - Custom (User defined protection policy)
- **Real-time data protection**: Real-time protection is the continuous transfer of redo changes from a protected database to **Recovery Service**. This reduces data loss and provides a recovery point objective (RPO) near 0. This is an extra cost option.
- **Deletion options after database termination**: You can use the following options to retain managed database backups after the database is terminated. These options can also help restore the database from backups in case of accidental or malicious damage to the database.

Copyright © 2022, 2026, Oracle and/or its affiliates.

  - **Retain backups according to the retention period**: When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be removed at the end of the specified retention period.
  - **Retain backups for 72 hours, then delete**: When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be retained for 72 hours and then deleted. The backups are retained for 72 hours to safeguard against accidental deletion by the user.
- **Scheduled day for initial backup**: Select a day of the week for the initial backup to begin.
- **Scheduled time for initial backup (UTC)**: Select a time for the initial backup to begin. The initial backup could start at any time or within the chosen two-hour scheduling window.
- **Scheduled time for daily backup (UTC)**: Select a time for the daily backup to begin. The daily backup could start at any time or within the chosen two-hour scheduling window.
- **Take the first backup immediately**: A full backup is an operating system backup of all data files and the control file that constitute an Oracle Database. A full backup must also include the parameter files associated with the database. You can take a database backup when the database is shut down or while the database is open. You must typically not take a backup after an instance failure or other unusual circumstances. If you select to defer the initial backup, your database may not be recoverable in the event of a database failure.

Object Storage

If **Object Storage** is selected as the **Backup destination**, you can configure the following options:

- **Backup retention period**: If you select to enable automatic backups, you can select a policy with one of the preset retention periods. The system automatically deletes your incremental backups at the end of your chosen retention period. You can change the backup retention period after provisioning.

  The following retention periods are available for **Object Storage**.

  - 7 days
  - 15 days
  - **30 days** (default)
  - 45 days
  - 60 days
- **Deletion options after database termination**: You can use the following options to retain managed database backups after the database is terminated. These options can also help restore the database from backups in case of accidental or malicious damage to the database.
  - **Retain backups according to the retention period**: When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be removed at the end of the specified retention period.
  - **Retain backups for 72 hours, then delete**: When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be retained for 72 hours and then deleted. The backups are retained for 72 hours to safeguard against accidental deletion by the user.
- **Scheduled day for full backup**: Select a day of the week for the initial and future full backups to begin.
- **Scheduled time for full backup (UTC)**: Select a time for the full backup to begin. The full backup could start at any time or within the chosen two-hour scheduling window.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Scheduled time for incremental backup (UTC)**: Select a time for the incremental backup to begin. The incremental backup could start at any time or within the chosen two-hour scheduling window.
- **Take the first backup immediately**: A full backup is an operating system backup of all data files and the control file that constitute an Oracle Database. A full backup must also include the parameter files associated with the database. You can take a database backup when the database is shut down or while the database is open. You must typically not take a backup after an instance failure or other unusual circumstances. If you select to defer the initial backup, your database may not be recoverable in the event of a database failure.

Advanced options

Expand **Advanced options** to provide the advanced options for this resource.

Management

Expand **Management** to provide the following details:

- **Character set**: The character set for the database. The default is AL32UTF8.
- **National character set**: The national character set for the database. The default is AL16UTF16.

Encryption

Expand **Encryption** to configure the encryption key management option for your database. By default, the database is configured using Oracle-managed encryption keys. To configure the database with encryption based on encryption keys you manage, provide the following details:

- Select **Use customer-managed keys**. You must have a valid encryption key in Oracle Cloud Infrastructure Vault service. For more information, see the *Let security admins manage vaults, keys, and secrets* topic in Common Policies.

  > **Note**
  >
  > You must use AES-256 encryption keys for your database.

- Select a **Compartment** where the Vault is located.
- Select a **Vault**.
- Select a **Master encryption key**.
- To specify a key version other than the latest version of the selected key, check **Choose the key version** and enter the OCID of the key you want to use in the **Key version OCID** field.

  > **Note**
  >
  > The key version will only be assigned to the container database (CDB) and not to its pluggable database (PDB). The PDB will be assigned an automatically generated new key version.

Tags

Expand **Tags** to provide tag details for this resource.

Copyright © 2022, 2026, Oracle and/or its affiliates.

You can add free-form tags or defined tags to this resource. You must have permission to use the tag namespace for defined tags. For information about using tags to manage your OCI resources, see Resource Tags.

> **Tip**
>
> Tags can be applied later or at any time to a resource.

Create

Select **Create**. The DB system appears in the list with a status of Provisioning. The DB system's icon changes from yellow to green (or red to indicate errors).

After the DB system's icon turns green, with a status of Available, you can select the highlighted DB system name to display details about the DB system. Note the IP addresses. You'll need the private or public IP address, depending on network configuration, to connect to the DB system.

Create a DB System from a Backup

This article provides the details and procedure to create a new DB system from a backup in Oracle Base Database Service.

General Information

- You can create a DB system from a backup taken using either Autonomous Recovery Service or Object Storage.
- When you create a DB system from a backup, the availability domain can be the same as where the backup is hosted or a different one across regions.
- The shape you specify must be the same type as the database from which the backup was taken. For example, if you are using a backup of a single-node database, then the DB system you select as your target must also be a single-node DB system.
- The Oracle Database version you specify must be an equal or greater version than that of the backed-up database.
- If you specify a DB system shape, then the available storage size will default to the data size of the backup, rounded up to the closest storage size option. However, you can specify a larger storage size.
- If you are creating a new DB system from an Object Storage, you may choose any level 0 weekly backup or a level 1 incremental backup created after the most recent level 0 backup. For more information on backups, see *Back Up and Recovery in Base Database Service*.
- If the backup being used to create a DB system is in a security zone compartment, the DB system cannot be created in a compartment that is not in a security zone. For a full list of policies that affect the resources, see Security Zone Policies.
- When creating a DB system from a backup, the new DB system and database will default to parameters, which are recommended as well as best practices for Oracle Cloud. For example, the memory parameters will be set according to the chosen shape. They might therefore slightly differ from the ones within the backup. You can later modify the configuration manually after creation.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Procedure

Perform the following steps to create a DB system from a backup using the OCI Console:

Go to the **DB Systems** list page. If you need help finding the list page or the DB system, see *List the DB Systems*.

Select a Backup

Navigate to the backup or standalone backup you want to use to create the new DB system:

> **Note**
>
> If you are creating a database from an automatic backup, you may choose any level 0 weekly backup or a level 1 incremental backup created after the most recent level 0 backup.

- To select a daily automatic backup or on-demand full backup as the source:
  1. Find the DB system where the database is located, and select the system name to display details about it.
  2. From the **Databases** list, select the source database name.
  3. Find your desired backup in the **Backups** list. If you don't see the backups list on the **Database Details** page, select **Backups** in the **Resources** menu.
  4. From the **Actions** menu for the backup, select **Create database**.
- To select the last archived redo log automatic backup as the source:
  1. Find the DB system where the database is located, and select the system name to display details about it.
  2. Find the database associated with the backup you wish to use, and select its name to display details about it.
  3. On the **Database Details** page, select **Create database from backup**.
- To specify a timestamp for a point-in-time copy of the source:
  1. Find the DB system where the database is located, and select the system name to display details about it.
  2. Find the database associated with the backup you wish to use, and select its name to display details about it.
  3. On the **Database Details** page, select **Create database from backup**.
  4. In the **Create database from backup** dialog, do the following:
     a. Select **Create database from specified timestamp**.
     b. In the **Restore timestamp** field, enter a timestamp. The restore timestamp determines the most recent data that will be included in the restored version of the database.
     c. Select **Create**.
- To select a standalone backup or a long-term retention backup (LTR) as the source:
  1. Under **Resources**, select **Standalone backups**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

2. In the list of standalone backups, find the backup you want to use to create the database.
3. From the **Actions** menu for the backup you are interested in, select **Create database**.

Basic details

- **Add PDB**: Select the pluggable databases that must be included in the new database.
  - **All PDBs**: This will include all the pluggable databases from the backup.
  - **Choose the PDBs to restore**: Specify one or more pluggable databases to be included from the backup.
- **PDB name**: A comma-separated list of pluggable database names to be included from the backup. It is required when you select **Choose the PDBs to restore** option in **Add PDB**.
- **DB system name**: Enter a name for the DB system. Avoid entering confidential information. OCI assigns an unique Oracle Cloud Identifier (OCID) to the DB system, so the name you enter can be nonunique.
- **Compartment**: Select a compartment for your new DB system from the list. By default, the current compartment is selected.
- **Select a region**: The target region where you want to create the database.
- **Availability domain**: Select an availability domain in which the DB system must reside.
- **Oracle Database software edition**: The database edition supported by the DB system. The database edition cannot be changed later.

Shape

Shape: The shape determines the type of DB system created and the resources allocated to the system. For a complete list of shapes, see Available Shapes and How It Determines the Resources Allocated.

- By default, the **AMD VM.Standard.E4.Flex** shape with **4 OCPU**s is selected.
- The Enterprise Edition Developer is supported only on the Ampere A1 shape.
- If **Enterprise Edition Developer** is selected for the Oracle Database software edition, then by default, the **VM.Standard.A1.Flex** shape with **1 OCPU** is selected.
- The Standard Edition is not supported on the Ampere A1 shape.
- By default, the same shape as the primary is selected for standby.
- To specify a shape other than the default, select **Change shape**.

In the **Change shape** panel, provide the following details to select from available shapes:

Shape series

Select an AMD, Intel, or Ampere processor in the processor group.

- **AMD:** Shapes that use current-generation AMD processors. The AMD shapes are flexible.
- **Intel:** Standard and optimized shapes that use current-generation Intel processors. Both fixed and flexible Intel shapes are available.
- **Ampere**: Shapes that use Arm-based Ampere processors. The Ampere shapes are flexible.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> If you select Ampere A1, AMD E4, AMD E5, or Intel X9 flexible shapes, the memory, network bandwidth, and maximum theoretical IOPS scale proportionally.

Configure OCPU

Select the shape you want to use for this instance.

To change the number of OCPUs, perform the following steps:

- From the actions menu of the desired shape, select **Update OCPU count**.
- Select the desired number of OCPUs per node from the list.
- Select **Update**.

You can change the number of OCPUs for Ampere A1, AMD E4, AMD E5, and Intel X9 flexible shapes.

- For the Ampere A1 shape, a minimum of 1 OCPU and a maximum of 57 OCPUs can be selected.
- For AMD E4 and E5 shapes, a minimum of 1 OCPU and a maximum of 64 OCPUs can be selected.
- For the Intel X9 shape, a minimum of 1 OCPU and a maximum of 32 OCPUs can be selected.

The following resources scale proportionately to the number of **OCPUs** you selected.

- **Memory (GB)**: The amount of memory you want to allocate to this instance.
  For Ampere A1, AMD E4, AMD E5, and Intel X9 shapes, the memory will scale proportionally based on the number of OCPUs selected.
  - For the Ampere A1 shape, for each OCPU, 8 GB of memory is allocated. A minimum of 8 GB and a maximum of 456 GB of memory is allocated.
  - For AMD E4 and E5 shapes, for each OCPU, 16 GB of memory is allocated. A minimum of 16 GB and a maximum of 1024 GB of memory is allocated.
  - For the Intel X9 shape, for each OCPU, 16 GB of memory is allocated. A minimum of 16 GB and a maximum of 512 GB of memory is allocated.
- **Network bandwidth (Gbps)**: The amount of network bandwidth you want to allocate to this instance.
  For Ampere A1, AMD E4, AMD E5, and Intel X9 shapes, the bandwidth will scale proportionally based on the number of OCPUs selected. For each OCPU, 1 Gbps of network bandwidth is allocated.
  - For the Ampere A1 shape, a minimum of 1 Gbps and a maximum of 40 Gbps of network bandwidth is allocated.
  - For AMD E4 and E5 shapes, a minimum of 1 Gbps and a maximum of 40 Gbps of network bandwidth is allocated.
  - For the Intel X9 shape, a minimum of 1 Gbps and a maximum of 32 Gbps of network bandwidth is allocated.
- **Theoretical max IOPS**: The amount of input and output per second (IOPS) you want to allocate to this instance. Theoretical max IOPS is also dependent on the storage you select.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For standard x86 shapes, the theoretical max IOPS will scale proportionally based on the number of ECPUs selected. For each ECPU, 8K theoretical max IOPS is allocated.

For Ampere A1, AMD E4, AMD E5, and Intel X9 shapes, the theoretical max IOPS will scale proportionally based on the number of OCPUs selected. For each OCPU, 16K theoretical max IOPS is allocated.

- For the Ampere A1 shape, a minimum of 16K and a maximum of 640K theoretical max IOPS is allocated.
- For AMD E4 and E5 shapes, a minimum of 16K and a maximum of 640K theoretical max IOPS is allocated.
- For the Intel X9 shape, a minimum of 16K to a maximum of 512K theoretical max IOPS is allocated.

Storage

- **Storage management software**: The storage determines how much storage to provision and the performance settings. By default, **Logical Volume Manager** with **Higher performance** is selected.
- To specify a storage other than the default, select **Change storage**.

> **Note**
>
> Ampere A1 shape is only supported on Logical Volume Manager. When the Ampere A1 shape is selected, the storage management software type changes to Logical Volume Manager with the Higher Performance option.

In the **Change storage** panel, provide the following details to select from available storage:

Storage management software

For the storage management software type, select one of the following:

- **Oracle Grid Infrastructure** to use Oracle Automatic Storage Management (recommended for production workloads)
- **Logical Volume Manager** to quickly provision your DB system using Logical Volume Manager storage management software.

> **Note**
>
> The Available storage (GB) value you specify during provisioning determines the maximum total storage available through scaling. For total storage available for each choice, see Storage Scaling Considerations While Using Fast Provisioning.

Storage performance

For the storage performance type, select one of the following:

- **Balanced** for most workloads that require a good balance between performance and cost savings.
- **Higher performance** for large databases and workloads with high I/O requirements. It is the default performance level.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Provide the following details:

- **Available data storage (GB)**: Select the amount of Block Storage in GB to allocate to the DB system. Available storage can be scaled up as needed after provisioning your DB system.
- **Recovery area storage (GB)**: (Read-only) This field displays the amount of storage available for recovery log data (RECO storage). The recovery area storage is determined based on the storage selected. However, you can change the recovery area storage independently after provisioning. For more information about changing the recovery area storage, see *Scale the DB System*.
- **Expected theoretical max IOPS for data storage**: (Read-only) This field displays the maximum theoretical IOPS that is achievable for the storage you have selected.
  For Standard x86 shapes, the theoretical max IOPS for data storage will scale proportionally based on the amount of data storage selected. For each GB of data storage, 100 theoretical max IOPS are allocated.

DB system configuration

- **Total node count**: The number of nodes in the DB system. You can specify either one or two nodes. It also depends on the shape and storage you select.
  - Multi-node RAC DB systems require a minimum of two OCPUs per node and are not available on Logical Volume Manager.
  - Ampere A1 shape and VM.Standard2.1 shape are only available on single-node DB systems.
- **Total storage (GB)**: Read-only. Displays the total amount of storage that will be used by the DB system, including storage required by the DB system software. The size of the backup determines the minimum value for available storage.
- **Cluster name**: Optional. Displays only for multi-node DB systems to enable you to specify the cluster to store the node.
- **Theoretical max IOPS**: Read-only. Displays the maximum IOPS that is supported for the instance. It is the minimum of the network IOPS and storage IOPS you selected in the *Shape* and *storage* sections. Maximum theoretical IOPS is calculated based on a database with an 8K block size.
- **IOPS limiting factor**: Read-only. Displays either **Storage** or **Network** based on which the theoretical max IOPS is determined. It helps identify if you need to increase storage or increase the network bandwidth (by increasing the number of OCPUs proportionally) for your shape if more IOPS are required.
- **SSH keys**: Add the public key portion of each key pair you want to use for SSH access. Select one of the following options:
  - **Generate SSH key pair**: Use this option to create a new SSH key pair. Select both **Save private key** and **Save public key** when using this option. The private key is downloaded to your local system and must be stored in a safe location. You cannot download another copy of the private key generated during this operation after completing the operation.
  - **Upload SSH key files**: Select this option to browse or drag and drop your existing public key (.pub) files.
  - **Paste SSH keys**: Select this option to paste in individual public keys. To paste multiple keys, select **+ Another SSH key** and supply a single key for each entry.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **License type**: The type of license you want to use for the DB system. Your choice affects metering for billing.
  - **License included** means the cost of this OCI Database service resource will include both the Oracle Database software licenses and the service.
  - **Bring Your Own License (BYOL)** means you will use your organization's Oracle Database software licenses for this OCI Database service resource. For more information, see Bring Your Own License.

Network information

- Select **Compartment** to select a VCN in a different compartment.
- **Virtual cloud network**: From the list, select a VCN in which to create the DB system. Select **Change compartment** to select a VCN in a different compartment.
- Select **Compartment** to select a subnet in a different compartment.
- **Client subnet**: The subnet to which the DB system attaches. For both single-node and multi-node RAC DB systems, do not use a subnet that overlaps with 192.168.16.16/28, which is used by the Oracle Clusterware private interconnect on the database instance. Specifying an overlapping subnet causes the private interconnect to malfunction.

  Choose a dual stack subnet if you want to configure the DB system with both IPv4 and IPv6 addresses.
- **Network security groups**: Optionally, you can specify one or more network security groups (NSGs) for your DB system. NSGs function as virtual firewalls, enabling you to apply a set of ingress and egress security rules to your DB system. A maximum of five NSGs can be specified.

  For more information, see Access and Security and *Security Rules for the DB System*.

  > **Note**
  >
  > If you select a subnet with a security list, the security rules for the DB system will be a union of the rules in the security list and the NSGs.

  To use network security groups:
  - Switch on the **Use network security groups to control traffic** toggle. Note that you must have a virtual cloud network selected to be able to assign NSGs to your DB system.
  - Specify the NSG to use with the DB system. You may need to use more than one NSG. If you're not sure, contact your network administrator.
  - To use additional NSGs, select **+ Another network security group**.
- **Hostname prefix**: Enter a hostname prefix for the DB system. The host name must begin with an alphabetic character and can contain only alphanumeric characters and hyphens (-). The maximum number of characters allowed is 16.

  > **Caution**
  >
  > The host name must be unique within the subnet. If it is not unique, the DB system will fail to provision.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Host domain name**: The domain name for the DB system. If the selected subnet uses the Oracle-provided Internet and VCN Resolver for DNS name resolution, then this field displays the domain name for the subnet, and it can't be changed. Otherwise, you can provide your choice of a domain name. Hyphens (-) are not permitted.
- **Host and domain URL**: Combines the host and domain names to display the fully qualified domain name (FQDN) for the database. The maximum length is 64 characters.
- **Private IP type**: Optionally, for non-RAC DB systems, you can define the IP address of the new DB system. This is useful in development contexts where you create and delete a DB system over and over, and you need each new iteration of the DB system to use the same IP address. If you specify an IP address that is currently in use within the subnet, the provisioning operation will fail with an error message regarding the invalid IP address.

  If a dual stack subnet is selected, then both IPv4 and IPv6 address options are displayed.

  – **IPv4 address**: You can either automatically assign an IPv4 address or enter it manually.

    * Select the **Automatically assign IPv4 addresses from subnet** option to assign an address automatically.
    * Select the **Manually assign IPv4 addresses** option to manually enter a private IP address. The IP address should be within the subnet CIDR range.

  – **IPv6 address**: You can either automatically assign an IPv6 address or enter it manually.

    * Select the **Automatically assign IPv6 addresses from subnet** option to assign an address automatically.
    * Select the **Manually assign IPv6 addresses** option to manually enter an IP address. The IP address should be within the subnet CIDR range.

Diagnostics collection

The diagnostics collection and notifications feature enables Oracle Cloud Operations and you to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to events to get notified about resource state changes. You can enable or disable this feature at anytime.

By default the options are selected for enabling. However, you can select to uncheck the diagnostic collection check boxes if you do not require the diagnostic feature.

- **Enable diagnostic events**: Enables and allows Oracle to collect and send fault notifications about critical, warning, and information events for you.
- **Enable health monitoring**: This diagnostics collection for Oracle Cloud operations viewing is not available for the Base Database Service.
- **Enable incident logs and trace collection**: Enables and allows Oracle to receive event notifications and collect incident logs and traces for fault diagnosis and issue resolution.

> **Note**
>
> You are opting in with the understanding that the list of events and log files can change in the future. You can opt out of this feature at any time.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Advanced options

Expand **Advanced options** to provide the advanced options for this resource.

Management

Expand **Management** to provide the following details:

- **Time zone**: The default time zone for the DB system is UTC, but you can specify a different time zone. The time zone options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system. For more information, see DB System Time Zone. The following options are available:
  - **UTC:** configures your DB system to use coordinated universal time.
  - **Browser-detected:** The console displays the time zone detected by your browser for this option.
  - **Select another time zone:** To manually specify a time zone, first make a choice using the **Region or country selector** to select a geographic region, then use the **Time zone** selector to select your required time zone.

> **Tip**
>
> If you want to set a time zone other than UTC or the browser-detected time zone, and if you do not see the time zone you want, try selecting "Miscellaneous" in the **Region or country** list.

Security

Expand **Security** to provide security details. Optionally, you can specify one or more security attributes to configure Zero Trust Packet Routing (ZPR) for the DB system.

- Select the **Namespace** in which the required security attribute is available.
- Select the **Key** and **Value** of the required security attribute.
- Select **Add security attribute**.

> **Note**
>
> - Administrators must set up security attribute namespaces and security attributes in a tenancy before users can apply security attributes to the DB systems.
> - A security attribute is effective only with appropriate policies. If a security attribute without any policies is added, all access will be denied by default, even if allowed in the Security List or NSGs.
> - If you use the security attribute, the security rules for the DB system will be a union of the rules in the security attributes along with any rules in the security list and the NSGs.
> - You may need to use more than one security attribute. If you're not sure, contact your network administrator.
> - A maximum of 3 security attributes can be specified for a DB system.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For more information about:

- Add, edit, or remove a security attribute after a DB system is created, see *Manage Security Attributes for the DB System*.
- Defining appropriate ZPR policies, see *Zero Trust Packet Routing*.
- Creating a security attribute, see Creating a Security Attribute.

Tags

Expand **Tags** to provide tag details for this resource.

You can add free-form tags or defined tags to this resource. You must have permission to use the tag namespace for defined tags. For information about using tags to manage your OCI resources, see Resource Tags.

> **Tip**
>
> Tags can be applied later or at any time to a resource.

Database basic details

- **Database name**: Enter a name for the database. It is also known as the `DB_NAME`. The database name must begin with an alphabetic character and can contain a maximum of eight alphanumeric characters. Special characters are not permitted.
- **Database unique name suffix**: *Optional.* The second portion of the database's unique name. The complete database unique name is created by appending the database unique name suffix to the database name you specify.
- **Database unique name**: Read-only. Displays the complete database unique name (`DB_UNIQUE_NAME`). The database unique name is a globally unique name for the database. Primary and standby databases in a Data Guard association can share the same database name but must have different database unique names.
- **Unified Auditing**: Select this to enable the unified auditing framework. Unified auditing provides selective and effective auditing inside the Oracle database using policies and conditions. The unified auditing option needs to be specified during provisioning. Unified auditing is enabled by default for Oracle AI Database 26ai. For more information, see Introduction to Auditing.
- **Database image:** *Optional*. By default, the latest database software image as the source database is used. You can specify what Oracle Database version is used for the database. You can mix database versions on the DB system, but not editions.
  - Select **Change database image** to choose a custom database software image that you or someone in your organization has created in your tenancy.
  - Use the **Compartment** filter to select the custom database software images from a specific compartment.
  - Select a database image from the table of available images for the Oracle Database version you selected.
  - Select the **Select** button.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Administrator credentials

- **Username**: sys (This is a read-only field). A database administrator named sys will be created with the password you supply.
- **Password**: Enter a password for the administrator. The password must meet the following criteria:
  - A strong password for SYS, SYSTEM, TDE wallet, and PDB administrator.
  - The password must be 9 to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters.
  - The special characters must be _, #, or -.
  - The password must not contain the user name (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reversed order and regardless of casing.
- **Confirm password**: Reenter the password you specified.
- **Enter the source database's TDE wallet or RMAN password**: *(Applies only to databases using Oracle-managed encryption keys)*. Enter either the TDE wallet password or the RMAN encryption password for the backup, whichever is applicable. The TDE wallet password is the SYS password provided when the database was created by using the OCI Console, API, or CLI. The RMAN encryption password is typically required instead if the password was subsequently changed manually.

Create

Select **Create**. The DB system appears in the list with a status of Provisioning. The DB system's icon changes from yellow to green (or red to indicate errors).

After the DB system's icon turns green, with a status of Available, you can select the highlighted DB system name to display details about the DB system. Note the IP addresses. You'll need the private or public IP address, depending on network configuration, to connect to the DB system.

Copyright © 2022, 2026, Oracle and/or its affiliates.

4

Update

Upgrade a DB System

This article describes the procedures to upgrade the operating system (OS) and Grid Infrastructure (GI) in DB systems using the OCI Console.

You can now upgrade the operating system to Oracle Linux 8 (OL8) and the Oracle Grid Infrastructure to 19c in your DB system.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be granted security access in a **policy** by an administrator. This access is required whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you get a message that you don’t have permission or are unauthorized, verify with your administrator what type of access you have and which **compartment** to work in.

For administrators: The policy in Let database admins manage Oracle Cloud database systems lets the specified group do everything with databases and related Database resources.

If you're new to policies, see Getting Started with Policies and Common Policies. If you want to dig deeper into writing policies for databases, see Details for the Database Service.

Prerequisites

The following are required to upgrade a DB system:

- The DB system must use Oracle Linux 6 (OL6) or Oracle Linux 7 (OL7).
- Oracle recommends having a complete standalone backup of the database.

> **Note**
>
> Oracle Linux Release 7 (OL7) on Base Database Service is nearing its end of life (EOL). For more information, see Support Timeline for Oracle Linux Release 7 (OL7) on Base Database Service (Doc ID 3056652.1).

About Upgrading a DB System

For DB system upgrades, note the following:

- Oracle recommends having a complete standalone backup of the database.
- You can upgrade the Grid Infrastructure software version 19c or 21c to Grid Infrastructure software version 26ai. All earlier Grid Infrastructure versions must be first upgraded to 19c.
- The DB system must use Oracle Linux 8 (OL8) to upgrade the Grid Infrastructure software version to 26ai.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- The DB system must use Oracle Linux 6 (OL6) or Oracle Linux 7 (OL7) to upgrade the Grid Infrastructure software version to 19c.
- You could upgrade the operating system to OL8 from OL6 or OL7.
- When you are on OL7, only the OS will be upgraded to OL8. The GI will not be upgraded.
- Databases earlier than 19.21 or 21.12 will be launched with OL7, which includes restoring a backup or a Data Guard standby creation. Databases from 19.21 or 21.12 onwards will always be provisioned with OL8. The OS version of a new DB system (restore or Data Guard) will therefore depend on the database version, regardless of the OS of the source DB system. If you require OL7 with the latest (e.g., 19.21 or 21.12) database update, you need to launch with an older DB version (e.g., 19.20 or 21.11) and update to the latest version.
- DB system upgrades involve some downtime (even on RAC DB systems). Plan and schedule your upgrade accordingly.
- Oracle recommends disabling the automatic backups before executing the upgrade.

> **Note**
>
> The upgrade process will itself disable and re-enable the automatic backup. However, Oracle recommends you do it after careful consideration.

- The upgrade process creates a new compute instance and clones the data disks.
- The private IP addresses and hostnames are carried over to the upgraded DB system.
- The ephemeral public IP addresses are not carried over and will change in the upgraded DB system.
- Oracle recommends not to launch any new instances within the subnet of the DB system. There is a short interval when the private IP addresses and hostnames are "free" during the upgrade process. During this interval, launching any new instance could reserve the "free" IP addresses and block the upgrade process.

> **Note**
>
> If any instances block the upgrade process, the upgrade process will automatically continue after you terminate the blocking instance.

- **All changes to the OS will be removed and are not carried over to the new OS**.
- All changes to Clusterware other than the database and services will be removed and are not carried over to the new GI stack.
- For databases with a Data Guard association, the upgrade is only allowed on the standby site. It ensures no re-instantiation is required as the standby could be ahead of the primary in case of rollback.
- Oracle recommends switching the Data Guard association to "Maximum Performance" before the upgrade, as otherwise, the primary will run into `NET_TIMEOUT`, and a reset of the configuration is required afterward. After the upgrade, if required, you can turn it to "Maximum Availability".

Copyright © 2022, 2026, Oracle and/or its affiliates.

Roll Forward a Failed Upgrade

You can roll forward upgrades that did not complete successfully for the following reasons using the Oracle support team.

- Database registration failed.
- Database startup failed even after a successful upgrade of OS and GI.

Roll Back a Failed Upgrade

You can roll back upgrades that did not complete successfully. A rollback resets your DB system to the state before the upgrade. All changes to the DB system made during and after the upgrade will be lost.

> **Note**
>
> The public IP addresses present before the upgrade process would have changed and cannot be reclaimed.

After the rollback, you can try upgrading the DB system again after finding and fixing the cause of the earlier upgrade failure.

After your Upgrade is Complete

After a successful upgrade, note the following:

- Check that automatic backup is enabled for the database if you disabled them before upgrading.
- The upgrade exchanges the boot volume. So all customization of the OS will be removed (and have to be reapplied by the customer).
- As the GI stack is reinstalled, you have to reapply all Clusterware customizations (like services, additional virtual IPs, etc.).
- Check that automatic backup is enabled for the database if you had disabled them before upgrading.

Upgrade the Operating System of a DB System

Perform the following steps to upgrade the operating system (OS) of a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB System details page, select the **Updates (OS)** tab to view the list of available OS upgrades for the DB system.
3. Under the **Upgrades** section, from the **Actions** menu for the upgrade you are interested in, select the following actions:
   a. **View details**: View information about the upgrade, such as the description, version, and last successful precheck.

Copyright © 2022, 2026, Oracle and/or its affiliates.

b. **Precheck**: Check for any prerequisites to ensure that the upgrade can be successfully applied.
To review the precheck details, go to the **Update History** tab and, from the **Actions** menu, select **View details**.

c. **Apply OS Update**: Apply the selected upgrade. Oracle recommends that you run the precheck for an upgrade before applying it.

> **Warning**
>
> - For a single-node DB system, you can expect downtime for the DB system while the OS is upgraded.
> - Any customizations to the OS, including additional OS users, groups, and configurations, will not be available on the upgraded OS. You should re-apply these changes to the upgraded system.

In the list, the **State**displays the status of the operation. While an upgrade is being applied, the State changes to **Updating** and the status of the DB system changes to **Updating**. Lifecycle operations on the DB system and its resources might be temporarily unavailable. When the upgrade completes successfully, the State changes to **Applied** and the status of the DB system changes to **Available**.

Upgrade the Grid Infrastructure of a DB System

Perform the following steps to upgrade the Grid Infrastructure (GI) of a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the details page, select the **Updates (GI)** tab to view the list of available Grid Infrastructure upgrades for the DB system.
3. From the **Actions** menu for the upgrade you are interested in, select one of the following actions:
   - **View details**: View the details about this upgrade.
   - **Precheck**: Check for any prerequisites to ensure that the upgrade can be successfully applied.
   - **Apply update**: Applies the selected upgrade. Oracle recommends that you run the precheck operation for an upgrade before applying it.

In the list, the **State** displays the status of the operation. While an update is being applied, the State displays as **Updating**. The status of the DB system also displays as **Updating**. Lifecycle operations on the DB system and its resources might be temporarily unavailable. If the update completes successfully, the State changes to **Applied** and the status of the DB system changes to **Available**.

View the Upgrade History of a DB System

Perform the following steps to view the upgrade history of a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

2. On the details page, select the **Update History** tab to view the history of update and upgrade operations for that DB system.

Roll Back a Failed Upgrade

Perform the following steps to roll back an unsuccessful upgrade of a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the details page, select the **Update History** tab to view the history of update and upgrade operations for that DB system.
   - When there is an unsuccessful upgrade, an alert message will be displayed at the top of the page.
3. Select **Rollback** and confirm.

Update a DB System

This article describes the procedures to apply DB system updates and Database Home updates using the Console and the API.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be granted security access in a **policy** by an administrator. This access is required whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you get a message that you don’t have permission or are unauthorized, verify with your administrator what type of access you have and which **compartment** to work in.

For administrators: The policy in Let database admins manage Oracle Cloud database systems lets the specified group do everything with databases and related Database resources.

If you're new to policies, see Getting Started with Policies and Common Policies. If you want to dig deeper into writing policies for databases, see Details for the Database Service.

Prerequisites

DB systems require access to the Object Storage, including connectivity to the applicable Swift endpoint for Object Storage. We recommend using a service gateway with the VCN to enable this access. For more information, see these topics:

> **Note**
>
> In addition to the prerequisites listed in this section, ensure that the following conditions are met to avoid update failures:
>
> - The `/u01` directory on the database host file system has at least 15 GB of free space to execute update processes.
> - The Oracle Clusterware is running on the DB system.
> - All DB system nodes are running.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For more information, see:

- *Update the DB System Resources Using dbcli*
- *VCN and Subnets*
- *Troubleshoot Update Failures*
- Object Storage FAQ

Currently Available Updates

Version	Architecture Type	DB System Update	Database Update
23.0.0.0 (26ai)	Linux.x86	July 2026, April 2026	July 2026, April 2026, January 2026, October 2025
23.0.0.0 (26ai)	Linux.ARM	Not applicable	July 2026, April 2026, January 2026, October 2025
21.0.0.0	Linux.x86	July 2026, April 2026	July 2026, April 2026, January 2026, October 2025
19.0.0.0	Linux.x86	July 2026, April 2026	July 2026, April 2026, January 2026, October 2025
19.0.0.0	Linux.ARM	Not applicable	April 2026, January 2026, October 2025, July 2025
18.0.0.0	Linux.x86	October 2021 (final update)	October 2021 (final update)
12.2.0.1	Linux.x86	July 2023, April 2023	October 2023, July 2023, April 2023, January 2023
12.1.0.2	Linux.x86	Not applicable	October 2023, July 2023, April 2023, January 2023
11.2.0.4	Linux.x86	Not applicable	October 2023, July 2023, April 2023, January 2023

Note

- For DB systems running version 18.0.0.0 software, the October 2021 update is the final update that will be available.
- For bare metal DB systems, the October 2022 update is the final update that will be available.

For guidelines on how to update Oracle Database releases 12.2, 12.1, and 11.2 for the January 2024 Release Update, see Patching Guidelines - Oracle Cloud Customers with lower database release version 18c and lower (Doc ID 2997504.1).

For a list of Release Updates and Release Update Revisions for Oracle Database 19c, see Oracle Database 19c Proactive Patch Information (Doc ID 2521164.1).

For more information, see:

- Update the Operating System of a DB System Node
- *Upgrade a Database*

Copyright © 2022, 2026, Oracle and/or its affiliates.

About Updating a DB System

Updating a DB system can take several minutes. To minimize the impact on users, run the update at a time when the system has the fewest users. To avoid system interruption, consider implementing a high availability strategy such as Oracle Data Guard.

Oracle recommends backing up your database and testing the update on a test system before you apply the update.

Oracle recommends updating the DB system before updating the database within that DB system. For the Grid Infrastructure, updates are available for two most recent versions (N and N-1). For the Oracle Database, updates are available for four most recent versions (N through N-3). For example, if an instance is using 21c and the latest version of 21c offered is 21.10.0.0, then 21.10.0.0 and 21.9.0.0 are available for Grid Infrastructure updates, while 21.10.0.0, 21.9.0.0, 21.8.0.0, and 21.7.0.0 are available for Oracle Database updates. You can use any of these updates, but Oracle recommends using the latest.

Update the Operating System of a DB System

> **Tip**
>
> Run a precheck to determine whether updates are available for the operating system.

Perform the following steps to update the operating system (OS) of a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, select the **Updates (OS)** tab to view the list of available OS updates for the DB system.
3. Under the **OS update** section, from the **Actions** menu of the status, select the following actions:
   a. **Precheck**: Check for any prerequisites to ensure that the update can be successfully applied.
   b. **Apply OS Update**: Apply the selected update. The latest available OS update is always applied. Oracle recommends that you run the precheck for an update before applying it.
4. To check the status or review details of prechecks or OS updates, go to the **Update History** tab. From the **Actions** menu for the relevant operation, select **View details**.

For 2-node RAC DB systems, you can apply OS updates to individual nodes one at a time, or to both nodes simultaneously by selecting them during the precheck and update process.

> **Warning**
>
> For a single-node DB system, you can expect downtime for the DB system while the OS is updated.

In the list, the **State** displays the status of the operation. While an update is being applied, the State changes to **Updating** and the status of the DB system changes to **Updating**. Lifecycle

Copyright © 2022, 2026, Oracle and/or its affiliates.

operations on the DB system and its resources might be temporarily unavailable. When the update completes successfully, the State changes to **Applied** and the status of the DB system changes to **Available**.

Update the Grid Infrastructure of a DB System

Perform the following steps to update the Grid Infrastructure (GI) of a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the details page, select the **Updates (GI)** tab to view the list of available Grid Infrastructure updates for the DB system.
3. From the **Actions** menu for the update you are interested in, select one of the following actions:
   - **View details**: View the details about this update.
   - **Precheck**: Check for any prerequisites to ensure that the update can be successfully applied.
   - **Apply update**: Applies the selected update. Oracle recommends that you run the precheck operation for an update before applying it.

In the list, the **State** displays the status of the operation. While an update is being applied, the State displays as **Updating**. The status of the DB system also displays as **Updating**. Lifecycle operations on the DB system and its resources might be temporarily unavailable. If the update completes successfully, the State changes to **Applied** and the status of the DB system changes to **Available**.

View the Update History of a DB System

Perform the following steps to view the update history of a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the details page, select the **Update History** tab to view the history of update and upgrade operations for that DB system.

Update the DB System Resources Using dbcli

You can use the dbcli utility to update your DB system resources. This includes updating the DB system, the operating system (OS) on the nodes within the DB system, and Database Homes.

For more information on dbcli commands, see *Oracle Database CLI Reference*.

Prerequisites

1. Preparing for an OS update.
   Before you update the OS, review the following important guidelines and information:
   - Back up the database in the DB system prior to attempting an OS update.
   - Do not remove packages from a DB system. However, you might have to remove custom RPMs (packages that were installed after the system was provisioned) for the update to complete successfully.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Oracle recommends that you test any updates thoroughly on a non-production system before updating a production system.
- The image used to launch a DB system is updated regularly with the necessary updates. After you launch a DB system, you are responsible for applying the required OS security updates published through the Oracle public YUM server.
- To apply OS updates, the virtual cloud network (VCN) in the DB system must be configured to allow access to the YUM repository. For more information, see *VCN and Subnets*.

2. Requirements for using SSH to connect to a DB system.
   To connect to the DB system via SSH, you need the path to the private key associated with the public key used when the DB system was launched.

   You also need the public or private IP address of the DB system.

   Use the private IP address to connect to the system from your on-premises network, or from within the VCN. This includes connecting from a host located on-premises connecting through a VPN or FastConnect to your VCN, or from another host in the same VCN. Use the DB system's public IP address to connect to the system from outside the cloud (with no VPN). You can find the IP addresses in the Console as follows:

   a. On the **DB System Details** page, under **Resources**, click **Nodes**.
   b. View the values that are displayed in the **Public IP address** and **Private IP address & DNS name** columns of the table displaying the **Nodes** of the DB system.

Check for Installed and Available Updates

Perform the following steps to check for installed and available updates.

1. SSH to the DB system.

   ```
   ssh -i <private_key_path> opc@<db_system_ip_address>
   ```

2. Log in as `opc` and then `sudo` to the `root` user. Use `sudo su -` with a hyphen to invoke the `root` user's profile, which will set the `PATH` to the dbcli directory (`/opt/oracle/dcs/bin`).

   ```
   sudo su -
   ```

3. Display the installed update versions by using the Component Command. If the **Available Version** column indicates a version number for a component, you should update the component.

   ```
   dbcli describe-component
   ```

   Output:

   ```
   System Version
   ---------------
   12.1.2.10.0

   Component Name        Installed Version    Available Version
   --------------------- -------------------- --------------------
   OAK                   12.1.2.10.0          up-to-date
   GI                    12.1.0.2.161018      up-to-date
   ORADB12102_HOME1      12.1.0.2.160719      12.1.0.2.161018
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

4. Display the latest update versions available in Object Storage by using the Latestpatch Command.

```
dbcli describe-latestpatch
```

Output:

```
componentType   availableVersion
--------------- --------------------
gi              12.1.0.2.161018
db              11.2.0.4.161018
db              12.1.0.2.161018
oak             12.1.2.10.0
```

Update Server Components

Perform the following steps to update the Grid Infrastructure (GI) and storage management kit (OAK) server components.

1. SSH to the DB system.

```
ssh -i <private_key_path> opc@<db_system_ip_address>
```

2. Log in as `opc` and then `sudo` to the `root` user. Use `sudo su -` with a hyphen to invoke the `root` user's profile, which will set the `PATH` to the dbcli directory (`/opt/oracle/dcs/bin`).

```
sudo su -
```

3. Update the server components by using the Server Command.

```
dbcli update-server
```

Output:

```
{
  "jobId" : "9a02d111-e902-4e94-bc6b-9b820ddf6ed8",
  "status" : "Created",
  "reports" : [ ],
  "createTimestamp" : "January 19, 2017 09:37:11 AM PST",
  "resourceList" : [ ],
  "description" : "Server Patching",
  "updatedTime" : "January 19, 2017 09:37:11 AM PST"
}
```

Note down the `jobId` in the above example.

4. Check the job output by using the Job Commands with the `jobId`.

```
dbcli describe-job -i 9a02d111-e902-4e94-bc6b-9b820ddf6ed8
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Output:

```
Job details
----------------------------------------------------------------
                     ID:  9a02d111-e902-4e94-bc6b-9b820ddf6ed8
            Description:  Server Patching
                 Status:  Running
                Created:  January 19, 2017 9:37:11 AM PST
                Message:

Task Name                                Start
Time                          End Time                          Status
----------------------------------------
---------------------------------- ----------------------------------
----------
Create Patching Repository Directories   January 19, 2017 9:37:11 AM
PST     January 19, 2017 9:37:11 AM PST     Success
Download latest patch metadata           January 19, 2017 9:37:11 AM
PST     January 19, 2017 9:37:11 AM PST     Success
Update System version                    January 19, 2017 9:37:11 AM
PST     January 19, 2017 9:37:11 AM PST     Success
Update Patching Repository               January 19, 2017 9:37:11 AM
PST     January 19, 2017 9:38:35 AM PST     Success
oda-hw-mgmt upgrade                      January 19, 2017 9:38:35 AM
PST     January 19, 2017 9:38:58 AM PST     Success
Opatch updation                          January 19, 2017 9:38:58 AM
PST     January 19, 2017 9:38:58 AM PST     Success
Patch conflict check                     January 19, 2017 9:38:58 AM
PST     January 19, 2017 9:42:06 AM PST     Success
Apply cluster-ware patch                  January 19, 2017 9:42:06 AM
PST     January 19, 2017 10:02:32 AM PST    Success
Updating GiHome version                  January 19, 2017 10:02:32 AM
PST    January 19, 2017 10:02:38 AM PST    Success
```

5. Verify that the server components were updated successfully by using the Component Command. The **Available Version** column should indicate `update-to-date`.

Update Database Home Components

Perform the following steps to update the Database Home components.

1. SSH to the DB system.

   ```
   ssh -i <private_key_path> opc@<db_system_ip_address>
   ```

2. Log in as `opc` and then `sudo` to the `root` user. Use `sudo su -` with a hyphen to invoke the `root` user's profile, which will set the PATH to the dbcli directory (`/opt/oracle/dcs/bin`).

   ```
   sudo su -
   ```

3. Get the `ID` of the Database Home by using the Dbhome Commands.

   ```
   dbcli list-dbhomes
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Output:

```
ID                                       Name                 DB Version Home
Location
---------------------------------------- -----------------  ----------
----------------------------------------
b727bf80-c99e-4846-ac1f-28a81a725df6 OraDB12102_home1
12.1.0.2  /u01/app/orauser/product/12.1.0.2/dbhome_1
```

4. Update the Database Home components by using the Dbhome Commands and providing the `ID` from the previous step.

```
dbcli update-dbhome -i b727bf80-c99e-4846-ac1f-28a81a725df6
```

Output:

```
{
  "jobId" : "31b38f67-f993-4f2e-b7eb-5bccda9901ae",
  "status" : "Created",
  "message" : null,
  "reports" : [ ],
  "createTimestamp" : "January 20, 2017 10:08:48 AM PST",
  "resourceList" : [ ],
  "description" : "DB Home Patching: Home Id is 52e2e799-946a-4339-964b-
c203dee35328",
  "updatedTime" : "January 20, 2017 10:08:48 AM PST"
}
```

Note down the `jobId` in the above example.

5. Check the job output by using the Job Commands with the `jobId`.

```
dbcli describe-job -i 31b38f67-f993-4f2e-b7eb-5bccda9901ae
```

Output:

```
Job details
----------------------------------------------------------------
                     ID:  31b38f67-f993-4f2e-b7eb-5bccda9901ae
            Description:  DB Home Patching: Home Id is b727bf80-c99e-4846-
ac1f-28a81a725df6
                 Status:  Success
                Created:  January 20, 2017 10:08:48 AM PST
                Message:

Task Name                                Start
Time                          End Time                          Status
---------------------------------------
---------------------------------- ----------------------------------
----------
Create Patching Repository Directories   January 20, 2017 10:08:49 AM
PST    January 20, 2017 10:08:49 AM PST    Success
Download latest patch metadata           January 20, 2017 10:08:49 AM
PST    January 20, 2017 10:08:49 AM PST    Success
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
Update System version                    January 20, 2017 10:08:49 AM
PST    January 20, 2017 10:08:49 AM PST    Success
Update Patching Repository               January 20, 2017 10:08:49 AM
PST    January 20, 2017 10:08:58 AM PST    Success
Opatch updation                          January 20, 2017 10:08:58 AM
PST    January 20, 2017 10:08:58 AM PST    Success
Patch conflict check                     January 20, 2017 10:08:58 AM
PST    January 20, 2017 10:12:00 AM PST    Success
db upgrade                               January 20, 2017 10:12:00 AM
PST    January 20, 2017 10:22:17 AM PST    Success
```

6. Verify that the Database Home components were updated successfully by using the Component Command. The **Available Version** column should indicate `update-to-date`.

Check for Available Operating System Updates for DB System Nodes

> **Note**
>
> - Oracle recommends rebooting the DB system if any kernel update is present in OS update.
> - Oracle does not recommend installing OS packages or dependencies that are not part of the version lock that Oracle provides.

1. SSH to the DB system.

   ```
   ssh -i <private_key_path> opc@<db_system_ip_address>
   ```

2. Log in as `opc` and then `sudo` to the `root` user. Use `sudo su -` with a hyphen to invoke the root user's profile, which will set the `PATH` to the dbcli directory (`/opt/oracle/dcs/bin`).

   ```
   sudo su -
   ```

3. Use the `get-availableospatches` command to identify updates you want to apply to the OS. Note that if the `rebootIsRequired` field is "true", you must reboot the DB system after applying the update.

   ```
   dbcli get-availableospatches
   ```

   Output:

   ```
   Update Available     Reboot Required
   -------------------- --------------------
   Yes                  Yes
   ```

   To get the output in JSON, use the following command.

   ```
   dbcli get-availableospatches -j
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Output:

```
{
"updateAvailable" : true,
"rebootIsRequired" : true,
"updateableRpms" : [ "curl.x86_64::7.29.0-59.0.1.el7_9.1",
"freetype.x86_64::2.8-14.el7_9.1",
    "kernel-devel.x86_64::3.10.0-1160.11.1.el7", "kernel-
headers.x86_64::3.10.0-1160.11.1.el7",
    "kernel-uek.x86_64::4.1.12-124.45.6.el7uek", "kernel-uek-
firmware.noarch::4.1.12-124.45.6.el7uek",
    "libX11.x86_64::1.6.7-3.el7_9", "libX11-common.noarch::1.6.7-3.el7_9",
    "libcurl.x86_64::7.29.0-59.0.1.el7_9.1",
"libsmbclient.x86_64::4.10.16-9.el7_9",
    "libwbclient.x86_64::4.10.16-9.el7_9",
"python.x86_64::2.7.5-90.0.1.el7",
    "python-libs.x86_64::2.7.5-90.0.1.el7", "samba-client-
libs.x86_64::4.10.16-9.el7_9",
    "samba-common.noarch::4.10.16-9.el7_9", "samba-common-
libs.x86_64::4.10.16-9.el7_9",
    "sudo.x86_64::1.8.23-10.el7_9.1" ],
"installedRpms" : [ "curl.x86_64::7.29.0-59.0.1.el7",
"freetype.x86_64::2.8-14.el7",
    "kernel-devel.x86_64::3.10.0-1160.2.2.el7", "kernel-
headers.x86_64::3.10.0-1160.2.2.el7",
    "kernel-uek.x86_64::4.1.12-124.43.4.el7uek", "kernel-uek-
firmware.noarch::4.1.12-124.43.4.el7uek",
    "libX11.x86_64::1.6.7-2.el7", "libX11-common.noarch::1.6.7-2.el7",
    "libcurl.x86_64::7.29.0-59.0.1.el7",
"libsmbclient.x86_64::4.10.16-7.el7_9",
    "libwbclient.x86_64::4.10.16-7.el7_9",
"python.x86_64::2.7.5-89.0.1.el7",
    "python-libs.x86_64::2.7.5-89.0.1.el7", "samba-client-
libs.x86_64::4.10.16-7.el7_9",
    "samba-common.noarch::4.10.16-7.el7_9", "samba-common-
libs.x86_64::4.10.16-7.el7_9",
    "sudo.x86_64::1.8.23-10.el7" ]
}
```

Update the Operating System of a DB System Node

This topic explains how to use the `dbcli` to run a precheck and then apply an update to the OS running on a DB system node.

> **Note**
>
> Some OS update operations require a reboot after update is complete. Use the `dbcli get-availableospatches` command as described in the previous topic to determine if the update you are applying requires a reboot.

Perform the following steps to update the OS of a DB system node.

Copyright © 2022, 2026, Oracle and/or its affiliates.

1. SSH to the DB system.

   ```
   ssh -i <private_key_path> opc@<db_system_ip_address>
   ```

2. Log in as `opc` and then `sudo` to the `root` user. Use `sudo su -` with a hyphen to invoke the root user's profile, which will set the `PATH` to the dbcli directory (`/opt/oracle/dcs/bin`).

   ```
   sudo su -
   ```

3. Use the `update-server -c os -p` command to run a precheck:

   ```
   dbcli update-server -c os -p
   ```

   Output:

   ```
   {
     "jobId" : "7fc5cadd-d256-436a-be0d-c2bfe9fd4e95",
     "status" : "Created",
     "message" : null,
     "reports" : [ ],
     "createTimestamp" : "March 01, 2021 07:36:19 AM UTC",
     "resourceList" : [ ],
     "description" : "OS Patching Prechecks",
     "updatedTime" : "March 01, 2021 07:36:20 AM UTC",
     "percentageProgress" : "0%"
   }
   ```

4. If the precheck is successful and uncovers no issues that prevent a successful update operation, you can update the OS. If the precheck is not successful, address the issues identified by the precheck before trying to update the OS.

   > **Note**
   >
   > You can use the `-l` (`--local`) flag to update the server components only in the current node.

   To update the OS, use the `dbcli update-server -c os` command:

   ```
   dbcli update-server -c os
   ```

   Output:

   ```
   {
     "jobId" : "bee1c6d9-45fb-4e5b-8ee8-f02e7cd192ab",
     "status" : "Created",
     "message" : null,
     "reports" : [ ],
     "createTimestamp" : "March 01, 2021 07:37:43 AM UTC",
     "resourceList" : [ ],
     "description" : "OS Patching",
     "updatedTime" : "March 01, 2021 07:37:43 AM UTC",
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
  "percentageProgress" : "0%"
 }
```

5. If the OS update requires a reboot, reboot the server after the update operation is complete.

Upgrade a Database

This article describes the procedure to upgrade a database in a DB system by using the Console and the API.

For Oracle Database release and software support timelines, see Release Schedule of Current Database Releases (Doc ID 742060.1) in the My Oracle Support portal.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be granted security access in a **policy** by an administrator. This access is required whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you get a message that you don't have permission or are unauthorized, verify with your administrator what type of access you have and which **compartment** to work in.

For administrators: The policy in Let database admins manage Oracle Cloud database systems lets the specified group do everything with databases and related Database resources.

If you're new to policies, see Getting Started with Policies and Common Policies. If you want to dig deeper into writing policies for databases, see Details for the Database Service.

Prerequisites

Review the following prerequisites to upgrade an Oracle Database in a DB system.

- The database must be in archivelog mode.
- The database must have flashback enabled.

For detailed information, see:

- Oracle Database Documentation for your database's release version to learn more about these settings
- Create an On-Demand Full Backup of a Database
- *Create a DB System from a Backup*

About Upgrading Databases

Review the following information about database software version upgrades.

- You can upgrade Oracle Database 19c or 21c to Oracle AI Database 26ai. All earlier Oracle Database versions must be first upgraded to Oracle Database 19c.
- The DB system must use Grid Infrastructure software version 26ai to upgrade to Oracle AI Database 26ai.
- The DB system must use Oracle Linux 8 (OL8) to upgrade to Oracle AI Database 26ai.
- The DB system must use Oracle Linux 7 (OL7) to upgrade to Oracle Database 19c.
- If your DB System uses Automatic Storage Management (ASM) storage management software, the system must use Oracle Grid Infrastructure 19c or later.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- For databases on DB systems not meeting the minimum software version requirements, you can upgrade only after using the backup and restore operations to restore the database to a DB system that uses OL7 and GI 19c.
- Database upgrades involve database downtime. Oracle recommends considering this when scheduling your database upgrade.
- Oracle recommends that you back up your database and test the new software version on a test system before you upgrade.
  For more information about creating an on-demand manual backup, see On-Demand Full Backups.
- Oracle recommends running an upgrade precheck operation for your database before you upgrade so that you can discover any issues that need mitigation before the time you plan to perform the upgrade. The precheck operation does not affect database availability. So you can perform it anytime.
- An upgrade operation cannot take place while an automatic backup operation is underway. Before upgrading, Oracle recommends disabling automatic backups and performing a manual backup.
  For more information, see Configure Automatic Backups for a Database and Create an On-Demand Full Backup of a Database.
- After upgrading, you cannot use automatic backups taken before the upgrade to restore the database to an earlier point in time.
- If you are upgrading a database that uses version 11.2 software, the resulting version 19c database will be a non-container database (non-CDB). You can convert the resulting 19c database to a pluggable database (PDB) using the Console or APIs after your upgrade is complete.
  For more information about running a precheck and converting your non-CDB to a PDB, see Convert a Non-Container Database To an Oracle Database 19c PDB.
- For upgrades using generally-available Oracle Database software releases, you cannot use the dbcli utility to perform the upgrade. Use the Console to perform your database upgrade. If your organization must upgrade using a customized software version, contact Oracle to receive a pre-authenticated URL that you can use with the dbcli to download your software. Performing upgrades using dbcli is only possible if Oracle has provided this pre-authenticated request URL (PAR URL).

Upgrading Databases that have Data Guard Association

For databases with a Data Guard association, you must always upgrade the standby database first and then the primary database. The upgrade options are available in the Console for Data Guard associations created using the Console. However, if you have a database that does not have a Data Guard association but is configured as a primary or standby database manually, then the following apply for upgrading such databases.

- The database service will be able to detect and identify whether a database is a primary or standby.
- For database versions 11.2 and 12.1, you must disable the Data Guard configuration before upgrading.
- You must first upgrade the standby database and then the primary database.
- The database service will set a Guaranteed Restore Point (GRP) on the database you are upgrading. After a successful upgrade, the GRP will be removed automatically from the primary database. However, on the standby database, you must manually remove the GRP.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- While upgrading, the `DB_HOME` will be changed, the standby database will open in mount mode, and the primary will open in read/write mode.

After successful upgrade of both primary and standby databases, perform the following.

1. For database versions 11.2 and 12.1, enable the Data Guard configuration you previously disabled before the upgrade.
2. Check the open mode of the standby database.
3. Drop the GRP created on standby.

How the Database Upgrade Operation is Performed by the Database Service

During the database upgrade process, the following steps are automatically performed:

- Executes an automatic precheck. This enables the system to identify issues needing mitigation and to stop the upgrade operation.
- Sets a GRP, enabling it to perform a flashback in the event of an upgrade failure.
- Creates a new Database Home based on the specified Oracle-published or custom database software image.
- Runs the Database Upgrade Assistant (DBUA) software to perform the upgrade on the database. For databases in Data Guard association, this step is executed only on the primary database.

Roll Back a Failed Database Upgrade

> **Note**
>
> The rollback operation is available for Oracle Database Enterprise Editions only.

If your database upgrade does not complete successfully, then you have the option of performing a rollback. Following an unsuccessful database upgrade operation, the rollback option is provided in a banner message displayed on the **Database Details** page.

Review the following information before you begin rollback.

- Rollback resets your database to the state before the upgrade.
- All changes to the database made during and after the upgrade will be lost.

For databases in Data Guard associations, the rollback must be performed according to the following steps.

- If the standby database upgrade has failed and you want to rollback, perform the following steps.
  1. Rollback the standby database.
- If the standby database upgrade has failed and you want to retry, perform the following steps.
  1. Rollback the standby database.
  2. Upgrade the standby database.
- If the primary database upgrade has failed and you want to rollback, perform the following steps.
  1. Rollback the primary database.

Copyright © 2022, 2026, Oracle and/or its affiliates.

2. Rollback the standby database using the CLI. The Console does not provide an option to rollback a successful upgrade.
   For more information about CLI, see, Upgrade Rollback.
3. Reenable Data Guard configuration on primary databases running on 11.2 and 12.1 database versions.
4. After a successful rollback, verify the Data Guard configurations.

- If the primary database upgrade has failed and you want to retry, perform the following steps.
  1. Rollback the primary database.
  2. Execute flashback to GRP on the standby database.
     For more information about GRP, see Managing Guaranteed Restore Points.
  3. Upgrade the primary database.

> **Note**
>
> You must always rollback the primary database first and then the standby database.

Generally, when you rollback the database using the rollback option in the Console, the following steps are taken care of by the database service automatically.

1. Execute flashback.
2. Change Database Home.
3. Drop GRP.

You can rollback a successful standby database upgrade only using CLI.

For more information about the steps to perform a rollback using the Console, see Roll Back a Failed Database Upgrade.

After Your Database Upgrade is Complete

After a successful upgrade, note the following:

- Oracle recommends that you remove the old Database Home using the dbcli utility.
  For more information, see Dbhome Commands.
- Check that automatic backups are enabled for the database if you disabled them prior to upgrading.
- Edit the Oracle Database `COMPATIBLE` parameter to reflect the new Oracle Database version.
  For more information, see What Is Oracle Database Compatibility?.
- Ensure that the `.bashrc` file in the home directory of the Oracle User has been updated to point to the 19c Database Home.
- If you upgraded a database from Oracle Database 11.2 to Oracle Database 19c, you can convert the resulting non-container database to a pluggable database (PDB). You can perform a precheck prior to the conversion to identify problems prior to the conversion operation.
- The GRP that was created on the standby database must be dropped.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Managing Guaranteed Restore Points

Listing Guaranteed Restore Points

You can use the following statement to list all the GRPs using the `V$RESTORE_POINT` View.

```
SELECT NAME FROM V$RESTORE_POINT WHERE GUARANTEE_FLASHBACK_DATABASE='YES';
```

For more information, see Listing Restore Points Using the V$RESTORE_POINT view in *Oracle Database Backup and Recovery User's Guide*.

Dropping Guaranteed Restore Points

You can use the following statement to drop a GRP using the `DROP RESTORE POINT` statement.

```
DROP RESTORE POINT <grp name>;
```

where, `grp name` is the name of the GRP that must be dropped.

For more information, see Dropping Restore Points in *Oracle Database Backup and Recovery User's Guide*.

Upgrade a Database

Perform the following steps to upgrade the database from an older release to a higher major release version using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, select the **Updates** tab to view the list of available upgrades for the database.

> **Note**
>
> Only the most recent update level of Oracle Database and the next-most recent update level can be used for the upgrade operation.

4. From the **Actions** menu for the upgrade you are interested in, select one of the following actions:
   - **View details**: View the details about this upgrade.
   - **Precheck**: Check for any prerequisites to ensure that the upgrade can be successfully applied.
   - **Apply update**: Applies the selected upgrade. Oracle recommends that you run the precheck operation for an upgrade before applying it.

In the list, the **State** displays the status of the operation. While an upgrade is being applied, the State displays as **Upgrading**. The status of the database also displays as **Upgrading**. Lifecycle operations on the database and its resources might be temporarily unavailable. If the

Copyright © 2022, 2026, Oracle and/or its affiliates.

upgrade completes successfully, the State changes to **Applied** and the status of the database changes to **Available**.

View the Database Upgrade History

Perform the following steps to view the upgrade history of a database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the details page, select the **Update History** tab to view the history of update and upgrade operations for that database.

Roll Back a Failed Database Upgrade

Prerequisites

- The upgrade rollback operation is only available for Oracle Database Enterprise Edition softwares that were unsuccessfully upgraded and are currently in the "Failed" lifecycle state.
- Review the information in the Roll Back a Failed Database Upgrade topic before proceeding with the following steps.

Procedure

Perform the following steps to rollback a failed database upgrade to the previous Oracle Database version using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. When there is an unsuccessful upgrade, a banner with a **Rollback** button will be displayed on top of the database details page.
4. Select **Rollback** and confirm.

Convert a Non-Container Database To an Oracle Database 19c PDB

This topic applies to databases upgraded from Oracle Database 11.2 to Oracle Database 19c. As part of the conversion process, you create a new container database (CDB) to hold the pluggable database (PDB) created by the conversion of the non-CDB. To convert a non-container database to a PDB that uses a later version of Oracle Database than 19c, follow the steps in this topic, then upgrade the resulting 19c database to a later software version, as described in Upgrade a Database.

Prerequisites and Recommended Practices

- You must have the TDE wallet password of the non-CDB in order to convert it into a PDB.
- Oracle recommends creating a manual backup of the database before attempting the conversion. For more information, see *Back Up and Recovery in Base Database Service*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- The Console enables you to perform a precheck for your conversion operation to ensure that the conversion can complete successfully. Oracle recommends running the precheck before performing the conversion. To perform the precheck, follow the steps in this procedure, and for the final step, select **Precheck**.
- You can clone the DB system and test the conversion operation on the database in the cloned system before attempting the conversion on the source DB system. For more information, see *Clone a DB System*.

Convert a Non-CDB to a PDB Using the Console

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, on the **Database information** tab, review the **Database architecture** field to confirm that the database is a non-container database.
4. On the database details page, from the **Actions** menu, select **Convert to PDB**.
5. In the **Convert non-CDB database to pluggable database** panel, provide the following information in the **Container database details** section:
   - **Container database name**: Provide a name for the new CDB that will hold your converted PDB.
   - **Password**: Provide a password for the new CDB.
   - **Confirm password**: Reenter the CDB password.
   - **Use the administrator password for the TDE wallet**: Uncheck this option if you want to set a separate password for the TDE wallet. After you uncheck the option, the following fields display:
     - **Enter TDE wallet password**: Provide a TDE wallet password for the new CDB.
     - **Confirm TDE wallet password**: Reenter the TDE wallet password.
6. In the **Non container database details** section, enter the existing TDE wallet password of the non-CDB that you are converting.
7. Select **Precheck** to perform a precheck, or select **Convert to PDB** if you are ready to convert the database.

> **Note**
>
> After you run a precheck, you are returned to the database details page in the Console. To start the conversion operation, follow all the steps in this topic again, and select the **Convert to PDB** option in the final step.

After the database has been successfully converted, the **Database Details** page in the Console displays Container database in the **Database architecture** field. This field is located in the **Database information** tab of the database details page.

Troubleshooting Tips for Converting a Non-CDB to a PDB

If your conversion operation does not complete successfully, you can troubleshoot the issue using the database cli (dbcli) command line utility. For more information about database cli, see *Oracle Database CLI Reference*.

To troubleshoot:

Copyright © 2022, 2026, Oracle and/or its affiliates.

1. Login to the DB system as described in *Overview of Connecting to a DB System*.
2. Use the `dbcli list-jobs` command to determine the job ID and status of the unsuccessful database conversion operation.
3. Use the `dbcli describe-job` command to display details about the unsuccessful database conversion operation.

Based on the information retuned by the `dbcli describe-job` command, you can try to resolve the issue that caused the conversion operation to fail. For errors that occur during the plugging in process (when the new PDB is being plugged into the new CDB), contact Oracle Support for assistance in completing the conversion after the issue that caused the failure has been resolved.

If a conversion operation fails, the console might display either 2 databases in the DB system, or display only a terminated database. The DB system can take up to 2 hours to reset itself. If the Console no longer shows 2 databases or a single terminated database, you can try the conversion again. If the DB system does not reset itself and allow you to try again, contact Oracle Support.

Update a Database

This article describes the procedure to update a database in a DB system by using the Console and the API.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be granted security access in a **policy** by an administrator. This access is required whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you get a message that you don't have permission or are unauthorized, verify with your administrator what type of access you have and which **compartment** to work in.

For administrators: The policy in Let database admins manage Oracle Cloud database systems lets the specified group do everything with databases and related Database resources.

If you're new to policies, see Getting Started with Policies and Common Policies. If you want to dig deeper into writing policies for databases, see Details for the Database Service.

About Updating a Database

Oracle recommends updating the DB system before updating the database within that DB system.

For the Oracle Database, updates are available for four most recent versions (N through N-3). For example, if an instance is using 21c and the latest version of 21c offered is 21.10.0.0, then 21.10.0.0, 21.9.0.0, 21.8.0.0, and 21.7.0.0 are available for Oracle Database updates. You can use any of these updates, but Oracle recommends using the latest.

You can also update your Oracle Database using a custom database software image. When updating with a software image, Oracle supports updating with any image based on any of the four most recent versions (N through N-3).

The following applies only to Oracle based updates starting April 2022 updates for 12.1 and 12.2 and July 2022 updates for 19c.

If you have any interim updates (previously known as a *one-off patch*) installed, then those interim updates will be rolled back automatically before installing the new update. If the new

Copyright © 2022, 2026, Oracle and/or its affiliates.

update does not include the interim updates, you could install the interim update again after you complete installing the new update.

If you have any interim updates installed in your database, then to avoid installing interim updates after the new update, Oracle recommends using a custom database software image and updating your database using that image.

After running precheck, you can identify if you have installed any additional interim updates by executing the following command in dbcli.

```
dbcli describe-job -i <job id>
```

where `job id` (`dbcli list-jobs`) is the job id of the precheck run.

All the changes in the Oracle Home will be copied to the newly updated and installed database.

> **Note**
>
> - Oracle does not recommend changing any files within the Database Home.
> - It is essential to ensure that all the changes you made to the Database Home are copied correctly. If you have copied any files or folders in the Database Home, like `tnsnames.ora`, `listener.ora` etc., you should back up those files manually.

The DB system should have access to the Identity and Object Storage endpoints within OCI. If the DB system running on a private subnet was only configured with a service gateway and no additional gateway (e.g., NAT Gateway), you should ensure that the service gateway is set to allow the access to all Oracle Services and not just the Object Storage. No change is required if the Identity and Object Storage endpoints can be reached by other means.

For more information on the list of currently available database updates, see Currently Available Updates.

> **Note**
>
> OJVM updates have to be applied manually using the OPATCH tool.

Apply Interim Updates Using a Database Software Image

We do not recommend applying interim updates to the DB systems. Instead, you should create and deploy a database software image that contains the necessary fixes. Only in cases where you cannot use a database software image, for example, to apply an interim update to a GI Home, you can file a Service Request (SR) and request support assistance to download an interim update.

For more information on database software images, see *Software Images*.

Update a Database

Perform the following steps to update the database from an older update release using the OCI Console:

Copyright © 2022, 2026, Oracle and/or its affiliates.

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, select the **Updates** tab to view the list of available updates for the database.
4. From the **Actions** menu for the update you are interested in, select one of the following actions:
   - **View details**: View the details about this update.
   - **Precheck**: Check for any prerequisites to ensure that the update can be successfully applied.
   - **Apply update**: Applies the selected update. Oracle recommends that you run the precheck operation for an update before applying it.

In the list, the **State** displays the status of the operation. While an update is being applied, the State displays as **Updating**. The status of the database also displays as **Updating**. Lifecycle operations on the database and its resources might be temporarily unavailable. If the update completes successfully, the State changes to **Applied** and the status of the database changes to **Available**.

View the Database Update History

Perform the following steps to view the update history of a database using the OCI Console:

> **Note**
>
> Update history in the Console do not show updates that were applied by using command line tools like dbcli or the Patch utility.

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the details page, select the **Update History** tab to view the history of update and upgrade operations for that database.

Apply an Interim Update

> **Note**
>
> This topic applies only to database homes in single-node and multi-node RAC DB systems.

To apply an interim update to fix a specific defect, follow the procedure in this section. Use the `Opatch` utility to apply an interim update to the Database Home.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> In the procedure example, the Database Home directory is **/u02/app/oracle/product/12.1.0.2/dbhome_1** and the update number is **26543344**.

1. Obtain the applicable interim update from My Oracle Support.
2. Review the information in the update `README.txt` file. This file might contain additional and/or custom instructions to follow to apply the update successfully.
3. Use SCP or SFTP to place the update on your target database.
4. Shut down each database that is running in the Database Home.

   ```
   srvctl stop database -db <db name> -stopoption immediate -verbose
   ```

5. Set the Oracle Home environment variable to point to the target Oracle Home.

   ```
   sudo su - oracle
   export ORACLE_HOME=/u02/app/oracle/product/12.1.0.2/dbhome_1
   ```

6. Change to the directory where you placed the update, and unzip the update.

   ```
   cd <working directory where opatch is stored>
   unzip p26543344_122010_Linux-x86-64.zip
   ```

7. Change to the directory with the unzipped update, and check for conflicts.

   ```
   cd 26543344
   $ORACLE_HOME/OPatch/opatch prereq CheckConflictAgainstOHWithDetail -ph ./
   ```

8. Apply the update.

   ```
   $ORACLE_HOME/OPatch/opatch apply
   ```

9. Verify that the update was applied successfully.

   ```
   $ORACLE_HOME/OPatch/opatch lsinventory -detail -oh $ORACLE_HOME
   ```

10. If the Database Home contains databases, restart them.

    ```
    $ORACLE_HOME/bin/srvctl start database -db <db name>
    ```

    Otherwise, run the following command as `root` user.

    ```
    # /u01/app/<db version>/grid/bin/setasmgidwrap o=/u01/app/oracle/
    product/<db version>/dbhome_1/bin/oracle
    ```

11. If the readme indicates that the update has a `sqlpatch` component, run the `datapatch` command against each database.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Before you run `datapatch`, ensure that all pluggable databases (PDBs) are open. To open a PDB, you can use SQL*Plus to execute the following against the PDB.

```
ALTER PLUGGABLE DATABASE <pdb name> OPEN READ WRITE
```

```
$ORACLE_HOME/OPatch/datapatch
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

5

Manage

Pluggable Databases

About Pluggable Databases

This article provides details about pluggable databases and managing their various features, such as backup, restore, relocate, and clone.

The multitenant architecture enables an Oracle Database to be a container database. A container database (CDB) contains one or more user-created, pluggable databases and application containers. A pluggable database (PDB) is a portable collection of schemas, schema objects, and nonschema objects that appears to an application as a separate database. At the physical level, each pluggable database has its own set of data files that store the data for the pluggable database. The container database includes all the data files for the pluggable databases contained within it and a set of system data files that store metadata for the container database itself.

Oracle 19c or later databases created in a DB system include an initial pluggable database that you can access from the database details page in the OCI Console. Using the OCI Console or APIs, you can start, stop, clone, and delete pluggable databases. You can also create additional pluggable databases in the container database. All pluggable database operations performed using the OCI Console or APIs can be monitored using the work request generated by the operation. For more information, see Work Requests.

You can create and manage pluggable databases in the DB system using the OCI Console and APIs.

> **Note**
>
> Generally, the term 'database' refers to the container database (CDB).

Create

You can have more than one pluggable database in a pluggable database. Pluggable databases must be created one at a time, and creating a new pluggable database has no effect on existing pluggable databases in the pluggable database.

To create a pluggable database using the OCI Console, see *Create a Pluggable Database*.

Restore

A pluggable database can be restored from a backup. The following two types of restore are supported:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Out-of-place restore**: You can restore a pluggable database by creating a container database from the backup, then selecting a pluggable database or a subset of them you want to restore on the new database.
- **In-place restore**: You can restore a pluggable database within the same container database to its last known good state or to a specified timestamp.

You can perform an in-place restore when you want to move a pluggable database back to a specified state or time. Both the container database and pluggable database must be up and running, and only one pluggable database can be restored at a time.

- If you have multiple pluggable databases in your container database and want to restore multiple of them to the same container database, then you could restore each individual pluggable database, one pluggable database at a time, from the container database backup.
- When the container database is down, you can restore the complete container database, and all the pluggable databases in that container database will also be restored.
- You could either restore the database to the specified time stamp or to its last known good state.

To restore a pluggable database using the OCI Console, see *Restore a Pluggable Database*.

Relocate

You can relocate a pluggable database from one container database to another container database within the same availability domain (AD) to the same or a later database version across compartments, DB systems, or VCNs. If two different VCNs are used, then both VCNs must be peered before relocating. During relocation, the pluggable database will be removed from the source container database and moved to the destination container database that is up and running. In a Data Guard association, a pluggable database relocated to the primary will be synchronized with the standby as well.

To relocate a pluggable database using the OCI Console, see *Relocate a Pluggable Database*.

Clone

A clone is an independent and complete copy of the given database as it existed at the time of the cloning operation. You can create clones of your pluggable database within the same container database or a different container database and refresh the cloned pluggable database.

The following types of clones are supported:

- **Local clone**: A copy of the pluggable database is created within the same container database.
- **Remote clone**: A copy of the pluggable database is created on a different container database.
  You can perform a remote clone of a pluggable database from one container database to another container database within the same availability domain (AD) to the same or a later database version across compartments, DB systems, or VCNs. If two different VCNs are used, then both VCNs must be peered before cloning.
- **Refreshable clone**: A copy of the pluggable database is created on a different container database, and you will be able to refresh the cloned pluggable database.
  You can perform a refreshable clone of a pluggable database from one container database to another container database within the same availability domain (AD) to the same or a later database version across compartments, DB systems, or VCNs. If two different VCNs are used, then both VCNs must be peered before cloning.

Copyright © 2022, 2026, Oracle and/or its affiliates.

To clone a pluggable database using the OCI Console, see *Clone a Pluggable Database*.

Refreshable Clone

A refreshable clone enables you to keep your remote clone updated with the source pluggable database.

- You can only refresh while the pluggable database is in mount mode.
- The only open mode you can have is read-only, and refresh cannot be performed while it is in read-only mode.
- A database link user credential is required for creating a refreshable clone.
- Clone, relocate, and in-place restore operations are not supported in the refreshable clone. Relocate and in-place restore operations are not supported in the source, and the source can only be deleted after disconnecting or deleting the refreshable clone.
- When cloning a pluggable database from 19c to 26ai, the cloned pluggable database is automatically upgraded to 26ai. For example, if you use refreshable clones to clone to 26ai and then convert it to a regular pluggable database, all necessary upgrade steps are automatically handled, converting the refreshable clone into a fully upgraded 26ai pluggable database.
- In a Data Guard association, a refreshable clone cannot be created on standby, but it can be created on the primary. However, the primary will not be synced to the standby.
- A pluggable database in standby cannot be used as the source for a refreshable pluggable database.

To create a refreshable clone using the OCI Console, see *Clone a Pluggable Database*.

To refresh a clone using the OCI Console, see *Refresh a Pluggable Database*.

Convert Refreshable Clone to Regular Pluggable Database

You can convert a refreshable clone to a regular pluggable database by disconnecting the refreshable clone (destination pluggable database) from the source pluggable database at any time. If the refresh pluggable database is in a Data Guard association, when it is converted to a regular pluggable database, then the pluggable database will be synced to the standby as part of the conversion process.

To convert a pluggable database using the OCI Console, see *Convert a Refreshable Clone to a Regular Pluggable Database*.

SQL Worksheets

The SQL Worksheet provides a web-based SQL workspace where you can enter SQL statements directly in the browser using a database connection. From the SQL Worksheet, you can run SQL statements or scripts against the database and create database objects.

You need to create a connection to use the SQL Worksheet. The SQL Worksheet utilizes the connection you create with the service to provide you the ability to run SQL commands and scripts from the Console. Scripts used in the SQL Worksheet can reside in either OCI Object Storage or on your local drive. Using the Connection selection menu, you can change the connection that the SQL Worksheet is using instantly.

Connections are resources that contain the necessary information for accessing an Oracle Database in OCI. Connections are created by simply providing information about the location

Copyright © 2022, 2026, Oracle and/or its affiliates.

of the database. The connection also contains the user used to access the database and the location of the password that is stored in the OCI Vault.

To create a connection, launch and use the SQL worksheet using the OCI Console, see *SQL Worksheets*.

Backup

You can take a backup of the pluggable database optionally during create, clone, or relocate operations when the container database is configured with the auto-backup feature. The pluggable database backup destination will always be the same as the container database, and the backups cannot be accessed directly or created on demand. Oracle recommends immediately backing up the pluggable database after you create or clone it. This is because the pluggable database will not be recoverable until the next daily auto-backup completes successfully, leading to a possible data loss.

Open Modes

On the OCI Console, you can see the open modes of a pluggable database, such as read-write, read-only, and mounted. If the pluggable database status is the same across all nodes, the system displays the same status for all pluggable databases. If the pluggable database statuses are different across the nodes, the system displays a message indicating on which nodes the pluggable databases are opened in read-write mode. You cannot change the open mode of a pluggable database through the API or OCI Console. However, you can start or stop a pluggable database. Starting the pluggable database will start it in read-write mode. Stopping the pluggable database will close it, and it will remain in mount mode.

Limitations of Pluggable Database Management

- New pluggable databases created with SQL are not immediately discovered and displayed in the OCI Console. However, OCI performs sync operations regularly to discover manually created pluggable databases, and they should be visible in the OCI Console and API-based tools within 6 hours of creation. Oracle recommends using the OCI Console or API-based tools (including the OCI CLI, SDKs, and Terraform) to create pluggable databases.
- Pluggable database operations are supported only for databases using Oracle Database 19c and later.
- Pluggable databases are backed up at the container database level, and each backup includes all the pluggable databases in the container database. OCI does not support the creation of backups for individual pluggable databases.

Create a Pluggable Database

This article provides the details and procedure to create a pluggable database.

You can have more than one pluggable database in a pluggable database. Pluggable databases must be created one at a time, and creating a new pluggable database has no effect on existing pluggable databases in the pluggable database.

Procedure

Perform the following steps to create a pluggable database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the database details page, select the **Pluggable Databases** tab.
4. Select **Create**.
5. In the **Create pluggable database** panel, enter the following values:

Basic details

- **PDB name**: Enter a name for the pluggable database. The name must begin with an alphabetic character and can contain a maximum of 30 alphanumeric characters.

PDB admin account

- **Unlock the PDB admin account**: Optional. Switch on this toggle to specify a pluggable database admin password and configure the pluggable database to be unlocked at creation.
  - **PDB admin password**: Create and enter a pluggable database admin password. The password must contain:
    * A minimum of 9 and a maximum of 30 characters.
    * At least two uppercase characters.
    * At least two lowercase characters.
    * At least two special characters. The valid special characters are: underscore (_), a hash sign (#), and a dash (-). You can use two of the same characters or any combination of two of the same characters.
    * At least two numeric characters (0 - 9).
  - **Confirm PDB admin password**: Reenter the pluggable database admin password.

Database admin credential

- **TDE wallet password of database**: Enter the TDE wallet password for the source container database.
- **Take a backup of the PDB immediately**: You must enable auto-backup on the container database to back up a pluggable database immediately. This toggle is switched on by default if auto-backup is enabled on the container database.

> **Note**
>
> If the toggle is switched off, the system displays a warning stating that the pluggable database cannot be recovered until the next daily backup has been successfully completed.

Advanced options

Expand **Advanced options** to provide the advanced options for this resource.

Tags

Expand **Tags** to provide tag details for this resource.

Copyright © 2022, 2026, Oracle and/or its affiliates.

You can add free-form tags or defined tags to this resource. You must have permission to use the tag namespace for defined tags. For information about using tags to manage your OCI resources, see Resource Tags.

> **Tip**
>
> Tags can be applied later or at any time to a resource.

Create

Select **Create**.

During the pluggable database creation operation, the source container database is in the 'updating' status.

Restore a Pluggable Database

This article provides the details and procedure to restore a pluggable database.

A pluggable database can be restored from a backup. The following two types of restore are supported:

- **Out-of-place restore**: You can restore a pluggable database by creating a container database from the backup, then selecting a pluggable database or a subset of them you want to restore on the new database.
- **In-place restore**: You can restore a pluggable database within the same container database to its last known good state or to a specified timestamp.

You can perform an in-place restore when you want to move a pluggable database back to a specified state or time. Both the container database and pluggable database must be up and running, and only one pluggable database can be restored at a time.

- If you have multiple pluggable databases in your container database and want to restore multiple of them to the same container database, then you could restore each individual pluggable database, one pluggable database at a time, from the container database backup.
- When the container database is down, you can restore the complete container database, and all the pluggable databases in that container database will also be restored.
- You could either restore the database to the specified time stamp or to its last known good state.

In-Place Restore

Perform the following steps to in-place restore a pluggable database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the pluggable database details page, from the **Actions** menu, select **Restore**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

5. In the **Restore PDB** panel, select one of the following restore options:
   - **Restore to latest**: Restores the database to its last known good state with the least possible data loss.
   - **Restore to a timestamp**: Restores the database to the timestamp specified.

   > **Note**
   >
   > SCN-based restore is not supported by the OCI Console but is available through the API and OCI CLI.

6. Select **Restore** and confirm.

Out-of-Place Restore

Perform the following steps to out-of-place restore a pluggable database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the database details page, select **Backups** tab to view the list of database backups.
4. From the **Actions** menu for the database backup that contains the pluggable databases you want to restore, select **Create database**.
5. In the **Create database from backup** panel, provide details about the new DB system by following the instructions provided in Create a DB System from a Backup.

Relocate a Pluggable Database

This article provides the details and procedure to relocate a pluggable database.

You can relocate a pluggable database from one container database to another container database within the same availability domain (AD) to the same or a later database version across compartments, DB systems, or VCNs. If two different VCNs are used, then both VCNs must be peered before relocating. During relocation, the pluggable database will be removed from the source container database and moved to the destination container database that is up and running. In a Data Guard association, a pluggable database relocated to the primary will be synchronized with the standby as well.

Procedure

Perform the following steps to relocate a pluggable database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.

Copyright © 2022, 2026, Oracle and/or its affiliates.

4. On the pluggable database details page, from the **Actions** menu, select **Relocate**.
5. In the **Relocate pluggable database** panel, provide the following details:

Destination

- **Compartment**: Select the required compartment.
- **DB System**: Select the destination DB system to which the pluggable database must be migrated.
- **Database**: Select the destination database to which the pluggable database must be migrated.

New PDB

- **PDB name**: Enter a name for the pluggable database. The name must begin with an alphabetic character and can contain a maximum of 30 alphanumeric characters.
- **Database SYS password**: Enter the admin password for the source container database.
- **Database TDE wallet password**: Enter the TDE wallet password for the source container database.
- **Unlock the PDB admin account**: Optional. Switch on this toggle to specify a pluggable database admin password and configure the pluggable database to be unlocked at creation.
  - **PDB admin password**: Create and enter a pluggable database admin password. The password must contain:
    * A minimum of 9 and a maximum of 30 characters.
    * At least two uppercase characters.
    * At least two lowercase characters.
    * At least two special characters. The valid special characters are: underscore (_), a hash sign (#), and a dash (-). You can use two of the same characters or any combination of two of the same characters.
    * At least two numeric characters (0 - 9).
  - **Confirm PDB admin password**: Reenter the pluggable database admin password.

Source

- **Source database SYS password:** Enter the database admin password of the source database.
- **Database link**: Optional. Enter the user name and password for the database link. A common user is created at the container database level. If you do not provide the details, the system creates the user and deletes it at the end of the operation. If the user name entered already exists in the database, relocate will reuse the same user name. However, the user's password will be reset to the one you entered.
  - **Database link user name**: Provide a user name for the database link.
  - **Database link password**: Provide a password for the database link.
  - **Confirm database link password**: Reenter the password for the database link.
- **Take a backup of the PDB immediately**: You must enable auto-backup on the container database to back up a pluggable database immediately. This toggle is switched on by default if auto-backup is enabled on the container database.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> If the toggle is switched off, the system displays a warning stating that the pluggable database cannot be recovered until the next daily backup has been successfully completed.

Advanced options

Expand **Advanced options** to provide the advanced options for this resource.

Tags

Expand **Tags** to provide tag details for this resource.

You can add free-form tags or defined tags to this resource. You must have permission to use the tag namespace for defined tags. For information about using tags to manage your OCI resources, see Resource Tags.

> **Tip**
>
> Tags can be applied later or at any time to a resource.

Relocate

Select **Relocate**.

- After the relocation is successful, the state of the pluggable database will change from **Available** to **Relocated** in the source container database.
- In the destination container database, the new pluggable database will be added, and its state will be displayed as **Available**.
- Relocation will incur downtime during the process, and the time required is based on the size of the pluggable database.

Clone a Pluggable Database

This article provides the details and procedure to clone a pluggable database.

A clone is an independent and complete copy of the given database as it existed at the time of the cloning operation. You can create clones of your pluggable database within the same container database or a different container database and refresh the cloned pluggable database.

The following types of clones are supported:

- **Local clone**: A copy of the pluggable database is created within the same container database.
- **Remote clone**: A copy of the pluggable database is created on a different container database.
  You can perform a remote clone of a pluggable database from one container database to another container database within the same availability domain (AD) to the same or a later database version across compartments, DB systems, or VCNs. If two different VCNs are used, then both VCNs must be peered before cloning.
- **Refreshable clone**: A copy of the pluggable database is created on a different container database, and you will be able to refresh the cloned pluggable database.

Copyright © 2022, 2026, Oracle and/or its affiliates.

You can perform a refreshable clone of a pluggable database from one container database to another container database within the same availability domain (AD) to the same or a later database version across compartments, DB systems, or VCNs. If two different VCNs are used, then both VCNs must be peered before cloning.

Prerequisite

To clone the pluggable database, you must have the TDE wallet password of the pluggable database's source database.

Procedure

Perform the following steps to clone a pluggable database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the pluggable database details page, from the **Actions** menu, select **Clone**.
5. In the **Clone pluggable database** panel, provide the following details:

Basic details

- Select a **clone type** according to your requirements from one of the following options:
  - **Local clone**: Create a copy of the source pluggable database on the same container database.
  - **Remote clone**: Create a copy of the source pluggable database on a different container database.
  - **Refreshable clone**: Create a copy of the source pluggable database on a different container database and be able to refresh the cloned DPB.

Destination

- **Compartment**: Select a compartment.
- **DB System**: Select the destination DB system to which the pluggable database must be cloned.
- **Database**: Select the destination database to which the pluggable database must be cloned.
- **PDB name**: Enter a name for the pluggable database. The name must begin with an alphabetic character and can contain a maximum of 30 alphanumeric characters.
- **Database SYS password**: Enter the admin password for the destination container database.
- **Database TDE wallet password**: Enter the TDE wallet password for the destination container database.
- **Unlock PDB admin account**: Optional for local clone and remote clone. Not applicable for refreshable clone. Switch on this toggle to specify a pluggable database admin password and configure the pluggable database to be unlocked at creation.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **PDB admin password**: Create and enter a pluggable database admin password. The password must contain:
  - A minimum of 9 and a maximum of 30 characters.
  - At least two uppercase characters.
  - At least two lowercase characters.
  - At least two special characters. The valid special characters are: underscore (_), a hash sign (#), and a dash (-). You can use two of the same characters or any combination of two of the same characters.
  - At least two numeric characters (0 - 9).
- **Confirm PDB admin password**: Reenter the pluggable database admin password.

Source

This section is applicable only when the selected clone type is a remote clone or refreshable clone.

- **Source database SYS password:** Enter the database admin password of the source database.
- **Database link**: Required for refreshable clone and optional for remote clone. A common user is created at the container database level. If you do not provide the details, the system creates the user and deletes it at the end of the operation. If the user name entered already exists in the database, the remote clone will reuse the same username. However, the user's password will be reset to the one you entered.
  - **Database link user name**: Provide a user name for the database link.
  - **Database link password**: Provide a password for the database link.
  - **Confirm database link password**: Reenter the password for the database link.
- **Take a backup of the PDB immediately**: You must enable auto-backup on the container database to back up a pluggable database immediately. This toggle is switched on by default if auto-backup is enabled on the container database.

> **Note**
>
> If the toggle is switched off, the system displays a warning stating that the pluggable database cannot be recovered until the next daily backup has been successfully completed.

Advanced options

Expand **Advanced options** to provide the advanced options for this resource.

Tags

Expand **Tags** to provide tag details for this resource.

You can add free-form tags or defined tags to this resource. You must have permission to use the tag namespace for defined tags. For information about using tags to manage your OCI resources, see Resource Tags.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Tip**
>
> Tags can be applied later or at any time to a resource.

Clone

Select **Clone**.

Refresh a Pluggable Database

This article provides the details and procedure to refresh a pluggable database.

A refreshable clone enables you to keep your remote clone updated with the source pluggable database.

- You can only refresh while the pluggable database is in mount mode.
- The only open mode you can have is read-only, and refresh cannot be performed while it is in read-only mode.
- A database link user credential is required for creating a refreshable clone.
- Clone, relocate, and in-place restore operations are not supported in the refreshable clone. Relocate and in-place restore operations are not supported in the source, and the source can only be deleted after disconnecting or deleting the refreshable clone.
- When cloning a pluggable database from 19c to 26ai, the cloned pluggable database is automatically upgraded to 26ai. For example, if you use refreshable clones to clone to 26ai and then convert it to a regular pluggable database, all necessary upgrade steps are automatically handled, converting the refreshable clone into a fully upgraded 26ai pluggable database.
- In a Data Guard association, a refreshable clone cannot be created on standby, but it can be created on the primary. However, the primary will not be synced to the standby.
- A pluggable database in standby cannot be used as the source for a refreshable pluggable database.

Procedure

Perform the following steps to refresh a pluggable database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the pluggable database details page, from the **Actions** menu, select **Refresh**, and then confirm.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Convert a Refreshable Clone to a Regular Pluggable Database

This article provides the details and procedure to convert a refreshable clone to a regular pluggable database.

You can convert a refreshable clone to a regular pluggable database by disconnecting the refreshable clone (destination pluggable database) from the source pluggable database at any time. If the refresh pluggable database is in a Data Guard association, when it is converted to a regular pluggable database, then the pluggable database will be synced to the standby as part of the conversion process.

Procedure

Perform the following steps to convert a refreshable clone to a regular pluggable database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the pluggable database details page, from the **Actions** menu, select **Convert to regular PDB**.
5. In the **Convert to regular PDB** panel, provide the following details:
   - **TDE wallet password of database**: Enter the TDE wallet password for the source container database.
   - **Take a backup of the PDB immediately**: You must enable auto-backup on the container database to back up a pluggable database immediately. This toggle is switched on by default if auto-backup is enabled on the container database.

     > **Note**
     >
     > If the toggle is switched off, the system displays a warning stating that the pluggable database cannot be recovered until the next daily backup has been successfully completed.

   - Select **Convert**.

Stop a Pluggable Database

This article provides the details and procedure to stop a pluggable database.

Prerequisite

The pluggable database must be available and running (started) to stop that pluggable database.

Procedure

Perform the following steps to stop a pluggable database using the OCI Console:

Copyright © 2022, 2026, Oracle and/or its affiliates.

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the pluggable database details page, from the **Actions** menu, select **Stop**, and then confirm.

Start a Pluggable Database

This article provides the procedure to start a pluggable database.

Procedure

Perform the following steps to start a pluggable database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the pluggable database details page, from the **Actions** menu, select **Start**, and then confirm.

Delete a Pluggable Database

This article provides the procedure to delete a pluggable database.

Procedure

Perform the following steps to delete a pluggable database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the pluggable database details page, from the **Actions** menu, select **Delete**, and then confirm.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Get Connection Strings for a Pluggable Database

This article provides the details and procedure to get connection strings for the administrative service of a pluggable database.

Procedure

Perform the following steps to get the connection strings for a pluggable database using the OCI Console:

> **Note**
>
> Oracle recommends that you connect applications to an application service using strings created for the application service.

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the pluggable database details page, from the **Actions** menu, select **PDB connection**.
5. In the **Pluggable database connection** panel, in the **Connection Strings** section, a list of existing connection strings are displayed.
   - From the **Actions** menu for the connection string, select **Copy connection string** to copy connection string to the clipboard.
   - Click **Close** to exit the panel.

SQL Worksheets

This article provides the details and procedure to create a connection, launch and use SQL worksheets.

The SQL Worksheet provides a web-based SQL workspace where you can enter SQL statements directly in the browser using a database connection. From the SQL Worksheet, you can run SQL statements or scripts against the database and create database objects.

You need to create a connection to use the SQL Worksheet. The SQL Worksheet utilizes the connection you create with the service to provide you the ability to run SQL commands and scripts from the Console. Scripts used in the SQL Worksheet can reside in either OCI Object Storage or on your local drive. Using the Connection selection menu, you can change the connection that the SQL Worksheet is using instantly.

Connections are resources that contain the necessary information for accessing an Oracle Database in OCI. Connections are created by simply providing information about the location of the database. The connection also contains the user used to access the database and the location of the password that is stored in the OCI Vault.

For more information about:

- connection, see Managing a Connection.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- SQL Worksheet, see Using the SQL Worksheet.

Create a Connection

Perform the following steps to create a connection using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the pluggable database details page, select the **Connections** tab.
5. Select **Create connection**.
6. In the **Create connection** panel, provide the following information about the connection:
7. **Name**: A user-friendly, informative name to describe the connection.
8. **Compartment**: Select a compartment from the list.
9. **Username**: The database user name you want to use for the connection.
10. **Role**: Select a role from the list. Use this menu to select a high-level, system-wide administrative privileged role to be granted to the user you provided. If no role is needed, you can leave the **default** value.
11. **User password secret:** This menu is populated with any secrets you have access to from the OCI Vault. Select **Change compartment** to find a secret in a different compartment.
12. **Create password secret:** If no secrets are listed or a new secret must be created, use this and provide the following information in the **Create password secret** panel:
    a. **Name:** Give the secret a name. Do not use the password or hints of the password in the name. For example, if a connection to the sales pluggable database is needed for the DBA user, the name could be `salesPDB-DBA`.
    b. **Description:** Optionally, provide a description of the secret.
    c. **Compartment:** Select a compartment in which you would like to create the secret.
    d. **Vault:** Choose an OCI vault that you have access to where the secret will be kept. Select **Change compartment** to find a vault in a different compartment.
    e. **Encryption key:** Select an encryption key to be used to encrypt the supplied password in the vault. Select **Change compartment** to find an encryption key in the same vault that is contained in a different compartment.
    f. **User password:** Provide the password for the user.
    g. **Confirm user password:** Retype the password previously entered.
    h. Select **Create** when done to create the secret in the vault.
13. The **Connection string** field is pre-populated.
14. If you want to designate that this connection will use a private endpoint, switch on the **Access database via a private network** toggle.
    - Select the endpoint using the **select private endpoint** menu. Select **Compartment** to find a private endpoint in a different compartment.

Copyright © 2022, 2026, Oracle and/or its affiliates.

15. A wallet must be provided when the use of mutual TLS (mTLS) authentication is required or when TLS authentication is used and the database returns a certificate not signed by a trusted certificate authority. Oracle recommends using an SSO wallet.

    Choose one of the following options in the **Wallet format** menu:

    - None
    - Java Key Store (e.g., keystore.jks, truststore.jks)
    - PKCS#12
    - SSO wallet (e.g, cwallet.sso)

16. Select **Create** to create a database connection.

For more information, see Managing a Connection.

Launch SQL Worksheet

Perform the following steps to launch the SQL Worksheet using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the pluggable database details page, select **Launch SQL Worksheet**.
5. In the **Launch SQL Worksheet** panel, select the connection you want to use and select **Launch SQL Worksheet**.

Use SQL Worksheet

For detailed steps about using SQL Worksheet, see Using the SQL Worksheet.

DB Systems

List the DB Systems

This article provides the details and procedure to list the DB systems.

Procedure

1. Open the **navigation menu**, select **Oracle Database**, and then select **Oracle Base Database Service**.

   The **DB Systems** list page opens. All DB systems in the selected compartment are displayed in a table.

2. To view the DB systems in a different compartment, use the **Compartment** filter to switch compartments.

   You must have permission to work in a compartment to see the resources in it. If you're not sure which compartment to use, contact an administrator. For more information, see *Understanding Compartments* in Learn Best Practices for Setting Up Your Tenancy.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Filter List Results

Use filters to limit the DB systems in the list.

Perform the following action to filter the list results:

- From the **Search and Filter** box above the list table, select one or more filters and specify the values that you want to use to narrow the list.

  In general, the filters correspond to the columns shown in the list table, although some filters represent attributes that aren't shown in the table.

  The **Compartment** filter is always displayed next to **Applied filters**.

For information about searching for resources and managing the columns in the list table, see OCI New Console Experience.

Change Sort Order

Change the order of the items in the list table by using the sort icons next to the column names.

Actions

In the list table, select the name of a DB system to open its details page, where you can view its status and perform other tasks.

To perform an action on a DB system directly from the list table, select any of the following options from the **Actions** menu in the row for that DB system:

- **View details**: Open the details page for the DB system.
- **Move resource**: Move the DB system to another compartment. See *Move a DB System to Another Compartment*.
- **Copy OCID**: Copy the OCID of the DB system to the clipboard.
- **Manage security attributes**: Add, update, or remove security attributes associated with the DB system. See *Manage Security Attributes for the DB System*.
- **Manage tags**: Add one or more tags to the DB system. See *Manage Tags for the DB System*.
- **Terminate**: Terminate the DB system. See *Terminate a DB System*.

To create a DB system, select **Create DB system**. See *Create a DB System*.

Check the Status of a DB System

This article provides the procedure to check the status of a DB system.

Procedure

Perform the following steps to check the status of a DB system using the OCI Console:

1. On the **DB Systems** list page, find the DB system you're interested in and check its state. If you need help finding the list page or the DB system, see *List the DB Systems*. The color of the icon and the text in it indicates the status of the system.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Provisioning**: Yellow icon. Resources are being reserved for the DB system, the system is booting, and the initial database is being created. Provisioning can take several minutes. The system is not ready to use yet.
- **Available**: Green icon. The DB system was successfully provisioned. A few minutes after the system enters this state, you can SSH to it and begin using it.
- **Terminating**: Gray icon. The DB system is being deleted by the terminate action in the OCI Console or API.
- **Terminated**: Gray icon. The DB system has been deleted and is no longer available.
- **Failed**: Red icon. An error condition prevented the provisioning or continued operation of the DB system.

Node Status

Perform the following steps to view the status of a database node using the OCI Console:

1. On the DB systems list page, select the DB system that contains the node that you want to check. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, select the **Nodes** tab to view the list of nodes.

In addition to the states listed for a DB system, a node's status can be one of the following:

- **Starting**: Yellow icon. The node is being powered on by the start or reboot action in the OCI Console or API.
- **Stopping**: Yellow icon. The node is being powered off by the stop or reboot action in the OCI Console or API.
- **Stopped**: Yellow icon. The node was powered off by the stop action in the OCI Console or API.

Scale the DB System

This article provides the details and procedure to scale up the storage for a DB system.

If a DB system requires more block storage, you can increase the storage at any time.

Procedure

Perform the following steps to scale up the storage for a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, select **Scale storage up**.
3. In the **Scale storage up** panel, provide the following details to select from available storage:

> **Note**
>
> The Available storage (GB) value you specify during provisioning determines the maximum total storage available through scaling. The total storage available for each choice is detailed in Storage Scaling Considerations for Databases Using Fast Provisioning.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- The **Storage management software** and **Storage performance** sections display the values you had selected while provisioning and cannot be changed.
- **Available data storage (GB)**: Select the amount of Block Storage in GB to allocate to the DB system. Available storage can be scaled up as needed after provisioning your DB system.
- **Recovery area storage (GB)**: Select the amount of storage required for recovery log data (RECO storage). The recovery area storage is determined based on the storage selected.

> **Tip**
>
> Oracle recommends keeping recovery storage at 20% of total storage or higher. Oracle charges for the total storage used, including data storage, recovery storage, and storage required for the system software.

- **Total storage (GB)**: Read-only. Displays the total amount of storage that will be used by the DB system, including storage required by the DB system software. The size of the backup determines the minimum value for available storage.
- **Expected theoretical max IOPS for data storage**: (Read-only) This field displays the maximum theoretical IOPS that is achievable for the storage you have selected. For Standard x86 shapes, the theoretical max IOPS for data storage will scale proportionally based on the amount of data storage selected. For each GB of data storage, 100 theoretical max IOPS are allocated.

4. Select **Update**.

> **Note**
>
> To change the number of OCPUs for the DB system, see *Change the Shape of a DB System*.

Change the Shape of a DB System

This article provides the details and procedure to change the shape of a DB system.

After you provision a DB system, you can change the shape at any time to adapt to changes in performance requirements. For example, you might require a system with more OCPUs, or you might want to reduce costs by reducing the number of OCPUs.

> **Note**
>
> The shape-changer operation takes place in a rolling fashion for multi-node RAC DB systems, allowing you to change the shape with no database downtime.

Changing the shape does not affect the amount of storage available to the DB system. However, the new shape can have different memory and network bandwidth characteristics, and you can reapply any customizations to these aspects after the change.

Prerequisites

- The DB system and database are in the 'Available' state.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- The DB system is registered with the Cluster Ready Services (CRS) Grid Infrastructure stack. By default, the DB systems use CRS.
- Database can be successfully restarted.
- Database is configured to use SPFILE (server parameter file), not PFILE. By default, databases in the DB systems use the SPFILE configuration.
- The `SGA_TARGET` parameter for Automatic Shared Memory Management (ASMM) has a nonzero value. By default, the DB systems use this ASMM configuration.

Migration to BaseDB x86 Shape

You can migrate your DB system from Standard x86 shape to Base DB x86 shape using the **Change shape**operation. The following prerequisites and restrictions apply to the migration of DB systems.

- Migration from Base DB x86 to any other shape is not supported.
- The DB system must use Oracle Database version 26ai only.

Migration to AMD Shapes

You can migrate your DB system from Intel X7 shapes and AMD E4 shapes to AMD E5 shapes using the **Change shape** operation. The following prerequisites and restrictions apply to the migration of DB systems.

- Migration from Ampere A1 and Intel X9 shapes to AMD E4 and E5 shapes and vice versa is not supported.
- The DB system must use the kernel UEK5. If the DB system is on UEK4, the migration operation will automatically update it to UEK5.
- For migration from Intel X7 shapes to Intel X9 or AMD E4 shapes, the DB system must use Oracle Linux 7 (OL7) or later. If a DB system is running on a version below OL7, you must upgrade the DB system to OL7 or later before migration. Oracle recommends OL8.
- For migration from Intel X7 or AMD E4 shape to AMD E5 shape, the DB system must use Oracle Linux 8 (OL8). If a DB system is running on a version below OL8, you must upgrade the DB system to OL8 before migration.
- The DB system must use an Oracle Grid Infrastructure version later than or equal to 19.15.
- The DB system must use the following Oracle Database versions or later versions for respective database version series.
  - 23.4.0.24.05
  - 21.6.0.0.0
  - 19.15.0.0.0
  - 12.2.0.1.220418
  - 12.1.0.2.220419
- The DB systems based on Oracle Database 11.2 cannot be migrated.
- You can migrate from Intel-based 2.X to AMD-based X OCPUs only. For example, if you are in the Intel VM.Standard2.2 shape, you will be able to migrate to the AMD 2 OCPU shape. However, after you have migrated, you will be able to change the AMD-based OCPUs according to the available options.
- After migrating to an AMD shape, you can scale up the storage to a maximum of 40 TB (the maximum allowed storage options available on Intel X7 shapes).

Copyright © 2022, 2026, Oracle and/or its affiliates.

- After migration to an AMD shape, you will be on the same storage performance option as the previous shape. You will not be able to change to a different option.
- While creating a clone for a migrated DB system, the clone will have the same characteristics as the migrated DB system.
- While creating Oracle Data Guard for a database in a migrated DB system, the standby will have the same characteristics as the migrated DB system.
- After successful migration, you will not be able to migrate back from AMD to Intel shapes.
- After successful migration, you will not be able to restore to an old boot volume backup.

Procedure

Perform the following steps to change the shape using the OCI Console:

> ⚠ **Caution**
>
> Changing shape requires a restart.

> ⓘ **Note**
>
> Shape change from AMD E5 flexible shape is currently not supported.

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, from the **Actions** menu, select **Change shape**.
3. In the **Change shape** panel, enter the following details. For a complete list of shapes, see Available Shapes and How It Determines the Resources Allocated.

Shape series

This section displays the values you had selected while provisioning and cannot be changed.

- If you are changing the shape from Intel-based X7 shapes to AMD-based flexible shape E4, then the number of OCPUs cannot be changed. You can migrate from Intel-based 2.X to AMD-based X OCPUs only. For example, if you are in the Intel VM.Standard2.2 shape, you will be able to migrate to the AMD 2 OCPU shape. However, after you have migrated, you will be able to change the AMD-based OCPUs according to the available options.
- You can migrate from **Standard x86** shape to **Base DB x86** shape.

Configure OCPU

Select the shape you want to use for this instance.

To change the number of OCPUs, perform the following steps:

- From the actions menu of the desired shape, select **Update OCPU count**.
- Select the desired number of OCPUs per node from the list.
- Select **Update**.

You can change the number of OCPUs for Ampere A1, AMD E4, AMD E5, and Intel X9 flexible shapes.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- For the Ampere A1 shape, a minimum of 1 OCPU and a maximum of 57 OCPUs can be selected.
- For AMD E4 and E5 shapes, a minimum of 1 OCPU and a maximum of 64 OCPUs can be selected.
- For the Intel X9 shape, a minimum of 1 OCPU and a maximum of 32 OCPUs can be selected.

The following resources scale proportionately to the number of **OCPUs** you selected.

- **Memory (GB)**: The amount of memory you want to allocate to this instance.
  For Ampere A1, AMD E4, AMD E5, and Intel X9 shapes, the memory will scale proportionally based on the number of OCPUs selected.
  - For the Ampere A1 shape, for each OCPU, 8 GB of memory is allocated. A minimum of 8 GB and a maximum of 456 GB of memory is allocated.
  - For AMD E4 and E5 shapes, for each OCPU, 16 GB of memory is allocated. A minimum of 16 GB and a maximum of 1024 GB of memory is allocated.
  - For the Intel X9 shape, for each OCPU, 16 GB of memory is allocated. A minimum of 16 GB and a maximum of 512 GB of memory is allocated.
- **Network bandwidth (Gbps)**: The amount of network bandwidth you want to allocate to this instance.
  For Ampere A1, AMD E4, AMD E5, and Intel X9 shapes, the bandwidth will scale proportionally based on the number of OCPUs selected. For each OCPU, 1 Gbps of network bandwidth is allocated.
  - For the Ampere A1 shape, a minimum of 1 Gbps and a maximum of 40 Gbps of network bandwidth is allocated.
  - For AMD E4 and E5 shapes, a minimum of 1 Gbps and a maximum of 40 Gbps of network bandwidth is allocated.
  - For the Intel X9 shape, a minimum of 1 Gbps and a maximum of 32 Gbps of network bandwidth is allocated.
- **Theoretical max IOPS**: The amount of input and output per second (IOPS) you want to allocate to this instance. Theoretical max IOPS is also dependent on the storage you select.
  For standard x86 shapes, the theoretical max IOPS will scale proportionally based on the number of ECPUs selected. For each ECPU, 8K theoretical max IOPS is allocated.

  For Ampere A1, AMD E4, AMD E5, and Intel X9 shapes, the theoretical max IOPS will scale proportionally based on the number of OCPUs selected. For each OCPU, 16K theoretical max IOPS is allocated.
  - For the Ampere A1 shape, a minimum of 16K and a maximum of 640K theoretical max IOPS is allocated.
  - For AMD E4 and E5 shapes, a minimum of 16K and a maximum of 640K theoretical max IOPS is allocated.
  - For the Intel X9 shape, a minimum of 16K to a maximum of 512K theoretical max IOPS is allocated.

Change shape

Select **Change** and confirm.

- If your shape change operation is not successful, see *Troubleshoot Shape Change Failures*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Clone a DB System

This article provides the details and procedure to clone a DB system.

Cloning creates a copy of a source DB system as it exists at the time of the cloning operation, including the storage configuration software and database volumes. When creating a clone, you can specify a new SSH key and admin password.

General Information

- To clone a DB system that has a Data Guard association, initiate the operation from the primary DB system. The clone operation does not clone Data Guard associations themselves, or Data Guard connections.
- When cloning a DB system that uses customer-managed encryption keys, the cloned database will be configured to use the same key version as the source database. For information on using customer-managed keys, see *Manage Encryption Keys*.
- The SSH keys selection only applies to this clone and does not affect the source DB system.
- The license type selection only applies to this clone and does not affect the source DB system.
- The clone can use a different VCN and subnet from the source DB system.
- If the clone is created in a different subnet from the source, the same host name can be used for both the clone and the source DB system.
- The TDE wallet password is inherited from the source DB system for databases using Oracle-managed encryption keys. When cloning a DB system that uses customer-managed encryption keys, the cloned database will be configured to use the same key version as the source database.

Limitations

- When cloning a DB system that uses Real Application Clusters (RAC), a new Oracle Grid Infrastructure (GI) configuration is created. The new GI is required to avoid conflicts with the source DB system. Therefore, the clone DB system does not include the following from the source system:
  - manually added clusterware resources,
  - database application services,
  - customized settings from the source database, such as environment variables,
  - manually added application IP addresses (application virtual IPs),
  - additional listener ports (such as those configured for Transport Layer Security or other purposes),
  - or any other resource or customization that is not present after the creation of a new DB system
- Cloning a RAC DB system takes longer than cloning a single-node DB system due to the time needed to create a new GI stack. Expect a RAC DB system cloning operation to take at least an hour.
- For DB systems using Oracle Automatic Storage Management (ASM), the GI software must be 19.9 or later.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Cloning is not currently supported for DB systems using Oracle Database 21c with Oracle Automatic Storage Management.
- You can't clone a DB system in a security zone to create a DB system that isn't in a security zone. See the security zone policies topic for a full list of policies that affect Database service resources.

For more information, see Oracle Automatic Storage Management and Security Zone Policies.

Procedure

Perform the following steps to clone a DB system using the OCI Console:

- On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
- On the DB system details page, from the **Actions** menu, select **Clone**.

Basic details

- **DB system name**: Enter a name for the DB system. Avoid entering confidential information. OCI assigns an unique Oracle Cloud Identifier (OCID) to the DB system, so the name you enter can be nonunique.
- **Compartment**: Select a compartment for your new DB system from the list. By default, the current compartment is selected.
- **SSH keys**: Add the public key portion of each key pair you want to use for SSH access. Select one of the following options:
  - **Generate SSH key pair**: Use this option to create a new SSH key pair. Select both **Save private key** and **Save public key** when using this option. The private key is downloaded to your local system and must be stored in a safe location. You cannot download another copy of the private key generated during this operation after completing the operation.
  - **Upload SSH key files**: Select this option to browse or drag and drop your existing public key (.pub) files.
  - **Paste SSH keys**: Select this option to paste in individual public keys. To paste multiple keys, select **+ Another SSH key** and supply a single key for each entry.
- **License type**: The type of license you want to use for the DB system. Your choice affects metering for billing.
  - **License included** means the cost of this OCI Database service resource will include both the Oracle Database software licenses and the service.
  - **Bring Your Own License (BYOL)** means you will use your organization's Oracle Database software licenses for this OCI Database service resource. For more information, see Bring Your Own License.

Network information

- Select **Compartment** to select a VCN in a different compartment.
- **Virtual cloud network**: From the list, select a VCN in which to create the DB system. Select **Change compartment** to select a VCN in a different compartment.
- Select **Compartment** to select a subnet in a different compartment.
- **Client subnet**: The subnet to which the DB system attaches. For both single-node and multi-node RAC DB systems, do not use a subnet that overlaps with 192.168.16.16/28,

Copyright © 2022, 2026, Oracle and/or its affiliates.

which is used by the Oracle Clusterware private interconnect on the database instance. Specifying an overlapping subnet causes the private interconnect to malfunction.

Choose a dual stack subnet if you want to configure the DB system with both IPv4 and IPv6 addresses.

- **Network security groups**: Optionally, you can specify one or more network security groups (NSGs) for your DB system. NSGs function as virtual firewalls, enabling you to apply a set of ingress and egress security rules to your DB system. A maximum of five NSGs can be specified.

  For more information, see Access and Security and *Security Rules for the DB System*.

  > **Note**
  >
  > If you select a subnet with a security list, the security rules for the DB system will be a union of the rules in the security list and the NSGs.

  To use network security groups:

  – Switch on the **Use network security groups to control traffic** toggle. Note that you must have a virtual cloud network selected to be able to assign NSGs to your DB system.

  – Specify the NSG to use with the DB system. You may need to use more than one NSG. If you're not sure, contact your network administrator.

  – To use additional NSGs, select **+ Another network security group**.

- **Hostname prefix**: Enter a hostname prefix for the DB system. The host name must begin with an alphabetic character and can contain only alphanumeric characters and hyphens (-). The maximum number of characters allowed is 16.

  > **Caution**
  >
  > The host name must be unique within the subnet. If it is not unique, the DB system will fail to provision.

- **Host domain name**: The domain name for the DB system. If the selected subnet uses the Oracle-provided Internet and VCN Resolver for DNS name resolution, then this field displays the domain name for the subnet, and it can't be changed. Otherwise, you can provide your choice of a domain name. Hyphens (-) are not permitted.
- **Host and domain URL**: Combines the host and domain names to display the fully qualified domain name (FQDN) for the database. The maximum length is 64 characters.
- **Private IP type**: Optionally, for non-RAC DB systems, you can define the IP address of the new DB system. This is useful in development contexts where you create and delete a DB system over and over, and you need each new iteration of the DB system to use the same IP address. If you specify an IP address that is currently in use within the subnet, the provisioning operation will fail with an error message regarding the invalid IP address.

  If a dual stack subnet is selected, then both IPv4 and IPv6 address options are displayed.

  – **IPv4 address**: You can either automatically assign an IPv4 address or enter it manually.

    * Select the **Automatically assign IPv4 addresses from subnet** option to assign an address automatically.

Copyright © 2022, 2026, Oracle and/or its affiliates.

    * Select the **Manually assign IPv4 addresses** option to manually enter a private IP address. The IP address should be within the subnet CIDR range.
  - **IPv6 address**: You can either automatically assign an IPv6 address or enter it manually.
    * Select the **Automatically assign IPv6 addresses from subnet** option to assign an address automatically.
    * Select the **Manually assign IPv6 addresses** option to manually enter an IP address. The IP address should be within the subnet CIDR range.

Diagnostics collection

The diagnostics collection and notifications feature enables Oracle Cloud Operations and you to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to events to get notified about resource state changes. You can enable or disable this feature at anytime.

By default the options are selected for enabling. However, you can select to uncheck the diagnostic collection check boxes if you do not require the diagnostic feature.

- **Enable diagnostic events**: Enables and allows Oracle to collect and send fault notifications about critical, warning, and information events for you.
- **Enable health monitoring**: This diagnostics collection for Oracle Cloud operations viewing is not available for the Base Database Service.
- **Enable incident logs and trace collection**: Enables and allows Oracle to receive event notifications and collect incident logs and traces for fault diagnosis and issue resolution.

> **Note**
>
> You are opting in with the understanding that the list of events and log files can change in the future. You can opt out of this feature at any time.

Advanced options

Expand **Advanced options** to provide the advanced options for this resource.

Management

Expand **Management** to provide the following details:

- **Fault domain**: The fault domain(s) in which the DB system resides. You can select which fault domain to use for your DB system. For multi-node RAC DB systems, you can specify which two fault domains to use. Oracle recommends that you place each node of a multi-node RAC DB system in a different fault domain. For more information about fault domains, see About Regions and Availability Domains.

Security

Expand **Security** to provide security details. Optionally, you can specify one or more security attributes to configure Zero Trust Packet Routing (ZPR) for the DB system.

- Select the **Namespace** in which the required security attribute is available.
- Select the **Key** and **Value** of the required security attribute.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Select **Add security attribute**.

> **Note**
>
> - Administrators must set up security attribute namespaces and security attributes in a tenancy before users can apply security attributes to the DB systems.
> - A security attribute is effective only with appropriate policies. If a security attribute without any policies is added, all access will be denied by default, even if allowed in the Security List or NSGs.
> - If you use the security attribute, the security rules for the DB system will be a union of the rules in the security attributes along with any rules in the security list and the NSGs.
> - You may need to use more than one security attribute. If you're not sure, contact your network administrator.
> - A maximum of 3 security attributes can be specified for a DB system.

For more information about:

- Add, edit, or remove a security attribute after a DB system is created, see *Manage Security Attributes for the DB System*.
- Defining appropriate ZPR policies, see *Zero Trust Packet Routing*.
- Creating a security attribute, see Creating a Security Attribute.

Tags

Expand **Tags** to provide tag details for this resource.

You can add free-form tags or defined tags to this resource. You must have permission to use the tag namespace for defined tags. For information about using tags to manage your OCI resources, see Resource Tags.

> **Tip**
>
> Tags can be applied later or at any time to a resource.

Database basic details

- **Database name**: Enter a name for the database. It is also known as the `DB_NAME`. The database name must begin with an alphabetic character and can contain a maximum of eight alphanumeric characters. Special characters are not permitted.
- **Database unique name suffix**: *Optional.* The second portion of the database's unique name. The complete database unique name is created by appending the database unique name suffix to the database name you specify.
- **Database unique name**: Read-only. Displays the complete database unique name (`DB_UNIQUE_NAME`). The database unique name is a globally unique name for the database. Primary and standby databases in a Data Guard association can share the same database name but must have different database unique names.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Administrator credentials

- **Username**: sys (This is a read-only field). A database administrator named sys will be created with the password you supply.
- **Password**: Enter a password for the administrator. The password must meet the following criteria:
  - A strong password for SYS, SYSTEM, TDE wallet, and PDB administrator.
  - The password must be 9 to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters.
  - The special characters must be _, #, or -.
  - The password must not contain the user name (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reversed order and regardless of casing.
- **Confirm password**: Reenter the password you specified.

Advanced options

Expand **Advanced options** to provide the advanced options for this resource.

Tags

Expand **Tags** to provide tag details for this resource.

You can add free-form tags or defined tags to this resource. You must have permission to use the tag namespace for defined tags. For information about using tags to manage your OCI resources, see Resource Tags.

> **Tip**
>
> Tags can be applied later or at any time to a resource.

Clone

- Select **Clone**.

Manage Licenses on a DB System

This article provides the procedure to change the license type of a DB system.

Procedure

Perform the following steps to change the license type of a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, from the **Actions** menu, select **Update license type**.
3. In the **Update license type** panel, select the desired license type.
   - **License type**: The type of license you want to use for the DB system. Your choice affects metering for billing.

Copyright © 2022, 2026, Oracle and/or its affiliates.

  - **License included** means the cost of this OCI Database service resource will include both the Oracle Database software licenses and the service.
  - **Bring Your Own License (BYOL)** means you will use your organization's Oracle Database software licenses for this OCI Database service resource. For more information, see Bring Your Own License.

4. Select **Save** and confirm.

Move a DB System to Another Compartment

This article provides the details and procedure to move a DB system between the compartments.

Prerequisites

- To move resources between compartments, resource users must have sufficient access permissions on the compartment that the resource is being moved to, as well as the current compartment. For more information about permissions for database resources, see *Policy Details for Base Database Service*.
- If your DB system is in a security zone, the destination compartment must also be in a security zone. For a full list of policies that affect database service resources, see Security Zone Policies.

Procedure

Perform the following steps to move a DB system between the compartments using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, from the **Actions** menu, select **Move resource**.
3. In the **Move resource** panel, select a **Destination compartment** from the list.
4. Select **Move resource**.

Manage Tags for the DB System

This article provides the procedure to manage tags for the DB system.

Procedure

Perform the following steps to manage tags for the DB systems using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, from the **Actions** menu, select **Add tags**.
3. Alternatively, on the DB system details page, select the **Tags** tab to view the list of existing tags.
   - To add more tags to this resource, select **Add**.
   - To delete an existing tag, select the checkbox next to the tag and then select **Actions** and then select **Delete**.
   - To edit an existing tag, select the checkbox next to the tag and then select **Actions** and then select **Edit**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

You can add free-form tags or defined tags to this resource. You must have permission to use the tag namespace for defined tags. For information about using tags to manage your OCI resources, see Resource Tags.

> **Tip**
>
> Tags can be applied later or at any time to a resource.

Terminate a DB System

This article provides the details and procedure to terminate a DB system.

> **Caution**
>
> Terminating a DB system permanently deletes it and any databases running on it.

General Information

- The database data is local to the DB system and will be lost when the system is terminated. Oracle recommends that you back up any data in the DB system prior to terminating it.
- Terminating a DB system removes all automatic incremental backups of all databases in the DB system from the Recovery Service and Object Storage. Full backups remain in the Recovery Service and Object Storage as standalone backups which you can use to create a new DB system. For information on creating a new DB system from a backup, see *Create a DB System from a Backup*.
- If your DB system has Data Guard enabled, you must terminate the standby DB system before terminating the primary DB system. If you try to terminate a primary DB system that has a standby, the terminate operation will not complete. For more information on Data Guard, see *Use Oracle Data Guard on a DB System*.

Procedure

> **Caution**
>
> Immediately after you complete this procedure, you cannot connect to the system, and any open connections will be terminated.

Perform the following steps to terminate a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, from the **Actions** menu, select **Terminate** and confirm.
   - The database system's icon indicates **Terminating**.

Stop a DB System

This article provides the details and procedure to stop a DB system.

Before you begin, review the following:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- DB system nodes are stopped individually.
- For multi-node RAC DB systems, you might need to stop only one node, for example, when proactively rebooting a node for scheduled maintenance.
- Stopping a node stops billing for all OCPUs/ECPUs associated with that node. Billing resumes when the node is restarted.
- Stopping a DB system stops ECPU billing. Infrastructure and storage charges for the VM.BaseDB.x86 shape continue until the DB system is terminated.
- After a node is restarted or rebooted, the floating IP address might take several minutes to update and appear in the Console.

Procedure

Perform the following steps in the Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, select the **Nodes** tab to view the list of nodes.
3. From the **Actions** menu of the desired node, select **Stop**.
   This action shuts down the node. After the node is powered off, the **Start** action is enabled.

Start a DB System

This article provides the details and procedure to start a DB system.

The DB system nodes are started individually. For multi-node DB systems, you may need to act on only one node (as in the case of proactively rebooting a virtual machine node with scheduled maintenance).

Procedure

Perform the following steps to start a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, select the **Nodes** tab to view the list of nodes.
3. From the **Actions** menu of the desired node, select **Start**.
   This action restarts a stopped node. After the node is restarted, the **Stop** action is enabled.

> **Note**
>
> After you restart or reboot a node, the floating IP address might take several minutes to be updated and displayed in the OCI Console.

Reboot a DB System

This article provides the details and procedure to reboot a DB system.

The DB system nodes are rebooted individually. For multi-node DB systems, you may need to act on only one node (as in the case of proactively rebooting a virtual machine node with scheduled maintenance).

Copyright © 2022, 2026, Oracle and/or its affiliates.

Procedure

Perform the following steps to reboot a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, select the **Nodes** tab to view the list of nodes.
3. From the **Actions** menu of the desired node, select **Reboot**.
   This action shuts down the node and then restarts it.

> **Note**
>
> After you restart or reboot a node, the floating IP address might take several minutes to be updated and displayed in the OCI Console.

View Work Requests for the DB System

This article provides the details and procedure to view work requests for a DB system.

You can view the log messages, error messages, work request information, and status details on the work request page.

Procedure

Perform the following steps to view work requests for a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, select the **Work requests** tab to view the list of all work requests.

For more information about work requests, see Work Requests.

Oracle True Cache

Oracle True Cache is an in-memory, consistent, and automatically managed SQL and key-value (object or JSON) cache for Oracle Database.

True Cache is an in-memory, read-only cache in front of an Oracle database.

Like Oracle Active Data Guard, True Cache is a fully functional, read-only replica of the primary database, except that it's mostly disk less.

Modern applications often require a high number of connections and fast, low-latency access to the data. A popular approach is to place caches in front of the database because applications typically perform many more reads than updates, and they can read from the cache without affecting the database performance. For example, an airline reservation system reads data frequently as people shop for flights, and it's okay if the data isn't the most current compared to what's in the database. The application only needs to access the most current data when someone reserves a flight.

Unlike conventional caches, True Cache automatically keeps the most frequently accessed data in the cache, and it keeps the cache consistent with the primary database, other objects in

Copyright © 2022, 2026, Oracle and/or its affiliates.

the same cache, and other caches. It caches all Oracle Database objects and data types, including JSON.

Because True Cache implements Oracle Database security policies, you can control access to the cache. This enables you to cache sensitive data, like private personal information, that you might not ordinarily cache.

For detailed information, see Overview of Oracle True Cache.

Migrate to a Dual Stack Network

This article provides details about migrating a DB system with an IPv4-only network to IPv4/IPv6 dual stack network using Data Guard.

Perform the following steps to migrate a DB system with IPv4-only network to IPv4/IPv6 dual stack network using Data Guard.

1. Set up a Data Guard association between an existing DB system with IPv4-only network and a new DB system with dual stack network. Initially, the database in the IPv4-only network will be the primary database, and the database in the dual stack network will be the secondary database. For detailed steps, see *Enable Oracle Data Guard on a DB System*.
2. Perform a database switchover. After the database switchover, the database in the dual stack network will become the primary database. For detailed steps, see *Perform Database Switchover and Failover*.

   > **Note**
   >
   > Ensure that both databases are synchronized before performing the switchover. The primary and standby databases should be synchronized, meaning there should be no backlog of redo logs on the standby database. You can check synchronization by ensuring that no redo logs are pending to be applied on the standby. Use the `DGMGRL> SHOW CONFIGURATION;` command to check the synchronization status.

3. Once you are satisfied and sure that the new DB system in the dual stack network is working fine, you can terminate the Data Guard association by deleting the current standby DB system in the IPv4-only network. For detailed steps, see *Terminate an Oracle Data Guard Association on a DB System*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Connect

Overview of Connecting to a DB System

This article provides an introduction about various settings required to connect to an active DB system. How you connect depends on the client tool or protocol you use, the purpose of the connection, and how your cloud network is set up.

> **Note**
>
> You can find information on various networking scenarios in Networking Overview, but for specific recommendations on how you should connect to a database in the cloud, contact your network security administrator.

Prerequisites

- To use the Console or the API to get the default administration service connection strings, you must be given the required type of access in a **policy** written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which **compartment** you should work in.
- To connect to the database, you'll need the public or private IP address of the DB system.

  Use the private IP address to connect to the system from your on-premises network, or from within the Virtual Cloud Network (VCN). This includes connecting from a host located on-premises connecting through a VPN or FastConnect to your VCN, or from another host in the same VCN. Use the public IP address to connect to the system from outside the cloud (with no VPN). You can find the IP addresses in the Console as follows:

  – On the DB System details page, select the **Nodes** tab.

  – The **DNS name**, **Public IP address**, and **Private IP address** are displayed in the table columns.
- For Secure Shell (SSH) access to the DB system, you'll need the full path to the file that contains the private key associated with the public key used when the DB system was launched.

If you have problems connecting, see *Troubleshoot Connection Issues*.

Database Services and Connection Strings

Database services allow you to control client access to a database instance depending on the functionality needed. For example, you might need to access the database for administration purposes only or you might need to connect an application to the database. Connection strings are specific to a database service.

When you provision a DB system, a default database administration service is automatically created. For 12c and later Oracle Databases, this service is for administrating the database at the CDB level. Because this service provides limited functionality, it is not suitable for connecting an application. Oracle recommends that you create a default application service for the initial database after you create your DB system. For 12c and later Oracle Databases,

Copyright © 2022, 2026, Oracle and/or its affiliates.

application services connect at the PDB level. Here are some important functions an application service can provide:

- Workload identification
- Load balancing
- Application continuity and Transaction Guard
- Fast Application Notification
- Resource assignment based on the service name

For details about these and other High Availability capabilities, see Client Failover Best Practices for Highly Available Oracle Databases.

Create an Application Service

You use the `srvctl` utility to create an application service. Before you can connect to the service, you must start it.

You can create an application service for a PDB or an 11g Oracle database using the following steps.

1. Log in to the DB system host as `opc`.
2. Switch to the oracle user, and set your environment to the Oracle Database you want to administer.

```
sudo su - oracle
. oraenv


ORACLE_SID = [oracle] ? <database_name>
The Oracle base has been set to /u01/app/oracle
```

3. Create the application service for the database. Include the `pdb` option only if you are creating an application service for a PDB.

```
srvctl add service
  -db <DB_unique_name>
  -pdb <PDB_name>
  -service <app_service_name>
  -role PRIMARY
  -notification TRUE
  -session_state dynamic
  -failovertype transaction
  -failovermethod basic
  -commit_outcome TRUE
  -failoverretry 30
  -failoverdelay 10
  -replay_init_time 900
  -clbgoal SHORT
  -rlbgoal SERVICE_TIME
  -preferred <rac_node1>,<rac_node2>
  -retention 3600
```

Note that the preferred option is required only for multi-node databases to specify the hostname of the node in the RAC.

Copyright © 2022, 2026, Oracle and/or its affiliates.

4. Start the application service.

```
srvctl start service -db <DB_unique_name> -s <app_service_name>
```

For more information about services for a PDB, see Administering PDBs.

Database Connection Strings

You must use the appropriate connection string to access a database administration or application service. You can use the Console or the API to get the string for connecting to the default administration service from within a VCN. For 12c and later Oracle Databases, this service is for administrating the database at the CDB level. The string is provided in both the Easy Connect and in the full connect descriptor (long) format. Use the long format for the connection if hostname resolution is not available. You can also use the long format to create an alias in the tnsnames.ora file.

For accessing a database service within the VCN, the connection string for a Real Application Cluster (RAC) DB system uses the Single Client Access Name (SCAN) while the connection string for single instance DB system uses the hostname instead.

The private SCAN name is a Round Robin DNS entry created when you launch a 2-node RAC DB system. The private SCAN name is resolvable only within the VCN. If the client and the database are in the same VCN, the connection mechanism is the same as an on-premises RAC database; all the features provided by VIPs and SCAN VIPs, such as server side load balancing and VIP failover, are available.

> **Note**
>
> If you manually change the DB_UNIQUE_NAME, DB_DOMAIN, or listener port on the DB system, the connection strings you see in the Console or API will not reflect your changes. Ensure that you use the actual values of these parameters when you make a connection.

Get the Connection Strings

You can get the connection strings for the default administration service using the following steps using the Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, select the **Connections** tab. A list of connection strings is displayed.
4. Select the applicable link to **view** or **copy** the connection string.

You can derive the connection strings for other database services by replacing part of the default application service connection string with the applicable values.

Derive the Connection String

Perform the following steps to derive the connection string for a PDB administration service or an application service.

Copyright © 2022, 2026, Oracle and/or its affiliates.

1. Follow the procedure to get the Easy Connect string for the default administration service. That string should have the following format:

   `<hostname|SCAN>:1521/<DB_unique_name>.<DB_domain>`

2. Make the appropriate substitution:
   - For the PDB administration service, replace `DB_unique_name` with the PDB name.

     `<hostname|SCAN>:1521/<PDB_name>.<DB_domain>`

   - For an application service, replace `DB_UNIQUE_NAME` with the name of the application service.

     `<hostname|SCAN>:1521/<app_service_name>.<DB_domain>`

Connect to a Database by Using SQLNet

This section describes how to connect to a database service from a computer that has a SQL*Net client installed. Port 1521 must be open to support the SQL*Net protocol.

Connect From Within the VCN

For security reasons, Oracle recommends that you connect to your database services from within the VCN. You can use this method whether you are connecting to an administration service or to an application service.

To connect using SQL*Plus, you run the following command using the applicable connection string:

```
sqlplus system/<password>@<connection_string>
```

Consider the following:

- If your system is not using the VCN Resolver, ensure that the DB system's hostname (for single-node systems) or SCAN name (for multi-node systems) can be resolved. For information about DNS name resolution, see DNS in Your Virtual Cloud Network.
- For connecting to the administration service of a PDB, ensure that the PDB is open or the service will not be available.
- For connecting to an application service, ensure that the service is started. For Fast Application Notification to work, ensure that port 6200 can be reached. For information about Fast Application Notification, see Client Failover Best Practices for Highly Available Oracle Databases.

Connect From the Internet

Although Oracle does not recommend connecting to your database from the Internet, you can connect to a database service by using a public IP address if port 1521 is open to the public for ingress.

To use this method, you run the following command using the public IP address instead of the hostname or SCAN in the connection string:

```
sqlplus system/<password>@<public_IP>:1521/<service_name>.<DB_domain>
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Consider the following:

- SCANs and hostnames are not resolvable on the Internet, therefore load balancing and failover for multi-node DB systems, which rely on these names, cannot work.
- For multi-node DB systems, which normally use SCANs, you must specify the IP address of one of the RAC hosts to access the database.

> **Note**
>
> Do not use this method to connect to the database from within the VCN. Doing so negatively impacts performance because traffic to the database is routed out of the VCN and back in through the public IP address.

Example: Connecting in SQL Developer Using SQL*Net

Prerequisites:

- Ensure that port 1521 is open for the Oracle default listener. (You can do this by checking the DB system's security list.)
- If port 1521 is open only to hosts in the VCN, then you must run your SQL Developer client from a machine that has direct access to the VCN. If you are connecting to the database from the Internet instead, then the public IP address of your computer must be granted access to port 1521 in the security list. (Alternatively, the security list can grant full access to port 1521, however, this is not recommended for security reasons.) You must use the public IP address of the host because connecting from the Internet does not support SCAN name resolution.

Connect From Within the VCN Using a Private IP Address

After the prerequisites are met, start SQL Developer and create a connection by supplying the following connection details:

- **Username:** sys as sysdba
- **Password:** The **Database Admin Password** that was specified in the **Launch DB System** dialog in the Console.
- **Hostname:** The hostname as it appears in the Easy Connect format of the connection string.For help with getting the connection string and identifying the hostname, see *Overview of Connecting to a DB System*.
- **Port:** 1521
- **Service name:** The concatenated name of the service and host domain name, for example, db1_phx1tv.example.com. You can identify this value as the last part of the Easy Connect string, `<service_name>.<DB_domain>`.

Connect to a Database with a Public IP by Using SSH Tunneling

You can access the services of DB system databases with public IP addresses by using SSH tunneling.

The main advantage of this method is that port 1521 does not need to be opened to the public internet. However, just like accessing the database with a public IP using a SQL*Net client, load balancing and failover for multi-node DB systems cannot work because they rely on SCANs and hostnames.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Oracle SQL Developer and Oracle SQLcL and are two tools that facilitate the use of tunneling for Oracle Database access.

To open a tunnel, and then connect to a database service by using SQLcL, you run the following commands:

```
sshtunnel opc@<public_IP> -i <private_key> -L <local_port>:<private_IP>:1521

connect system/<password>@localhost:<local_port>/<service_name>.<DB_domain>
```

For more information about these tools, see Oracle SQL Developer and Oracle SQLcL.

Connect to a Database By Using SSH and the Bequeath Protocol

This method allows you to connect to the database without using the network listener. It should be used to connect only for administration purposes.

When connecting to a multi-node DB system, you'll SSH to each individual node in the cluster.

Connect From a UNIX-style System

Use the following SSH command to access the DB system:

```
ssh –i <private_key> opc@<DB_system_IP_address>
```

`<private_key>` is the full path and name of the file that contains the private key associated with the DB system you want to access.

Use the DB system's private or public IP address depending on your network configuration.

For more information, see *prerequisites* in *Overview of Connecting to a DB System*.

Connect From a Windows System

1. Open `putty.exe`.
2. In the **Category** pane, select **Session** and enter the following fields:
   - **Host Name (or IP address):** `opc@<DB_system_IP_address>`

     Use the DB system's private or public IP address depending on your network configuration.
   - **Connection type:** SSH
   - **Port:** 22
3. In the **Category** pane, expand **Connection**, expand **SSH**, and then click **Auth**, and browse to select your private key.
4. Optionally, return to the **Session** category screen and save this session information for reuse later.
5. Click **Open** to start the session.

For more information, see *prerequisites* in *Overview of Connecting to a DB System*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Access a Database After You Connect

1. Log in as `opc`.

   ```
   login as: opc
   ```

2. `sudo` to the grid user.

   ```
   sudo su - grid
   ```

3. List all the databases on the system.

   ```
   srvctl config database -v
   ```

   Output:

   ```
   cdbm01 /u02/app/oracle/product/12.1.0/dbhome_2 12.1.0.2.0
   exadb /u02/app/oracle/product/11.2.0/dbhome_2 11.2.0.4.0
   mmdb /u02/app/oracle/product/12.1.0/dbhome_3 12.1.0.2.0
   ```

4. Connect as the oracle user.

   ```
   [root@ed1db01 ~]# su - oracle
   [oracle@ed1db01 ~]$ . oraenv
   ORACLE_SID = [oracle] ? cdbm01
   The Oracle base has been set to /u02/app/oracle
   ```

5. Get the details about one of the databases by using the `srvctl` command.

   ```
   srvctl config database -d cdbm01
   ```

   Output:

   ```
   Database unique name: cdbm01 <<== DB unique name
   Database name:
   Oracle home: /u02/app/oracle/product/12.1.0/dbhome_2
   Oracle user: oracle
   Spfile: +DATAC1/cdbm01/spfilecdbm01.ora
   Password file: +DATAC1/cdbm01/PASSWORD/passwd
   Domain: data.customer1.oraclevcn.com
   Start options: open
   Stop options: immediate
   Database role: PRIMARY
   Management policy: AUTOMATIC
   Server pools:
   Disk Groups: DATAC1,RECOC1
   Mount point paths:
   Services:
   Type: RAC
   Start concurrency:
   Stop concurrency:
   OSDBA group: dba
   OSOPER group: racoper
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
Database instances: cdbm011,cdbm012 <<== SID
Configured nodes: ed1db01,ed1db02
Database is administrator managed
```

6. Set the `ORACLE_SID` and `ORACLE_UNIQUE_NAME` using the values from the previous step.

```
export ORACLE_SID=cdbm011
export ORACLE_UNIQUE_NAME=cdbm01
sqlplus / as sysdba


SQL*Plus: Release 12.1.0.2.0 Production on Wed Apr 19 04:10:12 2017

Copyright (c) 1982, 2014, Oracle. All rights reserved.

Connected to:
Oracle Database 12c EE Extreme Perf Release 12.1.0.2.0 - 64bit Production
With the Partitioning, Real Application Clusters, Automatic Storage
Management, Oracle Label Security,
OLAP, Advanced Analytics and Real Application Testing options
```

Troubleshoot Connection Issues

The following issues might occur when connecting to a DB system or database.

ORA-28365: Wallet is Not Open Error

For a 1-node DB system or 2-node RAC DB system, regardless of how you connect to the DB system, *before* you use OS authentication to connect to a database (for example, `sqlplus / as sysdba`) be sure to set the ORACLE_UNQNAME variable. Otherwise, commands that require the TDE wallet will result in the error `ORA-28365: wallet is not open`.

> **Note**
>
> This is not an issue when using a TNS connection because ORACLE_UNQNAME is automatically set in the database CRS resource.

SSH Access Stops Working

If the DB system's root volume becomes full, you might lose the ability to SSH to the system (the SSH command will fail with permission denied errors). Before you copy a large amount of data to the root volume, for example, to migrate a database, use the `dbcli create-dbstorage` command to set up storage on the system's NVMe drives and then copy the database files to that storage. For more information, see Setting Up Storage on the DB System.

Manage Serial Console Connection to the DB System

This article provides the details and procedure to create and delete serial console connections to a DB system.

Serial console connection enables you to manage and troubleshoot the DB system in single-user mode using an SSH connection.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Create a Serial Console Connection to the DB System

Procedure

Perform the following steps to create a serial console connection using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, select the **Console connections** tab to view the list of existing console connections.
3. Select **Create console connection**.
   - This is disabled if all nodes have existing console connections.
4. In the **Create console connection** panel, enter the following details:
   - **DB system node**: For multi-node DB systems, select which node or nodes for which you want to create a connection. No node selector will display if the DB system has only one node or if there is only one node in a multi-node system that currently lacks a connection.
   - **SSH keys**: Add the public key portion of each key pair you want to use for SSH access. Select one of the following options:
     - **Generate SSH key pair**: Use this option to create a new SSH key pair. Select both **Save private key** and **Save public key** when using this option. The private key is downloaded to your local system and must be stored in a safe location. You cannot download another copy of the private key generated during this operation after completing the operation.
     - **Upload SSH key files**: Select this option to browse or drag and drop your existing public key (.pub) files.
     - **Paste SSH keys**: Select this option to paste in individual public keys. To paste multiple keys, select **+ Another SSH key** and supply a single key for each entry.
5. Select **Create**.

Delete a Serial Console Connection to your DB System

Procedure

Perform the following steps to delete a serial console connection using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, select the **Console connections** tab to view the list of existing console connections.
3. From the **Actions** menu of the connection you want to delete, select **Delete**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Monitor

Monitor Base Database Service

You can monitor the health, capacity, and performance of your DB systems and databases with metrics, alarms, and notifications. You can use the OCI Console, Monitoring APIs, or Database Management APIs to view metrics.

Monitoring Service

You can use the metrics feature in the Monitoring service to monitor Oracle Cloud resources. These metrics are available by **default** for Base Database Service resources.

- For a complete list of available metrics, see *Available Metrics for Base Database Service Resources*.
- For detailed instructions about viewing metrics, see *View Metrics for Base Database Service Resources*.
- For more information about the Monitoring service, see Overview of Monitoring.

Database Management

You can use the Database Management service to monitor and manage Oracle Databases. You must enable Database Management for Base Database Service resources before using it.

- For more information about monitoring using Database Management, see *Monitor Using Database Management Service*.
- For more information about enabling, disabling, or editing Database Management, see *Manage Database Management for Base Database Service Resources*.
- For a complete list of available metrics, see Oracle Cloud Database Metrics.
- For detailed instructions about viewing metrics, see *View Metrics for Base Database Service Resources*.
- For more information about Database Management, see Database Management for Oracle Databases.

Performance Hub

You can use Performance Hub to monitor Oracle Databases for defined time periods and download statistical reports. You must enable Database Management for Base Database Service resources before using the Performance Hub.

- For detailed instructions about viewing Performance Hub metrics, see *View Performance Hub Metrics for Base Database Service Resources*.
- For more information about Performance Hub, see About Performance Hub.

Enterprise Manager

You can use Enterprise Manager to manage and monitor the Base Database Service.

- For more information about monitoring using Enterprise Manager, see *Monitor Using Oracle Enterprise Manager*.
- For more information about Enterprise Manager, see About Enterprise Manager Cloud Control 13c.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Available Metrics for Base Database Service Resources

This article describes the metrics emitted by the Base Database service in the `oci_database_cluster` and `oci_database namespaces`.

Dimensions

All the metrics discussed in this article include the following dimensions.

- RESOURCEID - The OCID of the DB system.
- RESOURCENAME - The name of the DB system.

> **Note**
>
> In addition to the metrics listed below, Oracle analyzes additional metrics to provide the highest level of service operations and support for ensuring high availability of services.

Metrics for the DB System in the `oci_database_cluster` Namespace

The metrics listed in the following table are automatically available for the DB system.

Table 5-1 Metrics in oci_database_cluster Namespace

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions	Whether Visible in the DB System Details page
`ASMDiskgroupUtilization`	**ASM Diskgroup Utilization**	percentage	Percentage of usable space used in a Disk Group. Usable space is the space available for growth. DATA disk group stores our Oracle database files. RECO disk group contains database files for recovery such as archives and flashback logs.	10 minutes	hostName deploymentType diskgroupName	Yes
`FilesystemUtilization`	**Filesystem Utilization**	percentage	Percent utilization of provisioned filesystem.	1 minute	hostName deploymentType filesystemName	No
`CpuUtilization`	**CPU Utilization**	percentage	Percent CPU utilization.	1 minute	hostName deploymentType	Yes

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-1 (Cont.) Metrics in oci_database_cluster Namespace

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions	Whether Visible in the DB System Details page
`MemoryUtilization`	**Memory Utilization**	percentage	Percentage of memory available for starting new applications, without swapping. The available memory can be obtained via the following command: `cat/proc/meminfo`.	1 minute	hostName deploymentType	Yes
`SwapUtilization`	**Swap Utilization**	percentage	Percent utilization of total swap space.	1 minute	hostName deploymentType	Yes
`LoadAverage`	**Load Average**	integer	System load average over 5 minutes.	1 minute	hostName deploymentType	Yes
`NodeStatus`	**Node Status**	integer	Indicates whether the host is reachable in RAC environments.	1 minute	hostName deploymentType	Yes
`OcpusAllocated`	**OCPU Allocated**	integer	The number of OCPUs allocated.	1 minute	deploymentType	No

> **Note**
>
> Some of the above metrics are not displayed on the **DB System Details** page. Instead, you can view them using the Monitoring service.

Metrics for the Database in the `oci_database` Namespace

The metrics listed in the following table are automatically available for the database.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-2 Metrics in oci_database Namespace

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions	Whether Visible in the DB System Details page
`CpuUtilization`	**CPU Utilization**	percentage	The 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 two times the number of OCPUs.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database|pdb} resourceName_{database|pdb}	Yes
`StorageUtilization`	**Storage Utilization**	percentage	The percentage of provisioned storage capacity currently in use. Represents the total allocated space for all tablespaces.	1 hour	deploymentType resourceId_{database|pdb} resourceName_{database|pdb}	Yes
`BlockChanges`	**DB Block Changes**	Changes per second	The Average number of blocks changed per second.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database|pdb} resourceName_{database|pdb}	Yes
`ExecuteCount`	**Execute Count**	Count	The number of user and recursive calls that executed SQL statements during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType	Yes

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-2 (Cont.) Metrics in oci_database Namespace

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions	Whether Visible in the DB System Details page
`CurrentLogons`	**Current Logons**	Count	The number of successful logons during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database|pdb} resourceName_{database|pdb}	Yes
`TransactionCount`	**Transaction Count**	Count	The combined number of user commits and user rollbacks during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database|pdb} resourceName_{database|pdb}	No
`UserCalls`	**User Calls**	Count	The combined number of logons, parses, and execute calls during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database|pdb} resourceName_{database|pdb}	No

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-2 (Cont.) Metrics in oci_database Namespace

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions	Whether Visible in the DB System Details page
`ParseCount`	**Parse Count**	Count	The number of hard and soft parses during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database|pdb} resourceName_{database|pdb}	Yes
`StorageUsed`	**Storage Space Used**	GB	Total amount of storage space used by the database at the collection time.	1 hour	deploymentType resourceId_{database|pdb} resourceName_{database|pdb}	No
`StorageAllocated`	**Storage Space Allocated**	GB	Total amount of storage space allocated to the database at the collection time.	1 hour	deploymentType resourceId_{database|pdb} resourceName_{database|pdb}	No
`StorageUsedByTablespace`	**Storage Space Used By Tablespace**	GB	Total amount of storage space used by tablespace at the collection time. In case of container database, this metric provides root container tablespaces.	1 hour	tablespaceName, tablespaceType deploymentType resourceId_{database|pdb} resourceName_{database|pdb}	No

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-2 (Cont.) Metrics in oci_database Namespace

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions	Whether Visible in the DB System Details page
StorageAllocatedByTablespace	**Allocated Storage Space By Tablespace**	GB	Total amount of storage space allocated to the tablespace at the collection time. In case of container database, this metric provides root container tablespaces.	1 hour	tablespaceName, tablespaceType deploymentType resourceId_{database|pdb} resourceName_{database|pdb}	No
StorageUtilizationByTablespace	**Storage Space Utilization By Tablespace**	percentage	This indicates the percentage of storage space utilized by the tablespace at the collection time. In case of container database, this metric provides root container tablespaces.	1 hour	tablespaceName, tablespaceType deploymentType	No

> **Note**
>
> Some of the above metrics are not displayed on the **Database Details** page. Instead, you can view them using the Monitoring service.

View Metrics for Base Database Service Resources

This article provides the details and procedure to view metrics for the Base Database Service resources.

The following topics are covered in this article:

- General Information
- Prerequisites for Viewing Metrics
- View Metrics for a DB System
- View Metrics for a Database
- View Metrics for a Database in a Compartment
- View Metrics for a Pluggable Database

Copyright © 2022, 2026, Oracle and/or its affiliates.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be granted security access in a **policy** by an administrator. This access is required whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you get a message that you don't have permission or are unauthorized, verify with your administrator what type of access you have and which **compartment** to work in.

For administrators: The policy in Let database admins manage Oracle Cloud database systems lets the specified group do everything with databases and related Database resources.

If you're new to policies, see Getting Started with Policies and Common Policies. If you want to dig deeper into writing policies for databases, see Details for the Database Service.

General Information

- By default, the metrics for the last one hour are displayed.
- By default, the metrics from the `oci_database namespace` are displayed. If Database Management is enabled, then the metrics from `oracle_oci_database` are displayed. To enable Database Management for databases, see Enable Database Management for a Database.
- When there is a network problem and Oracle Trace File Analyzer (TFA) is unable to post metrics, TFA will wait for one hour before attempting to retry posting the metrics. This is required to avoid creating a backlog of metrics processing on TFA.
- Potentially one hour of metrics will be lost between network restore and the first metric posted.
- If you don't see any metrics, check the network settings and AHF version listed in the prerequisites section.

> **Note**
>
> **Known Issue**: When the DB System is deployed using the Logical Volume Manager storage management software, there may be missing metric collections, resulting in graphs with missing data points.

Prerequisites for Viewing Metrics

The following prerequisites are necessary for the DB system to generate metrics.

1. Metrics on the DB system depends on TFA agent. Ensure that these components are up and running. AHF version **22.3.3** or higher is required for capturing metrics from the DB system. To start, stop, or check the status of TFA, see Manage Oracle Trace File Analyzer.
2. The following network configurations are required.
    a. **Egress rules for outgoing traffic**: The default egress rules are sufficient to enable the required network path: For more information, see *Default Security Lists* in Security Lists. If you have blocked the outgoing traffic by modifying the default egress rules on your Virtual Cloud Network (VCN), you will need to revert the settings to allow outgoing traffic. The default egress rule allowing outgoing traffic (as shown in the *Security Rules for the DB System* article) is as follows:
        - Stateless: No (all rules must be stateful)

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Destination Type: CIDR
- Destination CIDR: **All <region> Services in Oracle Services Network**
- IP Protocol: TCP
- Destination Port: 443 (HTTPS)

b. **Public IP or Service Gateway**: The database server host must have either a public IP address or a service gateway to be able to send database server host metrics to the Monitoring service.

   If the instance does not have a public IP address, set up a service gateway on the VCN. The service gateway lets the instance send database server host metrics to the Monitoring service without the traffic going over the internet. Here are special notes for setting up the service gateway to access the Monitoring service:

   i. When creating the service gateway, enable the service label called **All <region> Services in Oracle Services Network**. It includes the Monitoring service.

   ii. When setting up routing for the subnet that contains the instance, set up a route rule with **Target Type** set to **Service Gateway**, and the **Destination Service** set to **All <region> Services in Oracle Services Network**.

   For detailed instructions, see Access to Oracle Services: Service Gateway.

View Metrics for a DB System

Perform the following steps to view the metrics for a DB system using the Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, select the **Monitoring** tab.
3. A chart for each metric is displayed.
   - By default, the metrics for the last one hour are displayed.
   - You can apply the **Metric namespace** and **Date filter** to select the required start time and end time. The metrics are refreshed immediately for the selected filter.
   - For each metric, you can select the interval and statistic independently.
     - **Interval**: The time period for which the metric is calculated.
     - **Statistic**: The mathematical method for which the metric is calculated.
   - For each metric, you can select the following options from the **Actions** menu of that metric:
     - View query in Metrics Explorer
     - Copy chart URL
     - Copy query (MQL)
     - Create an alarm on this query
     - Table view

> **Note**
>
> If you don't see any metrics, check the network settings and AHF version listed in the prerequisites section.

Copyright © 2022, 2026, Oracle and/or its affiliates.

View Metrics for a Database

Perform the following steps to view the metrics for a database using the Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, select the **Monitoring** tab.
   - By default, the metrics from the `oci_database namespace` are displayed. If Database Management is enabled, then the metrics from `oracle_oci_database` are displayed. To enable Database Management for databases, see Enable Database Management for a Database.
4. A chart for each metric is displayed.
   - By default, the metrics for the last one hour are displayed.
   - You can apply the **Metric namespace** and **Date filter** to select the required start time and end time. The metrics are refreshed immediately for the selected filter.
   - For each metric, you can select the interval and statistic independently.
     - **Interval**: The time period for which the metric is calculated.
     - **Statistic**: The mathematical method for which the metric is calculated.
   - For each metric, you can select the following options from the **Actions** menu of that metric:
     - View query in Metrics Explorer
     - Copy chart URL
     - Copy query (MQL)
     - Create an alarm on this query
     - Table view

> **Note**
>
> If you don't see any metrics, check the network settings and AHF version listed in the prerequisites section.

View Metrics for a Database in a Compartment

Perform the following steps to view the metrics for databases in a compartment using the Console:

1. From the navigation menu, select **Observability & Management**, and then under **Monitoring**, select **Service Metrics**.
2. On the service metrics page, select a **Compartment**.
3. Select **`oci_database`** under the **Metric namespace**.
4. If there are multiple databases in the compartment, you can show metrics aggregated across the databases by selecting **Aggregate metric streams**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

5. If you want to limit the metrics you see, next to **Dimensions** select **Add** (select **Edit** if you have already added dimensions).
   - In the **Dimension name**, select a dimension.
   - In the **Dimension Value**, select a value.
   - Select **Done**.
   - In the **Edit dimensions** dialog, select **+Additional dimension** to add an additional dimension.
     - If not required, select **x** to remove a dimension.
6. To create an alarm on a specific metric, select **Options** and select **Create an alarm on this query**. See Managing Alarms for information on setting and using alarms.

> **Note**
>
> If you don't see any metrics, check the network settings and AHF version listed in the prerequisites section.

View Metrics for a Pluggable Database

Prerequisites

- The Database Management for databases with **Full Management** option must be enabled. To enable Database Management for databases, see Enable Database Management for a Database.
- The Database Management for pluggable databases must be enabled. To enable Database Management for pluggable databases, see Enable Database Management for a Pluggable Database.

Procedure

Perform the following steps to view the metrics for pluggable databases using the Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the pluggable database details page, select the **Metrics** tab. The metrics for the database are displayed.

Monitor Using Database Management Service

You can use the Database Management Service to manage and monitor the health of the Oracle Databases in the Base Database Service.

Using Database Management, you can monitor single instance and RAC databases, which include Container Databases (CDBs), Pluggable Databases (PDBs), and Non-Container Databases (Non-CDBs). Database Management supports Oracle Database version 11.2.0.4 and later.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Using the Database Management Service, you can:

- Monitor the key performance and configuration metrics.
- Compare and analyze database metrics over a selected period.
- Group your resources, which reside across compartments, into a group and monitor them.

Database Management features for Oracle Cloud Databases are available as part of the **Full management** and **Basic management** options. The **Basic management** option is available at no additional cost.

- **Full management**: This option includes all Database Management features at an additional service cost. The **Full management** option is available for the Oracle Database Enterprise Editions and the Oracle Database Standard Edition, however, for the Oracle Database Standard Edition, the **Full management** option does not include Performance Hub features.
- **Basic management**: This option is available for Oracle Cloud Databases at no additional cost.

To enable Database Management for the databases, see *Manage Database Management for Base Database Service Resources*.

After enabling Database Management, you can perform the following actions:

- *View Metrics for Base Database Service Resources*
- *View Performance Hub Metrics for Base Database Service Resources*

For a complete list of available metrics, see Oracle Cloud Database Metrics.

For more information about Database Management, see Database Management for Oracle Databases.

Manage Database Management for Base Database Service Resources

This article describes the procedure to enable, edit, and disable Database Management for Base Database Service resources.

The following topics are covered in this article:

- Enable Database Management for a Database
- Edit Database Management for a Database
- Disable Database Management for a Database
- Enable Database Management for a Pluggable Database
- Edit Database Management for a Pluggable Database
- Disable Database Management for a Pluggable Database

Prerequisites

You must perform the following tasks before enabling Database Management for your databases.

- Obtain the permissions required to enable Database Management as detailed in Permissions Required to Enable Database Management for Oracle Cloud Databases.
- Complete the prerequisite tasks listed in Oracle Cloud Database-related Prerequisite Tasks.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be granted security access in a **policy** by an administrator. This access is required whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you get a message that you don't have permission or are unauthorized, verify with your administrator what type of access you have and which **compartment** to work in.

For administrators: The policy in Let database admins manage Oracle Cloud database systems lets the specified group do everything with databases and related Database resources.

If you're new to policies, see Getting Started with Policies and Common Policies. If you want to dig deeper into writing policies for databases, see Details for the Database Service.

Use the API

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 the following APIs for managing Database Management:

- enableDatabaseManagement
- disableDatabaseManagement
- updateDatabaseManagement

For the complete list of APIs for the Database service, see Database Service API.

Enable Database Management for a Database

> **Note**
>
> You can also enable Database Management for a database from the Database Management Administration page. For more information, see Enable Database Management for Oracle Cloud Databases.

Perform the following steps to enable Database Management for a database using the Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the **Database information** tab, under the **Associated Services** section, check the status of **Database Management**.
4. If the **Database Management** is displayed as **Not Enabled**, perform the following steps to enable Database Management.
5. Select **Enable**.
6. In the **Enable Diagnostics & Management** page, in the **Database information** section, enter the following details.
   - **Database type:** *Read-only*. Type of the database.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Database system**: *Read-only*. Compartment in which the database is located.
- **Database home**: *Read-only*. Database home of the database.
- **Database name**: *Read-only*. Name of the database.
- Select from one of the following options:
  - Enable Diagnostics & Management for CDB and all PDBs.
  - Enable Diagnostics & Management for CDB only.
- **Service name**: The unique service name of the database. A default unique name is displayed, which can be changed if required.
- **Protocol**: Select either the TCP or TCPS protocol to connect to the database. By default, the TCP protocol is selected.

> **Note**
>
> If Oracle Data Guard is enabled after Database Management was enabled for a DB system using the TCPS protocol, then TCPS will have to be reconfigured. Enabling Oracle Data Guard is causing TCPS configuration to be overwritten, and it's recommended that TCPS is configured on a DB system after enabling Oracle Data Guard.

> **Note**
>
> Database Management currently does not support Oracle Data Guard configuration, and Database Management features are not available for standby databases.

- **Port**: Specify the port number. If TCP is selected in the Protocol field, then port number `1521` is displayed by default, and you can change it if required. You can select the port number from a range of 1 to 65535.
- **Database wallet secret**: This field is only displayed if TCPS is selected in the Protocol field.
  Select the secret that contains the database wallet from the drop-down list. If an existing database wallet secret is not available, then select **Create new secret...** in the drop-down list. The **Create database wallet secret** panel is displayed, and you can create a new secret. For information on database wallets and creating a secret in the Vault service, see Oracle Cloud Database-related Prerequisite Tasks.

  If the Database Management (`dpd`) service policy that grants Database Management the permission to read the secret that contains the database wallet is not created, then the *'System policies are required..'* message is displayed. You can select **Add policy** to view and automatically create the service policy. For information on Vault service permissions required to use existing secrets or create new secrets, see Permissions Required to Enable Database Management for Oracle Cloud Databases.

7. In the **Specify credentials for the connection** section, provide the following details.
   - **Database user name**: Enter the database user name.
   - **Database user password secret**:
     Select the secret that contains the database user password from the drop-down list. If the compartment in which the secret resides is different from the compartment displayed, then select **Change compartment** and select another compartment. If an

Copyright © 2022, 2026, Oracle and/or its affiliates.

existing secret with the database user password is not available, then select **Create new secret...** in the drop-down list. The **Create password secret** panel is displayed, and you can create a new secret. For information on database monitoring user credentials and saving the database user password as a secret in the Vault service, see Oracle Cloud Database-related Prerequisite Tasks.

If the Database Management (`dpd`) service policy that grants Database Management the permission to read the secret that contains the database user password is not created, then the *'System policies are required..'* message is displayed. You can select **Add policy** to view and automatically create the service policy. For information on Vault service permissions required to use existing secrets or create new secrets, see Permissions Required to Enable Database Management for Oracle Cloud Databases.

8. In the **Select connection mode** section, select the private endpoint that will act as a representation of Database Management in the VCN in which the database can be accessed. You can choose the private endpoint from a different compartment as well. You must ensure that the appropriate Database Management private endpoint is available. Here are the two types of Database Management private endpoints:
   - Private endpoint for single-instance databases in the DB systems.
   - Private endpoint for RAC databases in the DB system.
   - If a Database Management private endpoint is not available, then you must create one. For information on how to create a private endpoint, see Create a Database Management Private Endpoint.
9. In the **Management options** section, choose between the following options.
   - **Full management**: This includes fleet management, advanced Performance Hub, and other SKU features along with basic management capabilities.
   - **Basic management**: This includes basic monitoring metrics and the ASH Analytics and SQL Monitoring features in Performance Hub for databases.

   For more information on Database Management options, see the *About Management Options* section in Enable Database Management for Oracle Cloud Databases.
10. Select **Enable Diagnostics & Management**.
11. A confirmation message with a link to the **Work requests** section on the **Database information** page is displayed. Select the link to monitor the progress of the work request.
12. On the **Database information** tab, under the **Associated Services** section, verify if the status of **Database Management** is enabled.

If you encounter issues when enabling Database Management, to know about likely causes and solutions, see Issues Encountered When Enabling Database Management for Oracle Cloud Databases.

Edit Database Management for a Database

Perform the following steps to edit Database Management settings for a database using the Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the **Database information** tab, under the **Associated Services** section, check the status of **Database Management**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

4. If the Database Management is displayed as enabled, perform the following steps to edit Database Management.
5. Select **Edit**.
6. In the **Enable Diagnostics & Management** page, in the **Database information** section, enter the following details.
   - **Database type:** *Read-only*. Type of the database.
   - **Database system**: *Read-only*. Compartment in which the database is located.
   - **Database home**: *Read-only*. Database home of the database.
   - **Database name**: *Read-only*. Name of the database.
   - Select from one of the following options:
     - Enable Diagnostics & Management for CDB and all PDBs.
     - Enable Diagnostics & Management for CDB only.
   - **Service name**: The unique service name of the database. A default unique name is displayed, which can be changed if required.
   - **Protocol**: Select either the TCP or TCPS protocol to connect to the database. By default, the TCP protocol is selected.

     > **Note**
     >
     > If Oracle Data Guard is enabled after Database Management was enabled for a DB system using the TCPS protocol, then TCPS will have to be reconfigured. Enabling Oracle Data Guard is causing TCPS configuration to be overwritten, and it's recommended that TCPS is configured on a DB system after enabling Oracle Data Guard.

     > **Note**
     >
     > Database Management currently does not support Oracle Data Guard configuration, and Database Management features are not available for standby databases.

   - **Port**: Specify the port number. If TCP is selected in the Protocol field, then port number `1521` is displayed by default, and you can change it if required. You can select the port number from a range of 1 to 65535.
   - **Database wallet secret**: This field is only displayed if TCPS is selected in the Protocol field.
     Select the secret that contains the database wallet from the drop-down list. If an existing database wallet secret is not available, then select **Create new secret...** in the drop-down list. The **Create database wallet secret** panel is displayed, and you can create a new secret. For information on database wallets and creating a secret in the Vault service, see Oracle Cloud Database-related Prerequisite Tasks.

     If the Database Management (`dpd`) service policy that grants Database Management the permission to read the secret that contains the database wallet is not created, then the *'System policies are required..'* message is displayed. You can select **Add policy** to view and automatically create the service policy. For information on Vault service permissions required to use existing secrets or create new secrets, see Permissions Required to Enable Database Management for Oracle Cloud Databases.

Copyright © 2022, 2026, Oracle and/or its affiliates.

7. In the **Specify credentials for the connection** section, provide the following details.
   - **Database user name**: Enter the database user name.
   - **Database user password secret**:
     Select the secret that contains the database user password from the drop-down list. If the compartment in which the secret resides is different from the compartment displayed, then select **Change compartment** and select another compartment. If an existing secret with the database user password is not available, then select **Create new secret...** in the drop-down list. The **Create password secret** panel is displayed, and you can create a new secret. For information on database monitoring user credentials and saving the database user password as a secret in the Vault service, see Oracle Cloud Database-related Prerequisite Tasks.

     If the Database Management (`dpd`) service policy that grants Database Management the permission to read the secret that contains the database user password is not created, then the *'System policies are required..'* message is displayed. You can select **Add policy** to view and automatically create the service policy. For information on Vault service permissions required to use existing secrets or create new secrets, see Permissions Required to Enable Database Management for Oracle Cloud Databases.
8. In the **Select connection mode** section, select the private endpoint that will act as a representation of Database Management in the VCN in which the database can be accessed. You can choose the private endpoint from a different compartment as well. You must ensure that the appropriate Database Management private endpoint is available. Here are the two types of Database Management private endpoints:
   - Private endpoint for single-instance databases in the DB systems.
   - Private endpoint for RAC databases in the DB system.
   - If a Database Management private endpoint is not available, then you must create one. For information on how to create a private endpoint, see Create a Database Management Private Endpoint.
9. In the **Management options** section, choose between the following options.
   - **Full management**: This includes fleet management, advanced Performance Hub, and other SKU features along with basic management capabilities.
   - **Basic management**: This includes basic monitoring metrics and the ASH Analytics and SQL Monitoring features in Performance Hub for databases.

   For more information on Database Management options, see the *About Management Options* section in Enable Database Management for Oracle Cloud Databases.
10. Select **Save Changes**.
11. A confirmation message with a link to the **Work requests** section on the **Database information** page is displayed. Select the link to monitor the progress of the work request.
12. On the **Database information** tab, under the **Associated Services** section, verify if the status of **Database Management** is enabled.

If you encounter issues when enabling Database Management, to know about likely causes and solutions, see Issues Encountered When Enabling Database Management for Oracle Cloud Databases.

Disable Database Management for a Database

Perform the following steps to disable Database Management for a database using the Console:

Copyright © 2022, 2026, Oracle and/or its affiliates.

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the **Database information** tab, under the **Associated Services** section, check the status of **Database Management**.
4. If the Database Management is displayed as enabled, then select **Disable** to disable Database Management.
5. A confirmation message with a link to the **Work requests** section on the **Database information** page is displayed. Select the link to monitor the progress of the work request.
6. On the **Database information** tab, under the **Associated Services** section, verify if the status of **Database Management** is disabled.

Enable Database Management for a Pluggable Database

> **Note**
>
> You can also enable Database Management for a database from the Database Management Administration page. For more information, see Enable Database Management for Oracle Cloud Databases.

Prerequisites

- The Database Management must be enabled for the associated database with **Full Management** option. To enable Database Management for databases, see Enable Database Management for a Database.
  For more information on Database Management options, see the *About Management Options* section in Enable Database Management for Oracle Cloud Databases.

Procedure

Perform the following steps to enable Database Management for pluggable databases using the Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the **Database information** tab, under the **Associated Services** section, check the status of **Database Management**.
5. If the **Database Management** is displayed as **Not Enabled**, perform the following steps to enable Database Management.
6. Select **Enable**.
7. In the **Enable Diagnostics & Management** page, in the **Database information** section, enter the following details.
   - **Database type:** *Read-only*. Type of the database.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Database system**: *Read-only*. Compartment in which the database is located.
- **Database home**: *Read-only*. Database home of the database.
- **Database name**: *Read-only*. Name of the database.
- Select from one of the following options:
  - Enable Diagnostics & Management for CDB and all PDBs.
  - Enable Diagnostics & Management for CDB only.
- **Service name**: The unique service name of the database. A default unique name is displayed, which can be changed if required.
- **Protocol**: Select either the TCP or TCPS protocol to connect to the database. By default, the TCP protocol is selected.

> **Note**
>
> If Oracle Data Guard is enabled after Database Management was enabled for a DB system using the TCPS protocol, then TCPS will have to be reconfigured. Enabling Oracle Data Guard is causing TCPS configuration to be overwritten, and it's recommended that TCPS is configured on a DB system after enabling Oracle Data Guard.

> **Note**
>
> Database Management currently does not support Oracle Data Guard configuration, and Database Management features are not available for standby databases.

- **Port**: Specify the port number. If TCP is selected in the Protocol field, then port number `1521` is displayed by default, and you can change it if required. You can select the port number from a range of 1 to 65535.
- **Database wallet secret**: This field is only displayed if TCPS is selected in the Protocol field.
  Select the secret that contains the database wallet from the drop-down list. If an existing database wallet secret is not available, then select **Create new secret...** in the drop-down list. The **Create database wallet secret** panel is displayed, and you can create a new secret. For information on database wallets and creating a secret in the Vault service, see Oracle Cloud Database-related Prerequisite Tasks.

  If the Database Management (`dpd`) service policy that grants Database Management the permission to read the secret that contains the database wallet is not created, then the *'System policies are required..'* message is displayed. You can select **Add policy** to view and automatically create the service policy. For information on Vault service permissions required to use existing secrets or create new secrets, see Permissions Required to Enable Database Management for Oracle Cloud Databases.

8. In the **Specify credentials for the connection** section, provide the following details.
   - **Database user name**: Enter the database user name.
   - **Database user password secret**:
     Select the secret that contains the database user password from the drop-down list. If the compartment in which the secret resides is different from the compartment displayed, then select **Change compartment** and select another compartment. If an

Copyright © 2022, 2026, Oracle and/or its affiliates.

existing secret with the database user password is not available, then select **Create new secret...** in the drop-down list. The **Create password secret** panel is displayed, and you can create a new secret. For information on database monitoring user credentials and saving the database user password as a secret in the Vault service, see Oracle Cloud Database-related Prerequisite Tasks.

If the Database Management (`dpd`) service policy that grants Database Management the permission to read the secret that contains the database user password is not created, then the *'System policies are required..'* message is displayed. You can select **Add policy** to view and automatically create the service policy. For information on Vault service permissions required to use existing secrets or create new secrets, see Permissions Required to Enable Database Management for Oracle Cloud Databases.

9. In the **Select connection mode** section, select the private endpoint that will act as a representation of Database Management in the VCN in which the database can be accessed. You can choose the private endpoint from a different compartment as well. You must ensure that the appropriate Database Management private endpoint is available. Here are the two types of Database Management private endpoints:
   - Private endpoint for single-instance databases in the DB systems.
   - Private endpoint for RAC databases in the DB system.
   - If a Database Management private endpoint is not available, then you must create one. For information on how to create a private endpoint, see Create a Database Management Private Endpoint.
10. Select **Enable Diagnostics & Management**.
11. A confirmation message with a link to the **Work requests** section on the **Database information** page is displayed. Select the link to monitor the progress of the work request.
12. On the **Database information** tab, under the **Associated Services** section, verify if the status of **Database Management** is enabled.

If you encounter issues when enabling Database Management, to know about likely causes and solutions, see Issues Encountered When Enabling Database Management for Oracle Cloud Databases.

Edit Database Management for a Pluggable Database

Perform the following steps to edit Database Management for pluggable databases using the Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the **Database information** tab, under the **Associated Services** section, check the status of **Database Management**.
5. If the Database Management is displayed as **enabled**, perform the following steps to edit Database Management.
6. Select **Edit**.
7. In the **Enable Diagnostics & Management** page, in the **Database information** section, enter the following details.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Database type:** *Read-only*. Type of the database.
- **Database system**: *Read-only*. Compartment in which the database is located.
- **Database home**: *Read-only*. Database home of the database.
- **Database name**: *Read-only*. Name of the database.
- Select from one of the following options:
  - Enable Diagnostics & Management for CDB and all PDBs.
  - Enable Diagnostics & Management for CDB only.
- **Service name**: The unique service name of the database. A default unique name is displayed, which can be changed if required.
- **Protocol**: Select either the TCP or TCPS protocol to connect to the database. By default, the TCP protocol is selected.

> **Note**
>
> If Oracle Data Guard is enabled after Database Management was enabled for a DB system using the TCPS protocol, then TCPS will have to be reconfigured. Enabling Oracle Data Guard is causing TCPS configuration to be overwritten, and it's recommended that TCPS is configured on a DB system after enabling Oracle Data Guard.

> **Note**
>
> Database Management currently does not support Oracle Data Guard configuration, and Database Management features are not available for standby databases.

- **Port**: Specify the port number. If TCP is selected in the Protocol field, then port number 1521 is displayed by default, and you can change it if required. You can select the port number from a range of 1 to 65535.
- **Database wallet secret**: This field is only displayed if TCPS is selected in the Protocol field.
  Select the secret that contains the database wallet from the drop-down list. If an existing database wallet secret is not available, then select **Create new secret...** in the drop-down list. The **Create database wallet secret** panel is displayed, and you can create a new secret. For information on database wallets and creating a secret in the Vault service, see Oracle Cloud Database-related Prerequisite Tasks.

  If the Database Management (`dpd`) service policy that grants Database Management the permission to read the secret that contains the database wallet is not created, then the *'System policies are required..'* message is displayed. You can select **Add policy** to view and automatically create the service policy. For information on Vault service permissions required to use existing secrets or create new secrets, see Permissions Required to Enable Database Management for Oracle Cloud Databases.

8. In the **Specify credentials for the connection** section, provide the following details.
   - **Database user name**: Enter the database user name.
   - **Database user password secret**:
     Select the secret that contains the database user password from the drop-down list. If the compartment in which the secret resides is different from the compartment

Copyright © 2022, 2026, Oracle and/or its affiliates.

displayed, then select **Change compartment** and select another compartment. If an existing secret with the database user password is not available, then select **Create new secret...** in the drop-down list. The **Create password secret** panel is displayed, and you can create a new secret. For information on database monitoring user credentials and saving the database user password as a secret in the Vault service, see Oracle Cloud Database-related Prerequisite Tasks.

If the Database Management (`dpd`) service policy that grants Database Management the permission to read the secret that contains the database user password is not created, then the *'System policies are required..'* message is displayed. You can select **Add policy** to view and automatically create the service policy. For information on Vault service permissions required to use existing secrets or create new secrets, see Permissions Required to Enable Database Management for Oracle Cloud Databases.

9. In the **Select connection mode** section, select the private endpoint that will act as a representation of Database Management in the VCN in which the database can be accessed. You can choose the private endpoint from a different compartment as well. You must ensure that the appropriate Database Management private endpoint is available. Here are the two types of Database Management private endpoints:
   - Private endpoint for single-instance databases in the DB systems.
   - Private endpoint for RAC databases in the DB system.
   - If a Database Management private endpoint is not available, then you must create one. For information on how to create a private endpoint, see Create a Database Management Private Endpoint.
10. Select **Save Changes**.
11. A confirmation message with a link to the **Work requests** section on the **Database information** page is displayed. Select the link to monitor the progress of the work request.
12. On the **Database information** tab, under the **Associated Services** section, verify if the status of **Database Management** is enabled.

If you encounter issues when enabling Database Management, to know about likely causes and solutions, see Issues Encountered When Enabling Database Management for Oracle Cloud Databases.

Disable Database Management for a Pluggable Database

Perform the following steps to disable Database Management for pluggable databases using the Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the **Database information** tab, under the **Associated Services** section, check the status of **Database Management**.
5. If the Database Management is displayed as **Enabled**, perform the following steps to disable Database Management.
6. Select **Disable**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

7. A confirmation message with a link to the **Work requests** section on the **Database information** page is displayed. Select the link to monitor the progress of the work request.
8. On the **Database information** tab, under the **Associated Services** section, verify if the status of **Database Management** is **Disabled**.

View Performance Hub Metrics for Base Database Service Resources

You can use the Performance Hub to monitor database activity, diagnose issues, and tune queries to improve the performance of Oracle Databases. This article describes the procedure to view Performance Hub in Base Database Service.

Performance Hub displays information about the performance of your database for the time period you specify. With this tool, you can view real-time and historical performance data. For more information about Performance Hub, see About Performance Hub.

The following topics are covered in this article:

- View Performance Hub Metrics for a Database
- View Performance Hub Metrics for a Pluggable Database

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be granted security access in a **policy** by an administrator. This access is required whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you get a message that you don’t have permission or are unauthorized, verify with your administrator what type of access you have and which **compartment** to work in.

For administrators: The policy in Let database admins manage Oracle Cloud database systems lets the specified group do everything with databases and related Database resources.

If you're new to policies, see Getting Started with Policies and Common Policies. If you want to dig deeper into writing policies for databases, see Details for the Database Service.

View Performance Hub Metrics for a Database

Prerequisites

- The Database Management for databases must be enabled. To enable Database Management for databases, see Enable Database Management for a Database.

Procedure

Perform the following steps to view the Performance Hub for databases using the Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, from the **Actions** menu, select **Performance Hub**. The Performance Hub for the database is displayed.

Copyright © 2022, 2026, Oracle and/or its affiliates.

View Performance Hub Metrics for a Pluggable Database

Prerequisites

- The Database Management for databases with **Full Management** option must be enabled. To enable Database Management for databases, see Enable Database Management for a Database.
- The Database Management for pluggable databases must be enabled. To enable Database Management for pluggable databases, see Enable Database Management for a Pluggable Database.

Procedure

Perform the following steps to view the performance hub for pluggable databases using the Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the pluggable database details page, from the **Actions** menu, select **Performance Hub**. The Performance Hub for the pluggable database is displayed.

Monitor Using Oracle Enterprise Manager

You can use Oracle Enterprise Manager to manage and monitor the Base Database Service.

Enterprise Manager is Oracle's management platform, providing a single pane of glass for managing all your Oracle deployments, whether in your data centers or the Oracle Cloud. Through deep integration with Oracle's product stack, Enterprise Manager provides market-leading management and automation support for Oracle applications, databases, middleware, hardware, and engineered systems.

For more information about the Enterprise Manager, see About Enterprise Manager Cloud Control 13c.

> **Note**
>
> Enterprise Manager versions 13.3 and above support monitoring and management of the Base Database Service.

To use Enterprise Manager to manage and monitor your database, you perform these high-level tasks:

- Configure connectivity between your Enterprise Manager deployment (whether on-premises or on the Oracle Cloud Infrastructure Marketplace) and your database.
- Use the Enterprise Manager console, CLI, or REST API to discover the database and add it as a target.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For detailed steps to perform these tasks, see Discovering an Oracle Public Cloud Machine.

After performing the discovery tasks, you can use Enterprise Manager to:

- Monitor the health and performance of your database and perform deep diagnostics on the Performance Hub.
- Perform database administration tasks such as storage management and schema management tasks such as creating database objects.

For details about how to use these features, see Administering and Monitoring a PaaS Cloud.

Monitor a Database with Enterprise Manager Express

This article explains how to set up an Enterprise Manager Express console to monitor the database.

On 1- and 2-node RAC DB Systems, by default, the EM Express console is not enabled on version 18.1.0.0, 12.2.0.1, and 12.1.0.2 databases. You can enable it for an existing database as described below, or you can enable it when you create a database by using the Database Commands with the `-co` parameter.

You must also update the security list and iptables for the DB system as described later in this topic.

When you enable the console, you'll set the port for the console. The procedure below uses port 5500, but each additional console enabled on the same DB system will have a different port.

Required IAM Policy

Some of the procedures below require permission to create or update security lists. For more information about security list policies, see Security Lists.

Enable the EM Express Console and Determine its Port Number

1. SSH to the DB system, log in as `opc`, sudo to the oracle user, and log in to the database as `SYS`.

   ```
   sudo su - oracle
   . oraenv
   <provide the database SID at the prompt>
   sqlplus / as sysdba
   ```

2. Do one of the following:
   - To enable the console and set its port, use the following command.

     ```
     exec DBMS_XDB_CONFIG.SETHTTPSPORT(<port>);
     ```

     For example:

     ```
     exec DBMS_XDB_CONFIG.SETHTTPSPORT(5500);
     ```

   - To determine the port for a previously enabled console, use the following command.

     ```
     select dbms_xdb_config.getHttpsPort() from dual;
     ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

For example:

```
select dbms_xdb_config.getHttpsPort() from dual;
```

Output:

```
DBMS_XDB_CONFIG.GETHTTPSPORT()
------------------------------
                          5500
```

3. Return to the operating system by typing `exit` and then confirm that the listener is listening on the port:

   ```
   lsnrctl status | grep HTTP
   ```

   Output:

   ```
   (DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)(HOST=xxx.xx.xxxxxx.xxx)(PORT=5500))
   (Security=(my_wallet_directory=/u01/app/oracle/admin/prod/xdb_wallet))
   (Presentation=HTTP)(Session=RAW))
   ```

4. If you are using a 2-node RAC DB system, see Set the Required Permissions On a 2-node RAC DB System.
5. Open the console's port as described in *Open Ports on the DB System*.
6. Update the security list for the console's port as described in *Update the Security List for the DB System*.

Set the Required Permissions On a 2-node RAC DB System

If you're using a 2-node RAC DB system, you'll need to add read permissions for the `asmadmin` group on the wallet directory on *both* nodes in the system.

1. SSH to one of the nodes in the DB system, log in as `opc`, and `sudo` to the grid user.

   ```
   sudo su - grid
   . oraenv
   ORACLE_SID = [+ASM1] ?
   The Oracle base has been set to /u01/app/grid
   ```

2. Get the location of the wallet directory by executing the following command.

   ```
   lsnrctl status | grep xdb_wallet
   ```

   Output:

   ```
   (DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)
   (HOST=dbsysHost1.sub04061528182.dbsysapril6.oraclevcn.com)(PORT=5500))
   (Security=(my_wallet_directory=/u01/app/oracle/admin/dbsys12_phx3wm/
   xdb_wallet))(Presentation=HTTP)(Session=RAW))
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

3. Return to the opc user, switch to the oracle user, and change to the wallet directory.

```
sudo su - oracle
cd /u01/app/oracle/admin/dbsys12_phx3wm/xdb_wallet
```

4. List the directory contents and note the permissions.

```
ls -ltr
```

Output:

```
total 8
-rw------- 1 oracle asmadmin 3881 Apr  6 16:32 ewallet.p12
-rw------- 1 oracle asmadmin 3926 Apr  6 16:32 cwallet.sso
```

5. Change the permissions:

```
chmod 640 /u01/app/oracle/admin/dbsys12_phx3wm/xdb_wallet/*
```

6. Verify that read permissions were added.

```
ls -ltr
```

Output:

```
total 8
-rw-r----- 1 oracle asmadmin 3881 Apr  6 16:32 ewallet.p12
-rw-r----- 1 oracle asmadmin 3926 Apr  6 16:32 cwallet.sso
```

7. Repeat the steps above on the other node in the cluster.

Connect to the EM Express Console

After you've enabled the Console and opened its port in the security list and iptables, you can connect as follows:

1. From a web browser, connect to the Console using the following URL format:

```
https://<ip_address>:<port>/em
```

For example, `https://129.145.0.164:5500/em`

Use the DB system's private or public IP address depending on your network configuration.

Use the private IP address to connect to the system from your on-premises network, or from within the Virtual Cloud Network (VCN). This includes connecting from a host located on-premises connecting through a VPN or FastConnect to your VCN, or from another host in the same VCN. Use the public IP address to connect to the system from outside the cloud (with no VPN). You can find the IP addresses in the Console as follows:

- On the **DB System Details** page, under **Resources**, click **Nodes**.
- The **Public IP address** and **Private IP address & DNS name** are displayed in the table columns.

Copyright © 2022, 2026, Oracle and/or its affiliates.

2. A login page is displayed and you can log in with any valid database credentials.

3. The Database Home page is displayed.

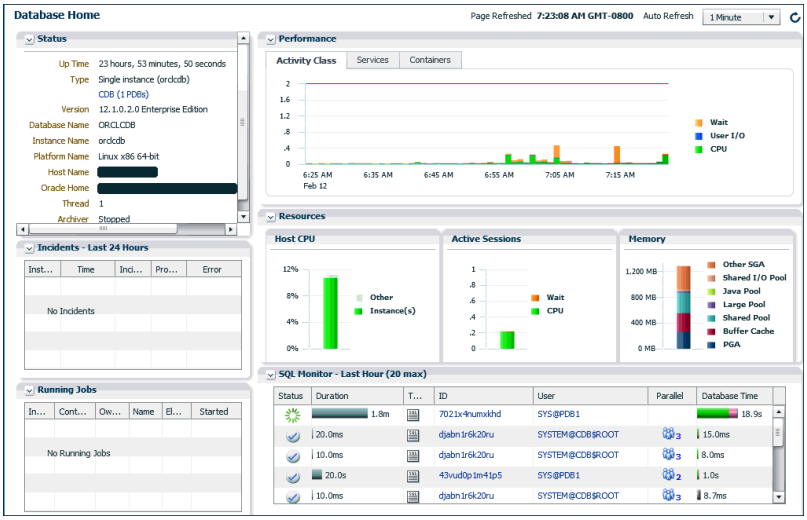

To learn more about EM Express, see Introduction to Oracle Enterprise Manager Database Express.

> **Note**
>
> If you're using a 1-node DB system, and you are unable to connect to the EM Express console, see Database Known Issues.

Monitor a Database with Enterprise Manager Database Control

This article explains how to set up an Enterprise Manager Database Control console to monitor the database.

By default, the Enterprise Manager Database Control console is not enabled on version 11.2.0.4 databases. You can enable the console:

- when you create a container database by using the Database Commands with the `-co` parameter.
- for an existing container database as described in Configuring Database Control Using EMCA.

Port 1158 is the default port used for the first console enabled on the DB system, but each additional console enabled on the DB system will have a different port.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> For a version 11.2.0.4 database on a 2-node RAC DB system, see Enable the Console For a Version 11.2.0.4 Database On a Multi-node DB System.

Required IAM Policy

Some of the procedures below require permission to create or update security lists. For more information about security list policies, see Security Lists.

Determine the Port For the Enterprise Manager Database Control Console

1. SSH to the DB system, log in as `opc`, and `sudo` to the oracle user.

   ```
   sudo su - oracle
   . oraenv
    <provide the database SID at the prompt>
   ```

2. Use the following command to get the port number.

   ```
   emctl status dbconsole
   ```

   The port is in the URL, as shown in the following output:

   ```
   Oracle Enterprise Manager 11g Database Control Release 11.2.0.4.0
   Copyright (c) 1996, 2013 Oracle Corporation.  All rights reserved.
   https://dbprod:1158/em/console/aboutApplication
   Oracle Enterprise Manager 11g is running.
   ------------------------------------------------------------------
   Logs are generated in directory /u01/app/oracle/product/11.2.0.4/dbhome_2/
   dbprod_db11/sysman/log
   ```

3. Open the console's port as described in *Open Ports on the DB System*.
4. Update the security list for the console's port as described in *Update the Security List for the DB System*.

Connect to the Enterprise Manager Database Control Console

After you've enabled the console and opened its port in the security list and iptables, you can connect as follows:

1. From a web browser, connect to the console using the following URL format:

   `https://<ip_address>:<port>/em`

   For example, `https://129.145.0.164:1158/em`

   Use the DB system's private or public IP address depending on your network configuration.

   Use the private IP address to connect to the system from your on-premises network, or from within the Virtual Cloud Network (VCN). This includes connecting from a host located on-premises connecting through a VPN or FastConnect to your VCN, or from another host

Copyright © 2022, 2026, Oracle and/or its affiliates.

in the same VCN. Use the public IP address to connect to the system from outside the cloud (with no VPN). You can find the IP addresses in the Console as follows:

- On the **DB System Details** page, under **Resources**, click **Nodes**.
- The **Public IP address** and **Private IP address & DNS name** are displayed in the table columns.

2. A login page will be displayed and you can log in with any valid database credentials.

To learn more about Enterprise Manager Database Control, see Introduction to Oracle Enterprise Manager Database Control.

Enable the Console For a Version 11.2.0.4 Database On a Multi-node DB System

A few extra steps are required to enable the console for a version 11.2.0.4 database on a multi-node DB system.

Configure SSH Equivalency Between the Two Nodes

You'll create SSH keys on each node and copy the key to the other node, so that each node has the keys for both nodes. The following procedure uses the sample names node1 and node2.

1. SSH to node1, log in as `opc`, and `sudo` to the oracle user.

   ```
   sudo su - oracle
   ```

2. Create a directory called `.ssh`, set its permissions, create an RSA key, and add the public key to the `authorized_keys` file.

   ```
   mkdir .ssh
   chmod 755 .ssh
   ssh-keygen -t rsa
   cat id_rsa.pub > authorized_keys
   ```

3. Repeat the previous steps on the other node in the cluster.
4. On each node, add the `id_rsa.pub` key for the *other* node to the `authorized_keys` file. When you're done, you should see both keys in `authorized_keys` on each node.
5. On node1, create the `known_hosts` file by doing the following:
   - SSH to node1 and reply yes to the authentication prompt.
   - SSH to node2 and reply yes to the authentication prompt.
6. On node2, create the `known_hosts` file by doing the following:
   - SSH to node2 and reply yes to the authentication prompt.
   - SSH to node1 and reply yes to the authentication prompt.
7. On node1, verify that SSH equivalency is now configured by using the following Cluster Verification Utility (CVU) command.

   ```
   cluvfy stage -pre crsinst -n all -verbose
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Configure the Console

1. On node1, create a file called `emca.rsp` with the following entries.

```
DB_UNIQUE_NAME=<pdb_unique_name>
SERVICE_NAME=<db_unique_name>.<db_domain>
PORT=<scan listener port>
LISTENER_OH=$GI_HOME
SYS_PWD=<admin password>
DBSNMP_PWD=<admin password>
SYSMAN_PWD=<admin password>
CLUSTER_NAME=<cluster name>     <=== to get the cluster name,
run: $GI_HOME/bin/cemutlo -n
ASM_OH=$GI_HOME
ASM_SID=+ASM1
ASM_PORT=<asm listener port>
ASM_USER_NAME=ASMSNMP
ASM_USER_PWD=<admin password>
```

2. On node1, run Enterprise Manager Configuration Assistant (EMCA) using the `emca.rsp` file as input.

```
$ORACLE_HOME/bin/emca
    -config dbcontrol db
    -repos create
    -cluster
    -silent
    -respFile <location of response file above>
```

3. On node2, configure the console so the agent in node1 reports to the console in node1, and the agent in node2 reports to the console in node2.

```
$ORACLE_HOME/bin/emca
    -reconfig dbcontrol
    -silent
    -cluster
    -EM_NODE <node2 host>
    -EM_NODE_LIST <node2 host>
    -DB_UNIQUE_NAME <db_unique_name>
    -SERVICE_NAME <db_unique_name>.<db_domain>
```

4. On each node, verify that console is working properly.

```
export ORACLE_UNQNAME=<db_unique_name>

emctl status agent
```

Output:

```
Oracle Enterprise Manager 11g Database Control Release 11.2.0.4.0
Copyright (c) 1996, 2013 Oracle Corporation.  All rights reserved.
------------------------------------------------------------------
Agent Version     : 10.2.0.4.5
OMS Version       : 10.2.0.4.5
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
Protocol Version  : 10.2.0.4.5
Agent Home        : /u01/app/oracle/product/11.2.0.4/dbhome_x/
<host>_<db_unique_name>
Agent binaries    : /u01/app/oracle/product/11.2.0.4/dbhome_x
Agent Process ID  : 26194
Parent Process ID : 25835
Agent URL         : https://<node host>:1831/emd/main
Repository URL    : https://<node host>:5501/em/upload/
Started at        : 2017-03-15 20:20:34
Started by user   : oracle
Last Reload       : 2017-03-15 20:27:00
Last successful upload                       : 2017-03-15 21:06:36
Total Megabytes of XML files uploaded so far :    22.25
Number of XML files pending upload           :
0                      <=== should be zero
Size of XML files pending upload(MB)         :     0.00
Available disk space on upload filesystem    :    42.75%
Data channel upload directory                : /u01/app/oracle/product/
11.2.0.4/dbhome_x/<host>_<db_unique_name>/sysman/recv
Last successful heartbeat to OMS             : 2017-03-15 21:08:45
---------------------------------------------------------------
```

Update iptables and Security List

1. On each node, edit iptables to open the console's port as described in *Open Ports on the DB System*.
2. Update the security list for the console's port as described in *Update the Security List for the DB System*.

Events

Manage Diagnostics Collection for the DB System

This article provides the details and procedure to edit diagnostic collection for the DB system.

The diagnostics collection and notifications feature enables Oracle Cloud Operations and you to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to events to get notified about resource state changes. You can enable or disable this feature at any time.

Diagnostic events: Allow Oracle to collect and publish critical, warning, error, and information events for you. For more information, see *Database Service Events*.

Incident logs and trace collection: Allow Oracle to collect incident logs and traces to enable fault diagnosis and issue resolution. For more information, see *Incident Logs and Trace Files*.

General Information

- You are opting in with the understanding that the list of events and log files can change in the future. You can opt out of this feature at any time.
- Disabling diagnostic events and health monitoring will only prevent the collection and notification of data/events from the time you opt out. However, historical data will not be purged from Oracle Cloud Operations data repositories.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- If you had previously opted in for incident log and trace file collection and then decide to opt out when Oracle Cloud operations run a log collection job, the job will run its course and will not be canceled. However, subsequent log collections won't happen till you opt in again to the incident logs and trace file collection option.
- The **Enable health monitoring** diagnostics collection for Oracle Cloud operations viewing is not available for the Oracle Base Database Service.

Procedure

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **DB system information** tab, under the **General information** section, the **Diagnostics collection** status is displayed. The status would be any of the following:
   - **Enabled**: Both the diagnostic events and incident logs and trace file collection are opted in.
   - **Disabled**: Both the diagnostic events and incident logs and trace file collection are opted out.
   - **Partially enabled**: Either the diagnostic events or incident logs and trace file collection are opted in.
3. Select the **Edit** button beside the **Diagnostic collection** status.
4. The **Edit diagnostics collection settings** panel is displayed.
5. Select the **Diagnostics collection** as per your requirements from the following options. Unchecking all the available options will disable diagnostics collection and notification.
   - **Enable diagnostic events**: Enables and allows Oracle to collect and send fault notifications about critical, warning, and information events for you.
   - **Enable incident logs and trace collection**: Enables and allows Oracle to receive event notifications and collect incident logs and traces for fault diagnosis and issue resolution.
6. Select **Save**.

Database Service Events

Database Service Events feature implementation enables you to get notified about health issues with your Oracle Databases or other components on the DB system.

It is possible that Oracle Database or Clusterware may not be healthy or various system components may be running out of space on the DB system. Customers are not notified of this situation. Database Service Events feature implementation generates events for Data Plane operations and conditions, as well as Notifications for customers by leveraging the existing OCI Events service and Notification mechanisms in their tenancy. Customers can then create topics and subscribe to these topics through email, functions, or streams.

> **Note**
>
> Events flow on the DB system depends on Oracle Trace File Analyzer (TFA) and Oracle Database Cloud Service (DBCS) agent. Ensure that these components are up and running.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Receive Notifications about Database Service Events

Subscribe to the Database Service Events and get notified. To receive notifications, subscribe to Database Service Events and get notified using the Oracle Notification service, see Notifications Overview. For more information about Oracle Cloud Infrastructure Events, see Overview of Events.

Events Service - Event Types

- Database - Critical
- DB Node - Critical
- DB Node - Error
- DB Node - Warning
- DB Node - Information
- DB System - Critical

Database Service Event Types

The following table lists the event types that the Database Service emits.

> **Note**
>
> - **Critical** events are triggered by various critical conditions and errors that disrupt the database or other essential components. For example, database unresponsive errors and availability errors for databases, database nodes, and DB systems indicate when a resource becomes unavailable.
> - **Information** events are triggered when the database and other critical components work as expected. For example, a clean shutdown of CRS, CDB, client, or scan listener, or a startup of these components will create an event with the severity of `INFO`.
> - Threshold limits reduce the number of notifications customers will receive for similar incident events whilst at the same time ensuring they receive the incident events and are reminded in a timely fashion.

Database Service Events

> **Note**
>
> In addition to the events listed below, Oracle analyzes additional events to provide the highest level of service operations and support for ensuring high availability of services.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-3 Database Service Events

Friendly Name	Event Name	Description	Remediation	Event Type	Threshold
Resource Utilization - Disk Usage	HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE	This event is reported when VM guest file system free space falls below 10% free, as determined by the operating system df(1) command, for the following file systems: • / • /u01 • /u02 • /var (X8M and later only) • /tmp (X8M and later only)	HEALTH-DB_GUEST-FILESYSTEM-FREE_SPACE	com.oraclecloud.databaseservice.dbnode.critical	Critical threshold : 90%
CRS status Up/Down	AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN.	An event of type CRITICAL is created when the Cluster Ready Service (CRS) is detected to be down.	AVAILABILITY-DB_GUEST-CRS_INSTANCE.DOWN	com.oraclecloud.databaseservice.dbnode.critical (if .DOWN and NOT "user_action")	NA
	AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN_CLEARED	An event of type INFORMATION is created once it is determined that the event for CRS down has cleared.	NA	com.oraclecloud.databaseservice.dbnode.information (if .DOWN_CLEARED)	NA
	AVAILABILITY.DB_GUEST.CRS_INSTANCE.EVICTION	An event of type CRITICAL is created.	AVAILABILITY-DB_GUEST-CRS_INSTANCE-EVICTION	com.oraclecloud.databaseservice.dbnode.critical	NA

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-3 (Cont.) Database Service Events

Friendly Name	Event Name	Description	Remediation	Event Type	Threshold
SCAN Listener Up/Down	AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN	A DOWN event is created when a SCAN listener goes down. The event is of type INFORMATION when a SCAN listener is shutdown due to user action, such as with the Server Control Utility (srvctl) or Listener Control (lsnrctl) commands, or any Oracle Cloud maintenance action that uses those commands, such as performing a grid infrastructure software update. The event is of type CRITICAL when a SCAN listener goes down unexpectedly. A corresponding DOWN_CLEARED event is created when a SCAN listener is started. There are three SCAN listeners per cluster called LISTENER_SCAN[1,2,3].	AVAILABILITY-DB_CLUSTER-SCAN_LISTENER-DOWN	com.oraclecloud.databaseservice.dbnode.critical (if .DOWN and NOT "user_action")	NA
	AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN_CLEARED	An event of type INFORMATION is created once it is determined that the event for SCAN Listener down has cleared.	NA	com.oraclecloud.databaseservice.dbnode.information (if .DOWN_CLEARED)	NA

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-3 (Cont.) Database Service Events

Friendly Name	Event Name	Description	Remediation	Event Type	Threshold
Net Listener Up/Down	AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN	A DOWN event is created when a client listener goes down. The event is of type INFORMATION when a client listener is shutdown due to user action, such as with the Server Control Utility (srvctl) or Listener Control (lsnrctl) commands, or any Oracle Cloud maintenance action that uses those commands, such as performing a grid infrastructure software update. The event is of type CRITICAL when a client listener goes down unexpectedly. A corresponding DOWN_CLEARED event is created when a client listener is started. There is one client listener per node, each called LISTENER.	AVAILABILITY-DB_GUEST-CLIENT_LISTENER.DOWN	com.oraclecloud.databaseservice.database.critical (if .DOWN and NOT "user_action")	NA
	AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN_CLEARED	An event of type INFORMATION is created once it is determined that the event for Client Listener down has cleared.	NA	com.oraclecloud.databaseservice.database.information (if .DOWN_CLEARED)	NA

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-3 (Cont.) Database Service Events

Friendly Name	Event Name	Description	Remediation	Event Type	Threshold
CDB Up/Down	AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN	A DOWN event is created when a database instance goes down. The event is of type INFORMATION when a database instance is shutdown due to user action, such as with the SQL*Plus (sqlplus) or Server Control Utility (srvctl) commands, or any Oracle Cloud maintenance action that uses those commands, such as performing a database home software update. The event is of type CRITICAL when a database instance goes down unexpectedly. A corresponding DOWN_CLEARED event is created when a database instance is started.	AVAILABILITY-DB_GUEST-CDB_INSTANCE-DOWN	com.oraclecloud.databaseservice.database.critical (if .DOWN and NOT "user_action")	NA
	AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN_CLEARED	An event of type INFORMATION is created once it is determined that the event for the CDB down has cleared.	NA	com.oraclecloud.databaseservice.database.information (if .DOWN_CLEARED)	NA
Critical DB Errors	HEALTH.DB_CLUSTER.CDB.CORRUPTION	Database corruption has been detected on your primary or standby database. The database alert.log is parsed for any specific errors that are indicative of physical block corruptions, logical block corruptions, or logical block corruptions caused by lost writes.	HEALTH-DB_CLUSTER-CDB-CORRUPTION	com.oraclecloud.databaseservice.database.critical	NA

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-3 (Cont.) Database Service Events

Friendly Name	Event Name	Description	Remediation	Event Type	Threshold Id
Other DB Errors	HEALTH.DB_CLUSTER.CDB.ARCHIVER_HANG	An event of type CRITICAL is created if a CDB is either unable to archive the active online redo log or unable to archive the active online redo log fast enough to the log archive destinations.	HEALTH-DB_CLUSTER-CDB-ARCHIVER_HANG	com.oraclecloud.databaseservice.database.critical	NA
	HEALTH.DB_CLUSTER.CDB.DATABASE_HANG	An event of type CRITICAL is created when a process or session becomes unresponsive in the CDB.	HEALTH-DB_CLUSTER-CDB-DATABASE_HANG	com.oraclecloud.databaseservice.database.critical	NA
Backup Failures	HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE	An event of type CRITICAL is created if there is a CDB backup with a FAILED status reported in the v$rman_status view.	HEALTH-DB_CLUSTER-CDB-BACKUP_FAILURE	com.oraclecloud.databaseservice.database.critical	NA
	HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE_CLEARED	An event of type INFORMATION is created.	NA	com.oraclecloud.databaseservice.database.information	NA
Disk Group Usage	HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE	An event of type CRITICAL is created when an ASM disk group reaches space usage of 90% or higher. An event of type INFORMATION is created when the ASM disk group space usage drops below 90%.	HEALTH-DB_CLUSTER-DISK_GROUP-FREE_SPACE	com.oraclecloud.databaseservice.dbsystem.critical com.oraclecloud.databaseservice.dbsystem.information (if < 90%)	Notifications are sent when the usage hits 70%, 80%, 90%, and 100% with a corresponding severity of 4, 3, 2, and 1.

Temporarily Restrict Automatic Diagnostic Collections for Specific Events

Use the `tfactl blackout` command to temporarily suppress automatic diagnostic collections.

Copyright © 2022, 2026, Oracle and/or its affiliates.

If you set blackout for a target, then Oracle Trace File Analyzer stops automatic diagnostic collections if it finds events in the alert logs for that target while scanning. By default, blackout will be in effect for 24 hours.

You can also restrict automatic diagnostic collection at a granular level, for example, only for `ORA-00600` or even only `ORA-00600` with specific arguments.

Syntax

```
tfactl blackout add|remove|print
    -targettype host|crs|asm|asmdg|database|dbbackup|db_dataguard|
        db_tablespace|pdb_tablespace|pdb|listener|service|os
    -target all|name
    [-container name]
    [-pdb pdb_name]
    -event all|"event_str1,event_str2"|availability
    [-timeout nm|nh|nd|none]
    [-c|-local|-nodes "node1,node2"]
    [-reason "reason for blackout"]
    [-docollection]
```

Parameters

Table 5-4 Parameters

Parameter	Description
`add|remove|print|`	Adds, removes, or prints blackout conditions.
`-targettype type`	Limits blackout only to the specified target type.
Target type:`host|crs|asm|asmdg|database|dbbackup|db_dataguard|db_tablespace|pdb_tablespace|pdb|listener|service|os`	`host`: The whole node is under blackout. If there is host blackout, then every blackout element that's shown true in the Telemetry JSON will have the reason for the blackout. `crs`: Blackout the availability of the Oracle Clusterware resource or events in the Oracle Clusterware logs. `asm`: Blackout the availability of Oracle Automatic Storage Management (Oracle ASM) on this machine or events in the Oracle ASM alert logs. `asmdg`: Blackout an Oracle ASM disk group. `database`: Blackout the availability of an Oracle Database, Oracle Database backup, tablespace, and so on, or events in the Oracle Database alert logs. `dbbackup`: Blackout Oracle Database backup events (such as CDB or archive backups). `db_dataguard`: Blackout Oracle Data Guard events. `db_tablespace`: Blackout Oracle Database tablespace events (container database). `pdb_tablespace`: Blackout Oracle pluggable database tablespace events (pluggable database). `pdb`: Blackout Oracle pluggable database events. `listener`: Blackout the availability of a listener. `service`: Blackout the availability of a service. `os`: Blackout one or more operating system records.
`-target all|name`	Specify the target for blackout. You can specify a comma-delimited list of targets. By default, the target is set to `all`.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-4 (Cont.) Parameters

Parameter	Description
`-container name`	Specify the database container name (`db_unique_name`) where the blackout will take effect (for PDB, DB_TABLESPACE, and PDB_TABLESPACE).
`-pdb pdb_name`	Specify the PDB where the blackout will take effect (for PDB_TABLESPACE only).
`-events all|"str1,str2"`	Limits blackout only to the availability events, or event strings, which should not trigger auto collections, or be marked as blacked out in telemetry JSON. `all`: Blackout everything for the target specified. string: Blackout for incidents where any part of the line contains the strings specified. Specify a comma-delimited list of strings.
`-timeout nh|nd|none`	Specify the duration for blackout in number of hours or days before timing out. By default, the timeout is set to 24 hours (24h).
`-c|-local`	Specify if blackout should be set to `cluster-wide` or `local`. By default, blackout is set to `local`.
`-reason comment`	Specify a descriptive reason for the blackout.
`-docollection`	Use this option to do an automatic diagnostic collection even if a blackout is set for this target.

Examples

The following are the examples to use tfactl blackout command.

To blackout event: `ORA-00600` on targettype: `database`, target: `mydb`

```
tfactl blackout add -targettype database -target mydb -event "ORA-00600"
```

To blackout event: `ORA-04031` on targettype: `database`, target: `all`

```
tfactl blackout add -targettype database -target all -event "ORA-04031" -
timeout 1h
```

To blackout db backup events on targettype: `dbbackup`, target: mydb

```
tfactl blackout add -targettype dbbackup -target mydb
```

To blackout db dataguard events on targettype: `db_dataguard`, target: `mydb`

```
tfactl blackout add -targettype db_dataguard -target mydb -timeout 30m
```

To blackout db tablespace events on targettype: `db_tablespace`, target: `system`, container: `mydb`

```
tfactl blackout add -targettype db_tablespace -target system -container mydb -
timeout 30m
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

To blackout `ALL` events on targettype: `host`, target: `all`

```
tfactl blackout add -targettype host -event all -target all -timeout 1h
    -reason "Disabling all events during patching"
```

To print blackout details:

```
tfactl blackout print
```

```
.--------------------------------------------------------------------------
------------------------------------------------------------------------------
-------------------.
|

myhostname
                 |
+--------------+--------------------+-----------
+-----------------------------+-----------------------------+--------
+--------------+-------------------------------------+
| Target Type   | Target              | Events    | Start
Time                    | End Time                      | Status | Do
Collection | Reason                              |
+--------------+--------------------+-----------
+-----------------------------+-----------------------------+--------
+--------------+-------------------------------------+
| HOST          | ALL                 | ALL       | Thu Mar 24 16:48:39 UTC
2022 | Thu Mar 24 17:48:39 UTC 2022 | ACTIVE | false         | Disabling all
events during patching |
| DATABASE      | MYDB                | ORA-00600 | Thu Mar 24 16:39:03 UTC
2022 | Fri Mar 25 16:39:03 UTC 2022 | ACTIVE | false         |
NA                                  |
| DATABASE      | ALL                 | ORA-04031 | Thu Mar 24 16:39:54 UTC
2022 | Thu Mar 24 17:39:54 UTC 2022 | ACTIVE | false         |
NA                                  |
| DB_DATAGUARD  | MYDB                | ALL       | Thu Mar 24 16:41:38 UTC
2022 | Thu Mar 24 17:11:38 UTC 2022 | ACTIVE | false         |
NA                                  |
| DBBACKUP      | MYDB                | ALL       | Thu Mar 24 16:40:47 UTC
2022 | Fri Mar 25 16:40:47 UTC 2022 | ACTIVE | false         |
NA                                  |
| DB_TABLESPACE | SYSTEM_CDBNAME_MYDB | ALL       | Thu Mar 24 16:45:56 UTC
2022 | Thu Mar 24 17:15:56 UTC 2022 | ACTIVE | false         |
NA                                  |
'--------------+--------------------+-----------
+-----------------------------+-----------------------------+--------
+--------------+-------------------------------------'
```

To remove blackout for event: `ORA-00600` on targettype: `database`, target: `mydb`

```
tfactl blackout remove -targettype database -event "ORA-00600" -target mydb
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

To remove blackout for db backup events on targettype: `dbbackup`, target: `mydb`

```
tfactl blackout remove -targettype dbbackup -target mydb
```

To remove blackout for db tablespace events on targettype: `db_tablespace`, target: `system`, container: `mydb`

```
tfactl blackout remove -targettype db_tablespace -target system -container
mydb
```

To remove blackout for host events on targettype: `all`, target: `all`

```
tfactl blackout remove -targettype host -event all -target all
```

Manage Oracle Trace File Analyzer

To check the run status of Oracle Trace File Analyzer, run the `tfactl status` command as `root` or a non-root user:

```
tfactl status
```

```
.-----------------------------------------------------------------------------
----------------.
| Host  | Status of TFA | PID    | Port | Version    | Build ID             |
Inventory Status |
+-------+---------------+--------+------+------------+----------------------
+------------------+
| node1 | RUNNING       | 41312  | 5000 | 22.1.0.0.0 | 22100020220310214615 |
COMPLETE         |
| node2 | RUNNING       | 272300 | 5000 | 22.1.0.0.0 | 22100020220310214615 |
COMPLETE         |
'-----------------------------------------------------------------------------
----------------'
```

To start the Oracle Trace File Analyzer daemon on the local node, run the `tfactl start` command as `root` user:

```
tfactl start
```

```
Starting TFA..
Waiting up to 100 seconds for TFA to be started..
. . . . .
Successfully started TFA Process..
. . . . .
TFA Started and listening for commands
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

To stop the Oracle Trace File Analyzer daemon on the local node, run the `tfactl stop` command as `root` user:

```
tfactl stop


Stopping TFA from the Command Line
Nothing to do !
Please wait while TFA stops
Please wait while TFA stops
TFA-00002 Oracle Trace File Analyzer (TFA) is not running
TFA Stopped Successfully
Successfully stopped TFA..
```

Manage Database Service Agent

View the `/opt/oracle/dcs/log/dcs-agent.log` file to identify issues with the agent.

To check the status of the Database Service Agent, run the `systemctl status` command:

```
systemctl status dbcsagent.service


dbcsagent.service
Loaded: loaded (/usr/lib/systemd/system/dbcsagent.service; enabled; vendor
preset: disabled)
Active: active (running) since Fri 2022-04-01 13:40:19 UTC; 6min ago
Process: 9603 ExecStopPost=/bin/bash -c kill `ps -fu opc |grep "java.*dbcs-
agent.*jar" |
    awk '{print $2}' ` (code=exited, status=0/SUCCESS)
Main PID: 10055 (sudo)
CGroup: /system.slice/dbcsagent.service
 10055 sudo -u opc /bin/bash -c umask 077; /bin/java -
Doracle.security.jps.config=/opt/oracle/...
```

To start the agent if it is not running, run the `systemctl start` command as the `root` user:

```
systemctl start dbcsagent.service
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Incident Logs and Trace Files

This article lists all the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.

> **Note**
>
> - Oracle will create a service request (SR) against the infrastructure Customer Support Identifier (CSI) when an issue is detected and needs customer interaction to resolve.
> - The customer's Oracle Cloud Infrastructure tenancy admin email will be used as the CSI contact to create SR and attach logs to it. Ensure tenancy admin email is added as a CSI contact in My Oracle Support (MOS).

Oracle Trace File Analyze (TFA) Component Driven Logs Collections

The directories are generally assigned to a component and that component can then be used to guide TFA to the files it needs to collect, for example, requesting the Cluster Ready Services (CRS) component would tell TFA to look at directories mapped to the CRS component and find files that match the required collection time frame.

> **Note**
>
> - If have previously opted-in for incident log and trace file collection and decide to opt-out when Oracle Cloud operations run a log collection job, then the job will run its course and will not cancel. Future log collections won't happen until you opt-in again to the incident logs and trace file collection option.
> - TFA is shipped with scripts that run when a particular component is requested, for example, for CRS component, `crscollect.pl` will run a number of `crsctl` commands and gather the input. By default, TFA does not redact collected logs.

Table 5-5 Oracle Trace File Analyze (TFA) Component Driven Logs Collections

Component	Script	Files/Directories
OS: Operating system logs	`oscollect.pl`	• `/var/log/messages` • OSWatcher archive

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-5 (Cont.) Oracle Trace File Analyze (TFA) Component Driven Logs Collections

Component	Script	Files/Directories
CRS: Grid Infrastructure and cluster logs	`crscollect.pl`	• `/etc/oracle` • `GIHOME/crf/db/HOSTNAME1` • `GIHOME/crs/log` • `GIHOME/css/log` • `GIHOME/cv/log` • `GIHOME/evm/admin/log` • `GIHOME/evm/admin/logger` • `GIHOME/evm/log` • `GIHOME/log/-/client` • `GIHOME/log/HOSTNAME1` • `GIHOME/log/HOSTNAME1/admin` • `GIHOME/log/HOSTNAME1/client` • `GIHOME/log/HOSTNAME1/crflogd` • `GIHOME/log/HOSTNAME1/crfmond` • `GIHOME/log/HOSTNAME1/crsd` • `GIHOME/log/HOSTNAME1/cssd` • `GIHOME/log/HOSTNAME1/ctssd` • `GIHOME/log/HOSTNAME1/diskmon` • `GIHOME/log/HOSTNAME1/evmd` • `GIHOME/log/HOSTNAME1/gipcd` • `GIHOME/log/HOSTNAME1/gnsd` • `GIHOME/log/HOSTNAME1/gpnpd` • `GIHOME/log/HOSTNAME1/mdnsd` • `GIHOME/log/HOSTNAME1/ohasd` • `GIHOME/log/HOSTNAME1/racg` • `GIHOME/log/HOSTNAME1/srvm` • `GIHOME/log/HOSTNAME1/xag` • `GIHOME/log/diag/asmtool` • `GIHOME/log/diag/clients` • `GIHOME/log/procwatcher/PRW_SYS_HOSTNAME1` • `GIHOME/network/log` • `GIHOME/opmn/logs` • `GIHOME/racg/log` • `GIHOME/scheduler/log` • `GIHOME/srvm/log` • `GRIDBASE/crsdata/@global/cvu` • `GRIDBASE/crsdata/HOSTNAME1/core` • `GRIDBASE/crsdata/HOSTNAME1/crsconfig` • `GRIDBASE/crsdata/HOSTNAME1/crsdiag` • `GRIDBASE/crsdata/HOSTNAME1/cvu` • `GRIDBASE/crsdata/HOSTNAME1/evm` • `GRIDBASE/crsdata/HOSTNAME1/output` • `GRIDBASE/crsdata/HOSTNAME1/ovmmwallets` • `GRIDBASE/crsdata/HOSTNAME1/scripts` • `GRIDBASE/crsdata/HOSTNAME1/trace` • `GRIDBASE/diag/crs/-/crs/cdump`

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-5 (Cont.) Oracle Trace File Analyze (TFA) Component Driven Logs Collections

Component	Script	Files/Directories
		• GRIDBASE/diag/crs/HOSTNAME1/crs/cdump • GRIDBASE/diag/crs/HOSTNAME1/crs/incident • GRIDBASE/diag/crs/HOSTNAME1/crs/trace
Database: Oracle Database logs	No database specific script - runs opatch lsinventory for the ORACLE_HOME the database runs from TFA will run ipspack based on the time range for certain database incidents.	• ORACLE_BASE/diag/rdbms/<dbname>/<instance_name>/cdump • ORACLE_BASE/diag/rdbms/<dbname>/<instance_name>/trace • ORACLE_BASE/diag/rdbms/<dbname>/<instance_name>/incident

DCS Agent Logs

- /opt/oracle/dcs/log/

Tooling-Related Grid Infrastructure/Database Logs

- **Grid Infrastructure**: GI_HOME/cfgtoollogs
- **Database alertlog**: /u02/app/oracle/diag/rdbms/*/*/alert*.log

Event Types for Base Database Service

Base Database service resources emit events, which are structured messages that indicate changes in resources.

For more information about Oracle Cloud Infrastructure Events, see Overview of Events. You may subscribe to events and be notified when they occur using the Oracle Notification service, see Notifications Overview.

Prerequisites

The following prerequisite is required to receive events for databases and DB systems.

- Telemetry must be enabled for databases and DB systems using the dbcli utility.

  For more information, see AHF Telemetry Commands.

Database Event Types

The following are the event types that the database in the DB system emits:

Table 5-6 Database Event Types

Friendly Name	Event Type
Database - Automatic Backup Begin	com.oraclecloud.databaseservice.automaticbackupdatabase.begin

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-6 (Cont.) Database Event Types

Friendly Name	Event Type
Database - Automatic Backup End	com.oraclecloud.databaseservice.automaticbackupdatabase.end
Database - Create Backup Begin	com.oraclecloud.databaseservice.backupdatabase.begin
Database - Create Backup End	com.oraclecloud.databaseservice.backupdatabase.end
Database - Critical	com.oraclecloud.databaseservice.database.critical
Database - Delete Backup Begin	com.oraclecloud.databaseservice.deletebackup.begin
Database - Delete Backup End	com.oraclecloud.databaseservice.deletebackup.end
Database - Information	com.oraclecloud.databaseservice.database.information
Database - Migrate to KMS Key Begin	com.oraclecloud.databaseservice.migratedatabasekmskey.begin
Database - Migrate to KMS Key End	com.oraclecloud.databaseservice.migratedatabasekmskey.end
Database - Move Begin	com.oraclecloud.databaseservice.movedatabase.begin
Database - Move End	com.oraclecloud.databaseservice.movedatabase.end
Database - Restore Begin	com.oraclecloud.databaseservice.restoredatabase.begin
Database - Restore End	com.oraclecloud.databaseservice.restoredatabase.end
Database - Rotate KMS Key Begin	com.oraclecloud.databaseservice.rotatedatabasekmskey.begin
Database - Rotate KMS Key End	com.oraclecloud.databaseservice.rotatedatabasekmskey.end
Database - Terminate Begin	com.oraclecloud.databaseservice.database.terminate.begin
Database - Terminate End	com.oraclecloud.databaseservice.database.terminate.end
Database - Update Begin	com.oraclecloud.databaseservice.updatedatabase.begin
Database - Update End	com.oraclecloud.databaseservice.updatedatabase.end
Database - Upgrade Begin	com.oraclecloud.databaseservice.upgradedatabase.begin
Database - Upgrade End	com.oraclecloud.databaseservice.upgradedatabase.end

Example 5-1 Database Example Event

This is a reference event for databases.

```
{
    "eventType" : "com.oraclecloud.databaseservice.backupdatabase.begin",
    "cloudEventsVersion" : "0.1",
    "eventTypeVersion" : "2.0",
    "source" : "DatabaseService",
    "eventTime" : "2020-01-08T17:31:43.666Z",
    "contentType" : "application/json",
    "data" : {
        "compartmentId" : "ocid1.compartment.oc1.<unique_ID>",
        "compartmentName": "example_compartment_name",
        "resourceName": "my_backup",
        "resourceId": "ocid1.dbbckup.oc1.<unique_ID>",
        "availabilityDomain": "<availability_domain>",
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
            "additionalDetails" : {
                "timeCreated" : "2020-01-08T17:31:44Z",
                "lifecycleState" : "CREATING",
                "dbSystemId" : "ocid1.dbsystem.oc1.<unique_ID>",
                "dbHomeId" : ocid1.dbhome.oc1.<unique_ID>",
                "dbUniqueName" : DB1115_iad1dv",
                "dbVersion" : "11.2.0.4.190716",
                "databaseEdition" : "ENTERPRISE_EDITION_HIGH_PERFORMANCE",
                "autoBackupsEnabled" : "false",
                "backupType" : "FULL",
                "databaseId" : "ocid1.database.oc1.<unique_ID>",
            },
            "definedTags" : {
                "My_example_tag_name" :
                    { "Example_key" : "Example_value" }
            },
        },
        "eventID": "<unique_ID>",
        "extensions" : {
            "compartmentId": "ocid1.compartment.oc1.<unique_ID>"
        }
    }
```

Database Information Event Details

The `com.oraclecloud.databaseservice.database.information` event delivers details about several types of information in the `additionalDetails` section of the database information event payload.

The following table lists the sub-types of the database information event.

Table 5-7 Database Information Event Details

EventName (in additionalDetails)	Description
`AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN_CLEARED`	An event of type INFORMATION is created once it is determined that the event for the CDB down has cleared.
`HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE_CLEARED`	An event of type INFORMATION is created once it is determined that the event for the backup failure has cleared.

Database Critical Event Details

The `com.oraclecloud.databaseservice.database.critical` event delivers details about several types of critical conditions and errors in the `additionalDetails` section of the Database critical event payload.

The following table lists the sub-types of the database critical event.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-8 Database Critical Event Details

EventName (in additionalDetails)	Description
AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN	A DOWN event is created when a database instance goes down. The event is of type INFORMATION when a database instance is shutdown due to user action, such as with the SQL*Plus (sqlplus) or Server Control Utility (srvctl) commands, or any Oracle Cloud maintenance action that uses those commands, such as performing a database home software update. The event is of type CRITICAL when a database instance goes down unexpectedly. A corresponding DOWN_CLEARED event is created when a database instance is started.
HEALTH.DB_CLUSTER.CDB.ARCHIVER_HANG	An event of type CRITICAL is created if a CDB is either unable to archive the active online redo log or unable to archive the active online redo log fast enough to the log archive destinations.
HEALTH.DB_CLUSTER.CDB.DATABASE_HANG	An event of type CRITICAL is created when a process or session becomes unresponsive in the CDB.
HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE	An event of type CRITICAL is created if there is a CDB backup with a FAILED status reported in the v$rman_status view.
HEALTH.DB_CLUSTER.CDB.CORRUPTION	Database corruption has been detected on your primary or standby database. The database alert.log is parsed for any specific errors that are indicative of physical block corruptions, logical block corruptions, or logical block corruptions caused by lost writes.

Example 5-2 Database Critical Example Event

The Database Critical, DB Node Critical, and DB System Critical events originate in the data plane and contain details about a critical condition in the additionalDetails section of the payload. See the preceding tables for details about these event subtypes.

The following is a reference "critical" data plane event for DB systems, DB system nodes, and databases.

```
{
    "eventType": "com.oraclecloud.databaseservice.database.critical",
    "cloudEventsVersion": "0.1",
    "eventTypeVersion": "2.0",
    "source": "DataPlane",
    "eventTime": "2020-11-10T19:52:15Z",
    "contentType": "application/json",
    "data": {
        "compartmentId": "ocid1.compartment.oc1.<unique_ID>",
        "compartmentName": "VMDBSI-Dev",
        "resourceName": "DB0422_iad3x7",
        "resourceId": "ocid1.database.oc1.iad.<unique_ID>",
        "availabilityDomain": "zvXp:US-ASHBURN-AD-3",
        "additionalDetails": {
            "serviceType": "dbcs",
            "hostName": "singlenodegi-sales",
            "component": "cdb",
            "instanceName": "db0422",
            "dbName": "db0422_iad3x7",
            "description": "Database : DB0422_iad3x7 Instance : DB0422,
status is offline",
            "eventName": "AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN",
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
            "dbSystemId": "ocid1.dbsystem.oc1.iad.<unique_ID>",
            "status": "offline"
        }
    },
    "eventID": "91653791-7aab-45dd-b57f-e2e9013acdb9",
    extensions": {
        "compartmentId": "ocid1.compartment.oc1.<unique_ID>"
    }
}
```

DB System Event Types

The following are the event types that the DB system emits.

Table 5-9 DB System Event Types

Friendly Name	Event Type
DB System - Change Compartment Begin	`com.oraclecloud.databaseservice.changedbsystemcompartment.begin`
DB System - Change Compartment End	`com.oraclecloud.databaseservice.changedbsystemcompartment.end`
DB System - Create Begin	`com.oraclecloud.databaseservice.launchdbsystem.begin`
DB System - Create End	`com.oraclecloud.databaseservice.launchdbsystem.end`
DB System - Critical	`com.oraclecloud.databaseservice.dbsystem.critical`
DB System - Information	`com.oraclecloud.databaseservice.dbsystem.information`
DB System - Terminate Begin	`com.oraclecloud.databaseservice.terminatedbsystem.begin`
DB System - Terminate End	`com.oraclecloud.databaseservice.terminatedbsystem.end`
DB System - Update IORM Begin	`com.oraclecloud.databaseservice.updateiormconfig.begin`
DB System - Update IORM End	`com.oraclecloud.databaseservice.updateiormconfig.end`

Example 5-3 DB System Example Event

This is a reference event for DB systems.

```
{
    "cloudEventsVersion": "0.1",
    "contentType": "application/json",
    "data": {
        "additionalDetails": {
            "cpuCoreCount": 1,
            "dataStoragePercentage": 80,
            "dataStorageSizeInGBs": 256,
            "exadataIormConfig": "null",
            "licenseType": "LICENSE_INCLUDED",
            "lifecycleMessage": null,
            "lifecycleState": "PROVISIONING",
            "nsgIds": "null",
            "patchHistoryEntries": "null",
            "sshPublicKeys": "...",
            "version": null
        },
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
        "availabilityDomain": "XXIT:US-ASHBURN-AD-1",
        "compartmentId": "ocid1.compartment.oc1.<unique_ID>",
        "compartmentName": "example_compartment_name",
        "resourceId": "ocid1.dbsystem.oc1.iad.<unique_ID>",
        "resourceName": "myDBsystem"
    },
    "eventID": "0c1f15b1-4bf2-4f27-8a78-a48d446aeb6f",
    "eventTime": "2019-10-25T20:30:46.836Z",
    "eventType": "com.oraclecloud.databaseservice.launchdbsystem.begin",
    "eventTypeVersion": "2.0",
    "extensions": {
        "compartmentId": "ocid1.compartment.oc1.<unique_ID>"
    },
    "source": "DatabaseService"
}
```

DB System Information Event Details

The `com.oraclecloud.databaseservice.dbsystem.information` event delivers details about information in the `additionalDetails` section of the information event payload.

The following table documents the sub-types of the DB system information event.

Table 5-10 DB System Information Event Details

EventName (in additionalDetails)	Description
`HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE`	An event of type INFORMATION is created when the ASM disk group space usage drops below 90%.

DB System Critical Event Details

The `com.oraclecloud.databaseservice.dbsystem.critical` event delivers details about critical conditions and errors in the `additionalDetails` section of the critical event payload.

The following table documents the sub-types of the DB system critical event.

Table 5-11 DB System Critical Event Details

EventName (in additionalDetails)	Description
`HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE`	An event of type CRITICAL is created when an ASM disk group reaches space usage of 90% or higher.

For a JSON example of a critical event, see Example 5-2.

DB Node Event Types

The following table lists the event types that database nodes emit.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-12 DB Node Event Types

Friendly Name	Event Type
DB Node - Update Begin	`com.oraclecloud.databaseservice.dbnodeaction.begin`
DB Node - Update End	`com.oraclecloud.databaseservice.dbnodeaction.end`
DB Node - Critical	`com.oraclecloud.databaseservice.dbnode.critical`
DB Node - Information	`com.oraclecloud.databaseservice.dbnode.information`

Example 5-4 DB System Node Example Event

This is a reference event for DB system nodes.

```
{
    "cloudEventsVersion": "0.1",
    "eventID": "<unique_ID>",
    "eventType": "com.oraclecloud.databaseservice.dbnodeaction.begin",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-07-29T04:43:24Z",
    "contentType": "application/json",
    "extensions": {
        "compartmentId": "ocid1.compartment.oc1.<unique_ID>"
    },
    "data": {
        "compartmentId": "ocid1.compartment.oc1.<unique_ID>",
        "compartmentName": "example_compartment",
        "resourceName": "",
        "resourceId": "ocid1.dbnode.oc1.phx.<unique_ID>",
        "availabilityDomain": "TGjA:PHX-AD-2",
        "freeFormTags": null,
        "definedTags": null,
        "additionalDetails": {
            "cpuCoreCount": null,
            "lifecycleState": "STARTING",
            "dataStorageSizeInTBs": null,
            "timeCreated": "2019-06-13T04:31:05.190Z",
            "timeUpdated": "2019-07-29T04:43:06.455Z",
            "hostName": "ora18c",
            "lifecycleDetails": null,
            "dbSystemId": "ocid1.dbsystem.oc1.phx.<unique_ID>",
            "dbHostId": "DbHost-<unique_ID>",
            "nodeNumber": null
        }
    }
}
```

DB Node Information Event Details

The `com.oraclecloud.databaseservice.dbnode.information` event delivers details about several types of information in the `additionalDetails` section of the DB node information event payload.

The following table documents the sub-types of the DB node information event.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-13 DB Node Information Event Details

EventName (in additionalDetails)	Description
AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN_CLEARED	An event of type INFORMATION is created once it is determined that the event for SCAN Listener down has cleared.
AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN_CLEARED	An event of type INFORMATION is created once it is determined that the event for Client Listener down has cleared.
AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN_CLEARED	An event of type INFORMATION is created once it is determined that the event for CRS down has cleared.
HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE	An event of type INFORMATION is created once the VM guest file system free space increases above 10%.

DB Node Critical Event Details

The `com.oraclecloud.databaseservice.dbnode.critical` event delivers details about several types of critical conditions in the `additionalDetails` section of the DB node critical event payload.

The following table documents the sub-types of the DB node critical event.

Table 5-14 DB Node Critical Event Details

EventName (in additionalDetails)	Description
AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN	A DOWN event is created when a SCAN listener goes down. The event is of type INFORMATION when a SCAN listener is shutdown due to user action, such as with the Server Control Utility (`srvctl`) or Listener Control (`lsnrctl`) commands, or any Oracle Cloud maintenance action that uses those commands, such as performing a grid infrastructure software update. The event is of type CRITICAL when a SCAN listener goes down unexpectedly. A corresponding DOWN_CLEARED event is created when a SCAN listener is started. There are three SCAN listeners per cluster called LISTENER_SCAN[1,2,3].
AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN	A DOWN event is created when a client listener goes down. The event is of type INFORMATION when a client listener is shutdown due to user action, such as with the Server Control Utility (`srvctl`) or Listener Control (`lsnrctl`) commands, or any Oracle Cloud maintenance action that uses those commands, such as performing a grid infrastructure software update. The event is of type CRITICAL when a client listener goes down unexpectedly. A corresponding DOWN_CLEARED event is created when a client listener is started. There is one client listener per node, each called LISTENER.
AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN	An event of type CRITICAL is created when the Cluster Ready Service (CRS) is detected to be down.
AVAILABILITY.DB_GUEST.CRS_INSTANCE.EVICTION	An event of type CRITICAL is created when the Cluster Ready Service (CRS) evicts a node from the cluster.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-14 (Cont.) DB Node Critical Event Details

EventName (in additionalDetails)	Description
`HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE`	This event is reported when VM guest file system free space falls below 10% free, as determined by the operating system `df(1)` command, for the following file systems: • / • /u01 • /u02 • /var • /tmp

For a JSON example of a critical event, see Example 5-2.

Oracle Database Home Event Types

The following table lists the event types that Oracle Database Homes emit.

Table 5-15 Oracle Database Home Event Types

Friendly Name	Event Type
DB Home - Create Begin	`com.oraclecloud.databaseservice.createdbhome.begin`
DB Home - Create End	`com.oraclecloud.databaseservice.createdbhome.end`
DB Home - Patch Begin	`com.oraclecloud.databaseservice.patchdbhome.begin`
DB Home - Patch End	`com.oraclecloud.databaseservice.patchdbhome.end`
DB Home - Terminate Begin	`com.oraclecloud.databaseservice.deletedbhome.begin`
DB Home - Terminate End	`com.oraclecloud.databaseservice.deletedbhome.end`
DB Home - Update Begin	`com.oraclecloud.databaseservice.updatedbhome.begin`
DB Home - Update End	`com.oraclecloud.databaseservice.updatedbhome.end`

Example 5-5 Oracle Database Home Example Event

This is a reference event for Database Homes.

```
{
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.createdbhome.begin",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-08-29T21:16:04Z",
    "contentType": "application/json",
    "extensions": {
        "compartmentId": "ocid1.compartment.oc1.<unique_ID>"
    },
    "data": {
        "compartmentId": "ocid1.compartment.oc1.<unique_ID>",
        "compartmentName": "example_compartment",
        "resourceName": "my_dbhome",
        "resourceId": "DbHome-unique_ID",
        "availabilityDomain": "all",
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
        "freeFormTags": {},
        "definedTags": {},
        "additionalDetails": {
            "id": "ocid1.id.oc1.<unique_ID>",
            "lifecycleState": "PROVISIONING",
            "timeCreated": "2019-08-29T12:00:00.000Z",
            "timeUpdated": "2019-08-29T12:30:00.000Z",
            "lifecycleDetails": "detail message",
            "dbSystemId": "DbSystem-unique_ID",
            "dbVersion": "19.0.0.0",
            "recordVersion": 4,
            "displayName": "example_display_name"
        }
    }
}
```

Pluggable Database Event Types

The following table lists the event types that pluggable databases emit.

Table 5-16 Pluggable Database Event Types

Friendly Name	Event Type
Pluggable Database - Create Begin	com.oraclecloud.databaseservice.createpluggabledatabase.begin
Pluggable Database - Create End	com.oraclecloud.databaseservice.createpluggabledatabase.end
Pluggable Database - Delete Begin	com.oraclecloud.databaseservice.deletepluggabledatabase.begin
Pluggable Database - Delete End	com.oraclecloud.databaseservice.deletepluggabledatabase.end
Pluggable Database - Local Clone Begin	com.oraclecloud.databaseservice.localclonepluggabledatabase.begin
Pluggable Database - Local Clone End	com.oraclecloud.databaseservice.localclonepluggabledatabase.end
Pluggable Database - Remote Clone Begin	com.oraclecloud.databaseservice.remoteclonepluggabledatabase.begin
Pluggable Database - Remote Clone End	com.oraclecloud.databaseservice.remoteclonepluggabledatabase.end
Start Pluggable Database - Begin	com.oraclecloud.databaseservice.startpluggabledatabase.begin
Start Pluggable Database - End	com.oraclecloud.databaseservice.startpluggabledatabase.end
Stop Pluggable Database - Begin	com.oraclecloud.databaseservice.stoppluggabledatabase.begin
Stop Pluggable Database - End	com.oraclecloud.databaseservice.stoppluggabledatabase.end

Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 5-6 Pluggable Database Example Event

This is a reference event for pluggable databases (PDBs).

```
{
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
        "compartmentId": "ocid1.compartment.oc1.<unique_ID>"
    },
    "eventType":
"com.oraclecloud.databaseservice.remoteclonepluggabledatabase.begin",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
        "compartmentId": "ocid1.compartment.oc1.<unique_ID>",
        "compartmentName": "MyCompartment",
        "resourceName": "11092020_PKS_PDB1",
        "resourceId": "ocid1.pluggabledatabases.oc1.phx.<unique_ID>",
        "availabilityDomain": "XXIT:PHX-AD-1",
        "freeFormTags": {},
        "definedTags": {},
        "additionalDetails": {
            "id": "ocid1.pluggabledatabases.oc1.phx.<unique_ID>",
            "timeCreated": "2021-03-13T21:15:59.000Z",
            "timeUpdated": "2021-03-13T21:15:59.000Z",
            "databaseId": "ocid1.database.oc1.<unique_ID>",
            "lifecycleState": "AVAILABLE",
            "lifecycleDetails": "Pluggable Database is available",
            "displayName": "Pluggable Database - Remote Clone Begin"
        }
    }
}
```

Data Guard Association Event Types

The following table lists the event types that Data Guard associations emit.

Table 5-17 Data Guard Association Event Types

Friendly Name	Event Type
Change Protection Mode Begin	`com.oraclecloud.databaseservice.changeprotectionmode.begin`
Change Protection Mode End	`com.oraclecloud.databaseservice.changeprotectionmode.end`
Data Guard Association - Create Begin	`com.oraclecloud.databaseservice.createdataguardassociation.begin`
Data Guard Association - Create End	`com.oraclecloud.databaseservice.createdataguardassociation.end`
Data Guard Association - Failover Begin	`com.oraclecloud.databaseservice.failoverdataguardassociation.begin`

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-17 (Cont.) Data Guard Association Event Types

Friendly Name	Event Type
Data Guard Association - Failover End	com.oraclecloud.databaseservice.failoverdataguardassociation.end
Data Guard Association - Reinstate Begin	com.oraclecloud.databaseservice.reinstatedataguardassociation.begin
Data Guard Association - Reinstate End	com.oraclecloud.databaseservice.reinstatedataguardassociation.end
Data Guard Association - Switchover Begin	com.oraclecloud.databaseservice.switchoverdataguardassociation.begin
Data Guard Association - Switchover End	com.oraclecloud.databaseservice.switchoverdataguardassociation.end

Example 5-7 Data Guard Association example event

This is a reference event for Data Guard associations.

```
{
    "cloudEventsVersion": "0.1",
    "contentType": "application/json",
    "data": {
        "additionalDetails": {
            "ApplyLag": null,
            "DGConfigId": "7e8eff2b-a4cd-474a-abd5-940b05c0b1fd",
            "DGConfigState": "null",
            "DatabaseId": "ocid1.database.oc1.iad.<unique_ID>",
            "DbHomeId": "ocid1.dbhome.oc1.iad.<unique_ID>",
            "DbSystemId": "ocid1.dbsystem.oc1.iad.<unique_ID>",
            "LastSyncedTime": null,
            "SyncState": "null",
            "dcsDgUpdateTimestamp": null,
            "lastUpdatedIdentifier": null,
            "lifeCycleMessage": null,
            "lifecycleState": "PROVISIONING",
            "timeCreated": "2019-10-25T21:42:19.041Z",
            "timeUpdated": "2019-10-25T21:42:19.041Z"
        },
        "availabilityDomain": "XXIT:US-ASHBURN-AD-1",
        "compartmentId": "ocid1.compartment.oc1.<unique_ID>",
        "compartmentName": "example_compartment",
        "resourceId": "ocid1.dgassociation.oc1.iad.<unique_ID>"
    },
    "eventID": "5b8b7fbf-2e9a-4730-9761-e52715b7bc79",
    "eventTime": "2019-10-25T21:42:16.579Z",
    "eventType":
"com.oraclecloud.databaseservice.createdataguardassociation.begin",
    "eventTypeVersion": "2.0",
    "extensions": {
        "compartmentId": "ocid1.compartment.oc1.<unique_ID>"
    },
    "source": "DatabaseService"
}
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Remediation for Database Service Events

This article describes the fixes needed for problems encountered while using Database Service Events.

The following remediations are available:

- HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE
- AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN
- AVAILABILITY.DB_GUEST.CRS_INSTANCE.EVICTION
- AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN
- AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN
- AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN
- HEALTH.DB_CLUSTER.CDB.CORRUPTION
- HEALTH.DB_CLUSTER.CDB.ARCHIVER_HANG
- HEALTH.DB_CLUSTER.CDB.DATABASE_HANG
- HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE
- HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE

HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE

Event Name

HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE

Event Description

This event is reported when VM guest file system free space falls below 10% free, as determined by the operating system `df(1)` command, for the following file systems:

- `/`
- `/u01`
- `/u02`
- `/var`
- `/tmp`

Problem Statement

One or more VM guest file systems has free space below 10% free.

Risk

Insufficient VM guest file system free space can cause disk space allocation failure, which can result in wide-ranging errors and failures in Oracle software (Database, Clusterware, Cloud Tooling).

Action/Repair

Oracle Cloud and DCS Agent run automatically to purge old log files and trace files created by cloud tooling to reclaim file system space.

Copyright © 2022, 2026, Oracle and/or its affiliates.

If the automatic file system space reclamation utilities cannot sufficiently purge old files to clear this event, then perform the following actions:

1. Remove unneeded files and/or directories created manually or by customer-installed applications or utilities. Files created by customer-installed software are outside the scope of Oracle's automatic file system space reclamation utilities. The following operating system command, run as the `root` user, is useful for identifying directories consuming excessive disk space:

   ```
   sudo du -hx <file system mount point> | sort -hr
   ```

   Only remove files or directories you are certain can be safely removed.
2. Set the automatic purging policy using cloud tooling. For more information, see Autologcleanpolicy Commands.
3. Open service request to receive additional guidance about reducing file system space use.

AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN

Event Name

AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN

Event Description

An event of type CRITICAL is created when the Cluster Ready Service (CRS) is detected to be down.

Problem Statement

The Cluster Ready Stack is in an offline state or has failed.

Risk

If the CRS is offline on a node, the node cannot provide database services for the application.

Action/Repair

1. Check if CRS was stopped by your administrator, as part of a planned maintenance event, or a scale up or down of local storage
   a. The following patching events will stop CRS
      i. GRID Update
      ii. Update of Guest
      iii. Update of Host
2. If CRS has stopped unexpectedly, the current status can be checked by issuing the `crsctl check crs` command.
   a. If the node is not responding, the VM node may be rebooting. Wait for the node reboot to finish, CRS will normally be started through the `init` process.
3. If CRS is still down, investigate the cause of the failure by referring to the `alert.log` found in `/u01/app/grid/diag/crs/<node_name>/crs/trace`. Review the log entries corresponding to the date/time of the down event and act on any potential remediation.
4. Restart the CRS, by issuing the `crsctl start crs` command.

Copyright © 2022, 2026, Oracle and/or its affiliates.

5. A successful restart of CRS will generate the clearing event: AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN_CLEARED

Clearing Event

AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN_CLEARED

Clearing Event Description

An INFORMATION event is created once the CRS is successfully started.

AVAILABILITY.DB_GUEST.CRS_INSTANCE.EVICTION

Event Name

AVAILABILITY.DB_GUEST.CRS_INSTANCE.EVICTION

Event Description

An event of type CRITICAL is created when the Cluster Ready Service (CRS) evicts a node from the cluster. The CRS `alert.log` is parsed for the `CRS-1632` error indicating that a node is being removed from the cluster.

Problem Statement

The Oracle Clusterware is designed to perform a node eviction by removing one or more nodes from the cluster if some critical problem is detected. A critical problem could be a node not responding via a network heartbeat, a node not responding via a disk heartbeat, a hung or severely degraded machine, or a hung `ocssd.bin` process. The purpose of this node eviction is to maintain the overall health of the cluster by removing unhealthy members.

Risk

During the time it takes to restart the evicted node, the node cannot provide database services for the application.

Action/Repair

A CRS node eviction could be caused by OCSSD (aka CSS daemon), CSSDAGENT or CSSDMONITOR processes. This requires determining which process was responsible for the node eviction and reviewing the relevant log files. Common causes of OCSSD eviction are network failures/latencies, IO issues with CSS voting disks, a member kill escalation. CSSDAGENT or CSSDMONITOR evictions could be OS scheduler problem or a hung thread within CSS daemon. Log files to review include clusterware alert log, cssdagent log, cssdmonitor log, ocssd log, lastgasp log, /var/log/messages, CHM/OS Watcher data, and opatch lsinventory detail.

For more information on collecting files together, see Autonomous Health Framework (AHF) Trace File Analyzer (TFA) & ORAchk/EXAchk . For more information on troubleshooting CRS node eviction, see Troubleshooting Clusterware Node Evictions (Reboots).

AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN

Event Name

AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN

Copyright © 2022, 2026, Oracle and/or its affiliates.

Event Description

A DOWN event is created when a SCAN listener goes down. The event is of type INFORMATION when a SCAN listener is shutdown due to user action, such as with the Server Control Utility (`srvctl`) or Listener Control (`lsnrctl`) commands, or any Oracle Cloud maintenance action that uses those commands, such as performing a grid infrastructure software update. The event is of type CRITICAL when a SCAN listener goes down unexpectedly. A corresponding DOWN_CLEARED event is created when a SCAN listener is started.

There are three SCAN listeners per cluster called LISTENER_SCAN[1,2,3].

Problem Statement

A SCAN listener is down and unable to accept application connections.

Risk

If all SCAN listeners are down, application connections to the database through the SCAN listener will fail.

Action/Repair

Start the SCAN listener to receive the DOWN_CLEARED event.

DOWN event of type INFORMATION

1. If the event was caused by an Oracle Cloud maintenance action, such as performing a grid infrastructure software update, then no action is required. The affected SCAN listener will automatically failover to an available instance.
2. If the event was caused by user action, then start the SCAN listener at the next opportunity.

DOWN event of type CRITICAL

1. Check SCAN status and restart the SCAN listener
   - Login to the VM as `opc` user and `sudo` to the `grid` user:

     ```
     [opc@vm ~] sudo su - grid
     ```

   - Check the SCAN listeners status on any node:

     ```
     [grid@vm ~] srvctl status scan_listener
     ```

   - Start the SCAN listener:

     ```
     [grid@vm ~] srvctl start scan_listener
     ```

   - Recheck the SCAN listeners status on any node: if the `scan_listener` is still down, investigate the cause of the scan listener failure:

Copyright © 2022, 2026, Oracle and/or its affiliates.

a. Collect both the CRS and OS logs 30 minutes prior and 10 minutes for the <hostName> indicated in log. Note the time in the event payload is always provided in UTC: For tfactl collection, adjust the time to the timezone of the VM cluster.

```
[grid@vm ~] tfactl diagcollect -crs -os -node <hostName>
    –from "<eventTime adjusted for local vm timezone> - 30 minute "
    -to "<eventTime adjusted for local vm timezone> + 10 minutes"
```

b. SCAN listener issues are logged:

```
/u01/app/grid/diag/tnslsnr/<hostName>/<listenerName>/trace
```

AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN

Event Name

AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN

Event Description

A DOWN event is created when a client listener goes down. The event is of type INFORMATION when a client listener is shutdown due to user action, such as with the Server Control Utility (srvctl) or Listener Control (lsnrctl) commands, or any Oracle Cloud maintenance action that uses those commands, such as performing a grid infrastructure software update. The event is of type CRITICAL when a client listener goes down unexpectedly. A corresponding DOWN_CLEARED event is created when a client listener is started.

There is one client listener per node, each called LISTENER.

Problem Statement

A client listener is down and unable to accept application connections.

Risk

If the node's client listener is down, the database instances on the node cannot provide services for the application.

If the client listener is down on all nodes, any application that connects to any database using the SCAN or VIP will fail.

Action/Repair

Start the client listener to receive the DOWN_CLEARED event.

DOWN event of type INFORMATION

1. If the event was caused by an Oracle Cloud maintenance action, such as performing a grid infrastructure software update, then no action is required. The affected client listener will automatically restart when maintenance affecting the grid instance is complete.
2. If the event was caused by user action, then start the client listener at the next opportunity.

DOWN event of type CRITICAL

1. Check client listener status and restart the client listener:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Login to the VM as `opc` user and `sudo` to the `grid` user:

  ```
  [opc@vm ~] sudo su - grid
  ```

- Check the client listener status on any node:

  ```
  [grid@vm ~] srvctl status listener
  ```

- Start the client listener:

  ```
  [grid@vm ~] srvctl start listener
  ```

- Recheck the client listener status on any node: if client listener is still down. Investigate the cause of the client listener failure:

  a. Use `tfactl` to collect both the CRS and OS logs 30 minutes prior and 10 minutes for the `hostName` indicated in log. Note the time in the event payload is always provided in UTC: For `tfactl` collection, adjust the time to the timezone of the VM cluster.

  ```
  [grid@vm ~] tfactl diagcollect -crs -os -node <hostName>
      –from "<eventTime adjusted for local vm timezone> - 30 minute "
      -to "<eventTime adjusted for local vm timezone> + 10 minutes"
  ```

  b. Review the listener log located under:

  ```
  /u01/app/grid/diag/tnslsnr/<hostName>/<listenerName>/trace
  ```

AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN

Event Name

AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN

Event Description

A DOWN event is created when a database instance goes down. The event is of type INFORMATION when a database instance is shutdown due to user action, such as with the SQL*Plus (`sqlplus`) or Server Control Utility (`srvctl`) commands, or any Oracle Cloud maintenance action that uses those commands, such as performing a database home software update. The event is of type CRITICAL when a database instance goes down unexpectedly. A corresponding DOWN_CLEARED event is created when a database instance is started.

Problem Statement

A database instance has gone down.

Risk

A database instance has gone down., which may result in reduced performance if database instances are available on other nodes in the cluster, or complete downtime if database instances on all nodes are down.

Action/Repair

Start the database instance to receive the DOWN_CLEARED event.

Copyright © 2022, 2026, Oracle and/or its affiliates.

DOWN event of type INFORMATION

1. If the event was caused by an Oracle Cloud maintenance action, such as performing a database home software update, then no action is required. The affected database instance will automatically restart when maintenance affecting the instance is complete.
2. If the event was caused by user action, then start the affected database instance at the next opportunity.

DOWN event of type CRITICAL

1. Check database status and restart the down database instance.
   a. Login to the VM as `oracle` user:
   b. Set the environment:

      ```
      [oracle@vm ~] . <dbName>.env
      ```

   c. Check the database status:

      ```
      [oracle@vm ~] srvctl status database -db <dbName>
      ```

   d. Start the database instance:

      ```
      [oracle@vm ~] srvctl start instance -db <dbName> -instance
      <instanceName>
      ```

2. Investigate the cause of the database instance failure.
   a. Review Trace File Analyzer (TFA) events for the database:

      ```
      [oracle@vm ~] tfactl events -database <dbName> -instance <instanceName>
      ```

   b. Review the database alert log located at:

      ```
      $ORACLE_BASE/diag/rdbms/<dbName>/<instanceName>/trace/
      alert_<instanceName>.log
      ```

HEALTH.DB_CLUSTER.CDB.CORRUPTION

Event Name

HEALTH.DB_CLUSTER.CDB.CORRUPTION

Event Description

Database corruption has been detected on your primary or standby database. The database `alert.log` is parsed for any specific errors that are indicative of physical block corruptions, logical block corruptions, or logical block corruptions caused by lost writes.

Problem Statement

Corruptions can lead to application or database errors and in worse case result in significant data loss if not addressed promptly.

A corrupt block is a block that was changed so that it differs from what Oracle Database expects to find. Block corruptions can be categorized as physical or logical:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- In a physical block corruption, which is also called a media corruption, the database does not recognize the block at all; the checksum is invalid or the block contains all zeros. An example of a more sophisticated block corruption is when the block header and footer do not match.
- In a logical block corruption, the contents of the block are physically sound and pass the physical block checks; however, the block can be logically inconsistent. Examples of logical block corruption include incorrect block type, incorrect data or redo block sequence number, corruption of a row piece or index entry, or data dictionary corruptions.

Block corruptions can also be divided into interblock corruption and intrablock corruption:

- In an intra-block corruption, the corruption occurs in the block itself and can be either a physical or a logical block corruption.
- In an inter-block corruption, the corruption occurs between blocks and can only be a logical block corruption.

Oracle checks for the following errors in the `alert.log`:

- `ORA-01578`
- `ORA-00752`
- `ORA-00753`
- `ORA-00600` [`3020`]
- `ORA-00600` [`kdsgrp1`]
- `ORA-00600` [`kclchkblk_3`]
- `ORA-00600` [`13013`]
- `ORA-00600` [`5463`]

Risk

A data corruption outage occurs when a hardware, software, or network component causes corrupt data to be read or written. The service-level impact of a data corruption outage may vary, from a small portion of the application or database (down to a single database block) to a large portion of the application or database (making it essentially unusable). If remediation action is not taken promptly, potential downtime and data loss can increase.

Action/Repair

The current event notification currently triggers on physical block corruptions (`ORA-01578`), lost writes (`ORA-00752`, `ORA-00753`, and `ORA-00600` with first argument `3020`) and logical corruptions (typical detected from `ORA-00600` with first argument of `kdsgrp1`, `kdsgrp1`, `kclchkblk_3`, `13013`, or `5463`).

We recommend the following steps:

1. Confirm that these corruptions were reported in the `alert.log` trace file. Log a Service Request (SR) with latest EXAchk report, excerpt of the `alert.log` and trace file containing the corruption errors, any history of recent application, database or software changes and any system, clusterware and database logs for the same time period. For all these cases, a TFA collection should be available and should be attached to the SR.
2. For more information on repair recommendations, see Primary Note for Handling Oracle Database Corruption Issues.

For physical corruptions or `ORA-1578` errors, the following notes will be helpful:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- OERR: ORA-1578 "ORACLE data block corrupted (file # %s, block # %s)" Primary Note (Doc ID 1578.1)
- How to identify all the Corrupted Objects in the Database with RMAN (Doc ID 472231.1)
- How to identify the corrupt Object reported by ORA-1578 / RMAN / DBVERIFY (Doc ID 819533.1)
- Primary Note for Handling Oracle Database Corruption Issues

> **Note**
>
> RMAN can be used to recover one or many data block that are physically corrupted. Also using Active Data Guard with real time apply, auto block repair of physical data corruptions would have occurred automatically.

For logical corruptions caused by lost writes (`ORA-00752`, `ORA-00753`, and `ORA-00600` with first argument `3020`) on the primary or standby databases, they will be detected on the primary or with standby's redo apply process. The following notes will be helpful:

- Primary Note for Handling Oracle Database Corruption Issues
- If you have a standby and lost write corruption on the primary or standby, see Resolving ORA-00752 or ORA-600 [3020] During Standby Recovery (Doc ID 1265884.1).

For logical corruptions (typical detected from `ORA-00600` with arguments of `kdsgrp1`, `kclchkblk_3`, `13013`, or `5463`)

- For more information on the error that was detected, see Primary Note for Handling Oracle Database Corruption Issues.
- If you have a standby and logical corruption on the primary, see Resolving Logical Block Corruption Errors in a Physical Standby Database (Doc ID 2821699.1).

HEALTH.DB_CLUSTER.CDB.ARCHIVER_HANG

Event Name

HEALTH.DB_CLUSTER.CDB.ARCHIVER_HANG

Event Description

An event of type CRITICAL is created if a container database (CDB) is either unable to archive the active online redo log or unable to archive the active online redo log fast enough to the log archive destinations.

Problem Statement

CDB RAC Instance may temporarily or permanently stall due to the log writer's (LGWR) inability to write the log buffers to an online redo log. This occurs because all online logs need archiving. Once the archiver (ARC) can archive at least one online redo log, LGWR will be able to resume writing the log buffers to online redo logs and the application impact will be alleviated.

Risk

If the archiver becomes temporarily unresponsive, applications may experience a brief brownout or stall when attempting to commit database changes. If the archiver is not restored, application processes can experience extended delays in processing.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Action/Repair

- To determine the hourly frequency for each thread/instance, see Script To Find Redolog Switch History And Find Archivelog Size For Each Instances In RAC (Doc ID 2373477.1).
  - If any hourly bucket is greater than 12, consider resizing the online redo logs. See item 2 below for resizing steps.
- If the database becomes temporarily unresponsive, the archiver may be unable to keep up with the redo log generated.
  - Check the `alert.log`, `$ORACLE_BASE/diag/rdbms/<dbName>/<instanceName>/trace/alert_<instanceName>.log`, for "All online logs need archiving", multiple events in a short period can indicate 2 possible solutions.
    1. If the number of redo logs groups per thread is less than 4, consider adding additional logs groups to reach 4, see item1 below for add redo log steps.
    2. The other possible solution is to resize the redo logs, see item 2 below for resizing steps.
- For sizing guidelines for Data Guard and non Data Guard, see Configure Online Redo Logs Appropriately.
- Add a redo log group for each thread. The additional redo log should equal the current log size.
  1. Use the following query:

     ```
     select max(group#) Ending_group_number, thread#,
         count(*) number_of_groups_per_thread,
         bytes redo_size_in_bytes from v$log group by thread#,bytes
     ```

  2. Add one new group per thread using the same size as the current redo logs.

     ```
     alter database add logfile thread <thread_number>
         group <max group + 1> ('<DATA_DISKGROUP>') size <redo_size_in_bytes>
     ```

- Resize the online redo logs by adding larger redo logs and dropping the current smaller redo logs.
  1. Use the following query:

     ```
     select max(group#) Ending_group_number, thread#,
         count(*) number_of_groups_per_thread,
         bytes redo_size_in_bytes from v$log group by thread#,bytes
     ```

  2. Add the same number of redo logs for each thread `number_of_groups_per_thread` that currently exist. The `new_redo_size_in_bytes` should be based on Configure Online Redo Logs Appropriately.

     a.
     ```
     alter database add logfile thread <thread_number>
         group <max group + 1> ('<DATA_DISKGROUP>') size
     <new_redo_size_in_bytes>
     ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

b. The original smaller redo logs should be deleted. A redo log can only be deleted if its status is inactive. To determine the status of a redo logs issue the following select.

```
select group#, thread#, status, bytes from v$log order by bytes,
group#, thread#;
```

Delete the original smaller redo logs:

```
alter database drop logfile <group#>;
```

- If the database becomes unresponsive, the primary log archive destination and alternate may be full.
  - For more information on freeing space in RECO and DATA disk groups, see HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE.

HEALTH.DB_CLUSTER.CDB.DATABASE_HANG

Event Name

HEALTH.DB_CLUSTER.CDB.DATABASE_HANG

Event Description

An event of type CRITICAL is created when a process or session becomes unresponsive in the container database (CDB).

Problem Statement

Blocker Resolver detected an unresponsive process or a block and generated an `ORA-32701` error message. Additionally, this event may be raised if the Diagnostic Process (`DIA0`) detects that a critical database process has become unresponsive.

Risk

An unresponsive process or a block can indicate resource, operating system, or application coding related issues.

Action/Repair

Investigate the cause of the session block.

- Review TFA events for the database for the following message patterns corresponding to the date/time of the event: `ORA-32701`, "*DIA0 Critical Database Process Blocked*", or "*DIA0 Critical Database Process As Root*".

  ```
  tfactl events -database <dbName> -instance <instanceName>
  ```

- Review the alert.log associated at the following location:

  ```
  $ORACLE_BASE/diag/rdbms/<dbName>/<instanceName>/trace/
  alert_<instanceName>.log
  ```

- For `ORA-32701`: An overloaded system can cause slow progress, which can be interpreted as a block. The Blocker Resolver may attempt to resolve the block by terminating the final blocker process.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- For DIA0 Critical Database Process messages: Review the related diagnostic lines indicating the process and the reason for the block.

HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE

Event Name

HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE

Event Description

An event of type CRITICAL is created if there is a CDB backup with a FAILED status reported in the v$rman_status view.

Problem Statement

A daily incremental BACKUP of the CDB failed.

Risk

A failure of the backup can compromise the ability to use the backups for restore/recoverability of the database. Recoverability Point Object (RPO) and the Recoverability Time Object (RTO) can be impacted.

Action/Repair

Review the RMAN logs corresponding to the date/time of the event. Note the event time stamp eventTime is in UTC, adjust as necessary for the VM's timezone.

For Oracle managed backups:

- RMAN output can be found at /opt/oracle/dcs/log/<hostname>/rman.
- Review the log for any failures:
  - If the failure is due to an external event outside of RMAN, for example the backup location was full or a networking issue, resolve the external issue.
  - For other RMAN script errors, collect the diagnostic logs and open a Service Request.

```
dbcli collect-diagnostics -h


Usage: collect-diagnostics [options]
  Options:
    --components, -c
       Supported components: [all, dcs, crs, acfs, asm, db]
       all -- Collects diagnosis for all supported components [all,
dcs, crs, acfs, asm, db]
       dcs -- Collects diagnosis for dcs
       crs -- Collects diagnosis for crs
       acfs -- Collects diagnosis for acfs
       asm -- Collects diagnosis for asm.
       db -- Collects diagnosis for db.
       For multiple parameter values, follow the example as "-c c1 c2"
       Default: [dcs]
    --dbNames, -d
       Comma separated database names. Valid only if 'db' or 'all'
specified in Components list.
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
    --endTime, -et
       End time of diagnostic logs. Please give time in yyyy-MM-dd
HH:mm:ss format
    --help, -h
       get help
    --json, -j
       json output
    --objectstoreuri, -ou
       Pre Authenticated Request URI
    --redaction, -r
       Diagnostic logs redaction. Might take longer time with some
components.
    --startTime, -st
    Start time of diagnostic logs. Please give time in yyyy-MM-dd
HH:mm:ss format
```

- If the issue is transient or is resolved, then take a new incremental backup. For more information, see *Back Up a Database Using the Console*.

For customer owned and managed backup taken through RMAN:

- Review the RMAN logs for the backup.

HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE

Event Name

HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE

Event Description

An event of type CRITICAL is created when an ASM disk group reaches space usage of 90% or higher. An event of type INFORMATION is created when the ASM disk group space usage drops below 90%.

Problem Statement

ASM disk group space usage is at or exceeds 90%.

Risk

Insufficient ASM disk group space can cause database creation failure, tablespace and data file creation failure, automatic data file extension failure, or ASM rebalance failure.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Action/Repair

ASM disk group used space is determined by the running the following query while connected to the ASM instance.

```
sudo su - grid
sqlplus / as sysasm


select 'ora.'||name||'.dg', total_mb, free_mb,
    round ((1-(free_mb/total_mb))*100,2) pct_used from v$asm_diskgroup;


NAME             TOTAL_MB    FREE_MB   PCT_USED
---------------- ---------- ---------- ----------
ora.DATAC1.dg     75497472    7408292      90.19
ora.RECOC1.dg     18874368   17720208       6.11
```

ASM disk group capacity can be increased in the following ways:

1. Scale the VM Cluster storage to add more ASM disk group capacity. For more information, see the Scale Up the Storage for a Virtual Machine DB System.

DATA disk group space use can be reduced in the following ways:

1. Drop unused data files and temp files from databases. For more information, see Dropping Data Files.

RECO disk group space use can be reduced in the following ways:

1. Drop unnecessary Guaranteed Restore Points. For more information, see Using Normal and Guaranteed Restore Points.
2. Delete archived redo logs or database backups already backed up outside the Flash Recovery Area (FRA). For more information, see Maintaining the Fast Recovery Area.

Back Up and Recovery

Back Up and Recovery in Base Database Service

Backing up your database is critical to ensure the safety of your data. Oracle offers several methods for backing up your database.

The Oracle-managed automatic backups feature is the preferred method for backing up Oracle Cloud databases because you can easily configure backup settings using the Console. The automatic backups feature supports Recovery Service and Object Storage as the backup destination to provide you with a fully automated cloud backup solution with the same cost. You do not need to perform any manual backups or backup storage administration tasks. You can also store backups in local storage. Each backup destination has its advantages and requirements that you should consider, as described below.

Recovery Service (Recommended)

A fully managed service based on the on-premises Oracle's Zero Data Loss Recovery Appliance technology which offers modern cybersecurity protection for Oracle Databases. Unique, automated capabilities protect Oracle Database changes in real time, validate backups

Copyright © 2022, 2026, Oracle and/or its affiliates.

without production database overhead, and enable fast, predictable recovery to any point in time.

If your backups are currently configured with Object Storage, you can seamlessly transition to Recovery Service to achieve advanced capabilities with the same cost.

For more information on Recovery Service, see About Oracle Database Autonomous Recovery Service.

Object Storage

A secure, scalable, on-demand storage solution for databases.

For more information on Object Storage, see Overview of Object Storage.

Local Storage

- Backups are stored locally in the Fast Recovery Area of the DB system.
- Durability: Low
- Availability: Medium
- Back Up and Recovery Rate: High
- Advantages: Optimized back up and fast point-in-time recovery.
- Disadvantages: If the DB system becomes unavailable, the backup is also unavailable.

Currently, Oracle does not provide the ability to attach block storage volumes to a DB system, so you cannot back up to network attached volumes.

For unmanaged backups, you can use `RMAN` or `dbcli`, and you must create and manage your own Object Storage buckets for backups.

> **Note**
>
> If you previously used `RMAN` or `dbcli` to configure backups and then you switch to using the Console or the API for backups, a new backup configuration is created and associated with your database. This means that you can no longer rely on your previously configured unmanaged backups to work.

For detailed instructions on creating backups, see:

- *Back Up a Database Using the Console*
- *Back Up a Database to Object Storage Using RMAN*

For detailed instructions on recovering from backups, see:

- *Recover a Database Using the Console*
- *Recover a Database from Object Storage Using RMAN Backup*
- *Recover a Database from the OCI Classic Object Store*

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be granted security access in a **policy** by an administrator. This access is required whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you get a message that you don’t have permission or are

Copyright © 2022, 2026, Oracle and/or its affiliates.

unauthorized, verify with your administrator what type of access you have and which **compartment** to work in.

If you're new to policies, see Getting Started with Policies and Common Policies.

Prerequisites

Review and ensure that the following prerequisites are met for the back up and recovery operation:

Recovery Service

For detailed information about Recovery Service prerequisites, see *Configure Recovery Service*.

Object Storage

- The DB system requires access to the Object Storage including connectivity to the applicable Swift endpoint. Oracle recommends using a service gateway with the VCN to enable this access. See *VCN and Subnets*.
- An existing Object Storage bucket to use as the backup destination. You can use the Console or the Object Storage API to create the bucket. See Managing Buckets.
- An auth token generated by OCI. You can use the Console or the IAM API to generate the password. See Managing User Credentials.
- The user name specified in the backup configuration file must have tenancy-level access to Object Storage. An easy way to do this is to add the user name to the Administrators group. However, that allows access to all of the cloud services. Instead, an administrator should create a policy like the following that limits access to only the required resources in Object Storage for backing up and restoring the database:

  ```
  Allow group <group_name> to manage objects in compartment
  <compartment_name> where target.bucket.name = '<bucket_name>'
  Allow group <group_name> to read buckets in compartment <compartment_name>
  ```

  For more information about adding a user to a group, see Managing Groups.

For more information on Object Storage, see Overview of Object Storage.

General Information

Your database and DB system must be in an "Available" state for a backup operation to run successfully. Oracle recommends that you avoid performing actions that could interfere with availability (such as patching and Data Guard operations) while a backup operation is in progress. If an automatic backup operation fails, the Database service retries the operation during the next day's backup window. If an on-demand full backup fails, you can try the operation again when the DB system and database availability are restored.

In addition to the prerequisites listed, ensure that the following conditions are met to avoid backup failures:

- The database's archiving mode is set to `ARCHIVELOG` (the default).
- The `/u01` directory on the database host file system has sufficient free space for the execution of backup processes.
- The .bash_profile file for the oracle user does not include any interactive commands (such as `oraenv` or one that could generate an error or warning message).

Copyright © 2022, 2026, Oracle and/or its affiliates.

- (For automatic backups) No changes were made to the default `WALLET_LOCATION` entry in the sqlnet.ora file.
- No changes were made to `RMAN` backup settings by using standard `RMAN` commands.

For more information on problems that can result from not following these guidelines, see *Troubleshoot Backup Failures*.

Managed Backup Features

The following information applies to managed backups configured using the Console or API.

> **Note**
>
> Databases in a security zone compartment must have automatic backups enabled. For a full list of policies that affect Base Database Service resources, see Security Zone Policies.

Currently, the following two types of automatic backups are supported:

- **Recovery Service**: A centralized, fully managed, and standalone backup solution for databases.
- **Object Storage**: A secure, scalable, on-demand storage solution for databases.

To align with the Oracle recommended practice of using SYSBACKUP administrative privileges for backup and recovery operations, cloud automation creates a common administrative user named C##DBLCMUSER with the SYSBACKUP role at the CDB$ROOT container level. Backup and recovery operations are therefore performed with the user having the least required privileges. Credentials for this user are randomly generated and securely managed by cloud automation. If the user is not found or is LOCKED and EXPIRED, then cloud automation will recreate or unlock this user during the backup or recovery operations.

Automatic Incremental and Archived Redo Log Backups

When you enable the Object Storage backup feature for a database, the service creates the following on an on-going basis:

- Weekly level 0 backup, generally created on a specified weekend day. A level 0 backup is the equivalent of a full backup. Note that in the Console, weekly level 0 backups appear in the list of backups with backup type "incremental", as do the daily level 1 backups.
- Daily level 1 backups, which are incremental backups created on each day for the six days following the level 0 backup day.
- Level 0 and level 1 backups are stored in Object Store and have an assigned OCID.
- Ongoing archived redo log backups (with a minimum frequency of every 60 minutes). The **Last backup time** field on the database details page in the OCI Console displays the time of the last archived redo logs. This backup differs from the level 0 and level 1 automatic backups in that it is based on log data and does not have an assigned OCID. The last archived redo log backup can be used to create a new database or to recover a database with minimal data loss.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Backup Retention

If you choose to enable automatic backups, you can choose from one of the provided retention periods or a custom policy. The system automatically deletes your incremental backups at the end of your chosen retention period.

The following retention periods are available for **Recovery Service**. The retention periods (in days) are defined in the Recovery Service protection policy.

- Bronze (14 days)
- **Silver (35 days)** (default)
- Gold (65 days)
- Platinum (95 days)
- Custom (User defined protection policy)

The following retention periods are available for **Object Storage**.

- 7 days
- 15 days
- **30 days** (default)
- 45 days
- 60 days

Long-Term Retention Backup with Recovery Service

Long-term retention backup (LTR) allows you to store full backups for periods up to ten years for compliance, regulatory, or other business needs with complete LTR lifecycle management and immutability.

For LTR with Recovery Service, the retention period should be for a minimum of **90 days** and a maximum of **3650 days** or **ten years** from when the backup was created.

To create an LTR backup with the required retention period, Recovery Service does not require creating a new full production backup but does so by utilizing already existing operational backups in the system within the defined recovery window in the policy. For detailed steps, see Create an On-Demand Backup of a Database.

You can change the retention period for a specific existing LTR backup within the retention period. For detailed steps, see Change the Retention Period of a Long-Term Retention Backup with Recovery Service.

You can restore an LTR backup to create a new database within the retention period. For detailed steps, see *Create a DB System from a Backup*.

Upon DB system termination, the LTR backups will be deleted as per the 'Deletion options after database termination' value.

- **Delete backups in 72 hours**: All backups, including LTR backups, will be deleted.
- **Delete based on policy**: LTR backups will be retained according to the retention policy of each LTR backup.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> Oracle recommends choosing the '**Delete based on policy**' option while terminating a DB system to ensure the LTR backups are retained.

Consider the following additional factors for LTR backups:

- LTR backups will continue to exist independently of any automatic backups configured on the database.
- **LTR backups will be automatically deleted after the specified retention period ends.**
- In-place restore is not supported for LTR.
- For databases in a Data Guard configuration, LTR backup will be created only for the database where it is requested.
- The database should be in the AVAILABLE state to create an LTR.
- LTR is supported for databases with file-based TDE or KMS-based keystores.
- All encryption keys will be maintained for the entire retention period of the LTR.
- You can cancel an LTR if it is in the 'creating' state.
- You can delete an LTR backup at any time after it has been created.
- During restore, if the LTR backup being restored is a supported DB Home major version, it will be restored to the latest RU of that version.
- During restore, if the LTR backup being restored is not a supported DB Home major version, it will be restored to any of the supported major versions, upon which the database must be upgraded to any of the supported major versions.

Restore Options

The following restore options are available for the database. This is not supported for LTR.

- **Restore to the latest**: Restores the database to the last known good state with the least possible data loss.
- **Restore to a timestamp**: Restores the database to the timestamp specified.
- **Restore to SCN**: Restores the database using the System Change Number (SCN) specified. This SCN must be valid.

> **Note**
>
> You can determine the SCN number to use either by accessing and querying your database host or by accessing any online or archived logs.

Protection Policies

Recovery Service uses protection policies to control database backup retention in Oracle Cloud.

Protection Policies provide automated retention management for protected databases, satisfying requirements for regulated environments. Each protected database must be associated with one protection policy.

Copyright © 2022, 2026, Oracle and/or its affiliates.

A protection policy determines the maximum period (in days) allowed to retain backups created by Recovery Service. Based on your business requirements, you can assign separate policies for each protected database or use a single policy across all protected databases in a VCN.

For more information, see Managing Protection Policies.

Protected Databases

Protected database refers to an Oracle Cloud database that uses Recovery Service for backup operations.

For more information, see Managing Protected Databases.

Real-time Data Protection

Recovery Service offers real-time data protection so that you can recover a database within a few sub-seconds of a database failure.

Real-time protection is the continuous transfer of redo changes from a protected database to **Recovery Service**. This reduces data loss and provides a recovery point objective (RPO) near 0. This is an extra cost option.

For more information, see Real-time Data Protection.

Backup Deletion Options After Database Termination

When you terminate a database, all of its resources are deleted, along with any automatic backups. Managed backups using the **Recovery Service** and **Object Storage** destination will be deleted according to the retention policy options selected.

You can use the following options to retain managed database backups after the database is terminated. These options can also help restore the database from backups in case of accidental or malicious damage to the database.

- **Retain backups according to the retention period**: When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be removed at the end of the specified retention period.
- **Retain backups for 72 hours, then delete**: When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be retained for 72 hours and then deleted. The backups are retained for 72 hours to safeguard against accidental deletion by the user.

Backup Scheduling

For Recovery Service backups, the automatic backup process starts at any time or within the assigned window.

For Object Storage backups, the automatic backup process used to create level 0 and level 1 backups can run at any time within the daily backup window (between midnight and 6:00 AM). You can optionally specify a 2-hour scheduling window for your database during which the automatic backup process will begin. There are 12 scheduling windows to choose from, each starting on an even-numbered hour (for example, one window runs from 4:00-6:00 AM, and the next from 6:00-8:00 AM). Backups jobs do not necessarily complete within the scheduling window.

For Object Storage backups, the default backup window of 00:00 to 06:00 in the time zone of the DB system's region is assigned to your database if you do not specify a window. Note that

Copyright © 2022, 2026, Oracle and/or its affiliates.

the default backup scheduling window is six hours long, while the windows you specify are two hours long.

Consider the following factors while scheduling backups.

- **Backup window time zone**: Automatic backups enabled for the first time after November 20, 2018 on any database will run between midnight and 6:00 AM in the time zone of the DB system's region. If you have enabled automatic backups on a database before this date, the backup window for the database will continue to be between midnight and 6:00 AM **UTC**. You can create a My Oracle Support service request to have your automatic backups run in a backup window of your choice.
- **Data Guard**: In a Data Guard association, you can configure automatic backups and create backups of the primary database. However, you cannot configure automatic backups or create backups of the standby database. Also, after a switchover operation, you must again configure automatic backups for the database that has assumed the primary role in the Data Guard association.
- **Retention period changes**: If you shorten your database's automatic backup retention period in the future, existing backups falling outside the updated retention period are deleted by the system.
- **Object Storage costs**: Automatic backups incur Object Storage usage costs.

On-Demand Full Backups

You can create a full backup of your database at any time unless your database is assuming the standby role in a Data Guard association.

Standalone Backups

When you terminate a DB system or a database, all of its resources are deleted, along with any automatic backups. Managed backups using the Recovery Service and Object Storage destination will be deleted according to the retention policy options selected. Full backups remain in Object Storage as standalone backups. You can use a standalone backup to create a new database.

> **Note**
>
> - The list of backups you see in the Console does not include any unmanaged backups (backups created directly by using `RMAN` or `dbcli`).
> - All backups are encrypted with the same master key used for Transparent Data Encryption (TDE) wallet encryption.

Cancel a Running Full or Incremental Backup

You can cancel an ongoing backup, allowing you to free up system resources. As part of the Create Database workflow and independently (after the database has been created), you may enable automatic backup and select the desired backup destination. Depending on the backup destination selected, you may have one or more full backups and several incremental backups. Once any of these backups have started, you have the option to cancel them midway. You can cancel any running backup (automatic or standalone) from the Console or the API. You can cancel a manual backup, which is triggered when you click the Create backup button. You can also delete a canceled manual backup.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Backup and Restore from a Standby Database in a Data Guard Association

You can backup and restore from a standby database in a Data Guard association. By using this feature, you can offload backups to the standby database, thereby freeing up resources in the primary database.

Before you begin, note the following:

- You can use the Recovery Service or Object Storage to backup the standby databases.
- You can schedule automatic backups and configure retention periods and backup schedules on the standby database.
- You can create a database in another availability domain (AD) within the same region or a different region from a backup of the standby database.
- Backups can be configured on only the primary database, on only the standby database, or on both primary and standby databases. However, when configured, both the primary and standby databases must have the same backup destination.
- For backups in the Recovery Service, the primary database can be restored or recovered from the backups of either the standby database or the primary database. Similarly, the standby database can be restored or recovered from the backups of either the primary database or the standby database.
- For backups in the Object Storage, the primary database and standby database can be restored or recovered from their respective backups only.
- The backup destination of the primary database and standby database in a Data Guard association must be the same. For example, if the backup destination of the primary database is Recovery Service, then the backup destination of the standby database must also be Recovery Service. Similarly, if the backup destination of the standby database is Recovery Service, then the backup destination of the primary database must also be Recovery Service.
- Backup destination can be changed only after disabling the backup on either the primary or standby database in a Data Guard association. For example, if the backup destination of both the primary and standby databases is Object Storage and you want to change the backup destination of the primary database to Recovery Service, you must first disable the backup on the standby database.
- If automatic backup was configured on the primary, upon switchover, the backups would continue on the new standby database.
- If automatic backup was configured on the standby, upon failover, the backups will continue on the new primary database, and the backups will be disabled on the new standby database.

For detailed steps to configure automatic backups using the Console, see Configure Automatic Backups for a Standby Database.

Cross Region Restore

You can use an existing backup and restore it to create a database (out-of-place restore) across availability domains within the same region or in a different region using a backup created with Object Storage or Autonomous Recovery Service.

You can also restore a backup taken on either a database that was configured using host-based wallets (local wallets) or OCI Vault. For more information on database encryption keys, see *Manage Encryption Keys*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

The following prerequisites are required for cross region restore (within the same tenancy) for the Oracle Database Autonomous Recovery Service.

1. **VCN peering:** Both the VCNs in local and remote regions must be peered across regions.
2. Add security rules on the source and target VCNs.
   a. Add Ingress rules on the source.
      i. Click **Add Ingress Rule**, and add these details to set up a rule that allows HTTPS traffic from anywhere:
         - Source Type: CIDR
         - Source CIDR: Specify the CIDR of the VCN where the database resides.
         - IP Protocol: TCP
         - Source Port Range: All
         - Destination Port Range: 8005
         - Description: Specify an optional description of the ingress rule to help manage the security rules.
      ii. Click Add Ingress Rule, and add these details to set up a rule that allows SQLNet traffic from anywhere:
         - Source Type: CIDR
         - Source CIDR: Specify the CIDR of the VCN where the database resides.
         - IP Protocol: TCP
         - Source Port Range: All
         - Destination Port Range: 2484
         - Description: Specify an optional description of the ingress rule to help manage the security rules.
      iii. Click Add Ingress Rule, and add these details to set up a rule that allows HTTPS traffic from anywhere:
         - Source Type: CIDR
         - Source CIDR: Specify the CIDR of the target VCN.
         - IP Protocol: TCP
         - Source Port Range: All
         - Destination Port Range: 8005
         - Description: Specify an optional description of the ingress rule to help manage the security rules.
      iv. Click Add Ingress Rule, and add these details to set up a rule that allows SQLNet traffic from anywhere:
         - Source Type: CIDR
         - Source CIDR: Specify the CIDR of the target VCN.
         - IP Protocol: TCP
         - Source Port Range: All
         - Destination Port Range: 2484

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Description: Specify an optional description of the ingress rule to help manage the security rules.

b. Add Egress rules on the target. These are optional if the egress traffic is opened for all IPs and ports.

   i. Click Add Egress Rule, and add these details to set up a rule that allows HTTPS traffic from anywhere:

      - Source Type: CIDR
      - Source CIDR: Specify the CIDR of the source VCN.
      - IP Protocol: TCP
      - Source Port Range: All
      - Destination Port Range: 8005
      - Description: Specify an optional description of the ingress rule to help manage the security rules.

   ii. Click Add Egress Rule, and add these details to set up a rule that allows SQLNet traffic from anywhere:

      - Source Type: CIDR
      - Source CIDR: Specify the CIDR of the source VCN.
      - IP Protocol: TCP
      - Source Port Range: All
      - Destination Port Range: 2484
      - Description: Specify an optional description of the ingress rule to help manage the security rules.

> **Note**
>
> Ensure that the Recovery Service subnet is present in the source region and is attached to the source VCN. For more information, see Creating a Recovery Service Subnet in the Database VCN.

3. Perform DNS peering between local and remote VCNs.

Audit and Trace File Retention for Databases Using Automatic Backups

Oracle Database writes audit and trace files to your database's local storage in the `/u01` directory. These files are retained for 30 days by default, though you can change this interval. Once a day, audit and trace files older than 30 days (or the user-specified interval, if applicable) are discarded by a Oracle Scheduler job. You can also disable the Scheduler job if you want to retain these files permanently. Use the following `dbcli` commands to make changes to this Scheduler job.

- To change the retention period from the default setting of 30 days:

```
dbcli update-database -in <dbName> -lr <number_of_days_to_retain_files>
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

For example:

```
dbcli update-database -in inventorydb -lr 15
```

- To disable the daily discard Scheduler job for older audit and trace files:

```
dbcli update-schedule -i <schedulerID> -d
```

For example:

```
dbcli update-schedule -i 5678 -d
```

Use the API

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 database backups:

- ListBackups
- GetBackup
- CreateBackup
- DeleteBackup
- RestoreDatabase
- UpdateDatabase - To enable and disable automatic backups.
- CreateDbHome - For creating a DB system database from a standalone backup.

For the complete list of APIs for the Database service, see Database Service API.

Configure Recovery Service

This article provides information about configuring the Oracle Database Autonomous Recovery Service for use with the Oracle Base Database Service.

Oracle Database Autonomous Recovery Service is a fully managed, standalone, and centralized cloud backup solution for Oracle Cloud Infrastructure (OCI) databases.

For more information about Recovery Service, see About Oracle Database Autonomous Recovery Service.

Prerequisite Configuration Checklist for Recovery Service

> **Note**
>
> The Data Guard broker is required when configuring Recovery Service as a backup destination for Data Guard enabled databases.

This checklist describes the prerequisite tasks that you must complete before you can use Recovery Service as the backup destination for the databases in your tenancy.

Copyright © 2022, 2026, Oracle and/or its affiliates.

1. Assign Policies to Allow Access to Recovery Service and Related Resources
2. Review Networking Service Permissions to Configure a Subnet
3. Review Subnet Size Requirements and Security Rules for Recovery Service Subnet
4. Create a Recovery Service Subnet in the Database VCN
5. Register Recovery Service Subnet
6. Ensure that the Recovery Service Subnet Can Communicate with Oracle Services
7. Ensure that Your Database Has TDE Fully Configured
8. Turn Off Any Manual Operational Backups

Assign Policies to Allow Access to Recovery Service and Related Resources

Assign policy statements such that the supported OCI Database Services can use Recovery Service for data protection.

In the Console, use the Policy Builder to quickly assign the policies required to use Recovery Service in your tenancy. In the **Policy Builder**, select **Autonomous Recovery Service** as the **Policy Use Case**, and then select these predefined policy templates:

- Ability to do all things with Autonomous Recovery Service
- Let users manage protection policies in Autonomous Recovery Service
- Let users manage Autonomous Recovery Service subnets

Ability to do all things with Autonomous Recovery Service

The **Ability to do all things with Autonomous Recovery Service** policy template includes all the policy statements required to provide permissions for the supported database services to use Recovery Service, and for Recovery Service to use the network resources to access databases in a VCN.

You can either select the policy template or add these policy statements using the manual editor in the Policy Builder.

Table 5-18 Policy Statements Required for Using Recovery Service

Policy Statement	Create In	Purpose
`Allow service database to manage recovery-service-family in tenancy`	Root compartment	Enables the OCI Database Service to access protected databases, protection policies, and Recovery Service subnets within your tenancy.
`Allow service database to manage tagnamespace in tenancy`	Root compartment	Enables the OCI Database Service to access the tag namespace in a tenancy.
`Allow service rcs to manage recovery-service-family in tenancy`	Root compartment	Enables Recovery Service to access and manage protected databases, Recovery Service subnets, and protection policies within your tenancy.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-18 (Cont.) Policy Statements Required for Using Recovery Service

Policy Statement	Create In	Purpose
`Allow service rcs to manage virtual-network-family in tenancy`	Root compartment	Enables Recovery Service to access and manage the private subnet in each database VCN within your tenancy. The private subnet defines the network path for backups between a database and Recovery Service.
`Allow group admin to manage recovery-service-family in tenancy`	Root compartment	Enables users in a specified group to access all Recovery Service resources. Users belonging to the specified group can manage protected databases, protection policies, and Recovery Service subnets.

Let users manage protection policies in Autonomous Recovery Service

The **Let users manage protection policies in Autonomous Recovery Service** policy template grants permissions for users in a specified group to create, update, and delete protection policy resources in Recovery Service.

You can either select the policy template or add this policy statement using the manual editor in the Policy Builder.

Table 5-19 Policy Statement for Managing Protection Policies

Policy Statement	Create In	Purpose
`Allow group {group name} to manage recovery-service-policy in compartment {location}`	Compartment that owns the protection policies.	Enables all users in a specified group to create, update, and delete protection policies in Recovery Service.

Consider this example. This policy grants the `RecoveryServiceUser` group with the permissions to create, update, and delete protection policies in `ABC` compartment.

```
Allow group RecoveryServiceUser to manage recovery-service-policy in
compartment ABC
```

Let users manage Autonomous Recovery Service Subnets

The **Let users manage Autonomous Recovery Service subnets** policy template grants permissions for users in a specified group to create, update, and delete Recovery Service subnet resources.

You can either select the policy template or add this policy statement in the Policy Builder.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-20 Policy Statement for Managing Recovery Service Subnets

Policy Statement	Create In	Purpose
`Allow Group {group name} to manage recovery-service-subnet in compartment {location}`	Compartment that owns the Recovery Service subnets.	Enables all users in a specified group to create, update, and delete Recovery Service subnets.

Consider this example. This policy grants the `RecoveryServiceAdmin` group with the permissions to manage Recovery Service subnets in `ABC` compartment.

```
Allow group RecoveryServiceAdmin to manage recovery-service-subnet in
compartment ABC
```

For more information about policies, see Managing Policies.

About Using a Private Subnet for Recovery Service

Recovery Service uses a private subnet inside a virtual cloud network (VCN) where your database resides. The private subnet defines the network path for backups between your database and Recovery Service.

Oracle recommends that your database VCN must have a single private subnet dedicated for backups to Recovery Service. Your Oracle Cloud database can reside in the same private subnet used by Recovery Service, or in a different subnet within the same VCN.

You can either create a private subnet or use a preexisting subnet in your database VCN. Oracle recommends that you use a subnet size of /24 (256 IP addresses).

> **Note**
>
> Select an IPv4-only subnet for Recovery Service in your database VCN. Do not select an IPv6-enabled subnet as Recovery Service does not support using an IPv6-enabled subnet.

The database VCN requires security rules to allow backup traffic between your database and Recovery Service. Security rules must include stateful ingress rules to allow destination ports 8005 and 2484. You can use these Networking service features to implement security rules:

- **Security Lists**: A security list allows you to add security rules at the subnet level. In your database VCN, select the security list that is used for the Recovery Service subnet, and add the ingress rules to allow destination ports 8005 and 2484.
- **Network Security Groups**: Network security groups (NSGs) enable granular control over security rules that apply to individual VNICs in a VCN. Recovery Service supports these options to configure security rules using NSGs:
  - To implement network isolation, create one NSG for the database VNIC (add egress rules to allow ports 2484 and 8005) and a separate NSG for Recovery Service (add ingress rules to allow ports 2484 and 8005).
  - Create and use a single NSG (with egress and ingress rules) for the database VNIC and Recovery Service.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> If you have configured a security list and an NSG within your database VCN, then the rules defined in the NSGs takes precedence over the rules defined in a security list.

After you create a private subnet in the database VCN, assign the security rules and then register the subnet as a Recovery Service subnet in Recovery Service. If you have created NSGs to implement security rules, then you must also ensure to associate the Recovery Service NSG with the Recovery Service subnet.

> **Note**
>
> Oracle recommends using a private subnet for your backups. However, it is possible to use a public subnet.

For more information, see:

- Security Lists
- Network Security Groups
- VCN and Subnet Management

Review Subnet Size Requirements and Security Rules for Recovery Service Subnet

The security rules are necessary to allow backup traffic between a database and Recovery Service.

> **Note**
>
> Select an IPv4-only subnet for Recovery Service in your database VCN. Do not select an IPv6-enabled subnet as Recovery Service does not support using an IPv6-enabled subnet.

Table 5-21 Subnet size requirements and ingress rules for a private subnet used by Recovery Service

Item	Requirements
Recommended subnet size	/24 (256 IP addresses)
General ingress rule 1: Allow HTTPS traffic from Anywhere	This rule allows backup traffic from your OCI Database to Recovery Service. • **Stateless**: No (all rules must be stateful) • **Source Type**: CIDR • **Source CIDR**: CIDR of the VCN where the database resides • **IP Protocol**: TCP • **Source Port Range**: All • **Destination Port Range**: 8005

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 5-21 (Cont.) Subnet size requirements and ingress rules for a private subnet used by Recovery Service

Item	Requirements
General ingress rule 2: Allows SQLNet Traffic from Anywhere	This rule allows recovery catalog connections and real-time data protection from your OCI Database to Recovery Service. • **Stateless**: No (all rules must be stateful) • **Source Type**: CIDR • **Source CIDR**: CIDR of the VCN where the database resides • **IP Protocol**: TCP • **Source Port Range**: All • **Destination Port Range**: 2484

> **Note**
>
> If you use network security groups (NSG) to implement security rules or if your database VCN restricts network traffic between subnets, then ensure to add an egress rule for ports 2484 and 8005 from the database NSG or subnet to the Recovery Service NSG or subnet that you create.

Review Networking Service Permissions to Configure a Subnet

Ensure that you have these Networking Service permissions required to create a subnet in the database VCN and to assign security rules for Recovery Service.

Table 5-22 Networking Service Permissions Required to Create a Private subnet and Configure Security Rules for Recovery Service

Operation	Required IAM Policies
Configure a private subnet in a database VCN	• `use vcns` for the compartment which the VCN is in • `use subnets` for the compartment which the VCN 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

Alternatively, you can create a policy that allows a specified group with broader access to networking components.

For example, use this policy to allow a `NetworkAdmin` group to manage all networks in any compartment in a tenancy.

Example 5-8 Policy for Network Administrators

```
Allow group NetworkAdmin to manage virtual-network-family in tenancy
```

Create a Recovery Service Subnet in the Database VCN

Use the OCI Console to configure a private subnet for Recovery Service in your database virtual cloud network (VCN).

Copyright © 2022, 2026, Oracle and/or its affiliates.

1. From the navigation menu, select **Networking**, and then select **Virtual cloud networks**. The VCN list page opens. All VCNs in the selected compartment are displayed in a table.
2. Select the VCN in which your database resides.
3. Use these steps to create a Recovery Service subnet with a security list. If you choose to use network security groups, then skip this step.
   a. In the VCN details page, select the **Security** tab.
   b. Select the security list that is used for the VCN. You must add two ingress rules to allow destination ports 8005 and 2484.
   c. In the security list details page, select the **Security rules** tab.
   Select **Add Ingress Rules** and add the following details to set up a stateful ingress rule that allows HTTPS traffic from anywhere:
      - Source Type: CIDR
      - Source CIDR: Specify the CIDR of the VCN where the database resides.
      - IP Protocol: TCP
      - Source Port Range: All
      - Destination Port Range: 8005
      - Description: Specify an optional description of the ingress rule to help manage the security rules.
   d. Select **Add Ingress Rule** and add these details to set up a stateful ingress rule that allows SQLNet traffic from anywhere:
      - Source Type: CIDR
      - Source CIDR: Specify the CIDR of the VCN where the database resides.
      - IP Protocol: TCP
      - Source Port Range: All
      - Destination Port Range: 2484
      - Description: Specify an optional description of the ingress rule to help manage the security rules.

      > **Note**
      >
      > Select an IPv4-only subnet for Recovery Service in your database VCN. Do not select an IPv6-enabled subnet as Recovery Service does not support using an IPv6-enabled subnet.

   e. Go back to the VCN details page and select the **Subnets** tab.
   f. Select **Create Subnet**.
   g. Create a private subnet or select a private subnet that already exists in the database VCN. Oracle recommends a subnet size of /24 (256 IP addresses) for the private subnet.
   h. In the subnet details page, select the **Security** tab.
   i. Select **Add Security List** to add the security list that includes the ingress rules to allow destination ports 8005 and 2484.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> If your database VCN restricts network traffic between subnets, then ensure to add an egress rule for ports 2484 and 8005 from the database subnet to the Recovery Service subnet that you create.

4. Use these steps to create a Recovery Service subnet with network security groups (NSG).
   a. Go back to the VCN details page and select the **Subnets** tab.
   b. In the **Network Security Groups** section, select **Create Network Security Group**.
   c. Use one of these supported methods to configure security rules using NSGs:
      - To implement network isolation, create one NSG for the database VNIC (add egress rules to allow ports 2484 and 8005) and a separate NSG for Recovery Service (add ingress rules to allow ports 2484 and 8005).
      - Create and use a single NSG (with egress and ingress rules) for the database VNIC and Recovery Service.

      The Network Security Group page lists the NSGs that you create.
5. After you create the Recovery Service subnet in the database VCN, proceed to register the subnet as a Recovery Service subnet. Oracle recommends that you register a single Recovery Service subnet per VCN. If you have implemented security rules using NSGs, then you must also ensure to add the Recovery Service NSG to the Recovery Service subnet.

Register Recovery Service Subnet

After you have created a private subnet for Recovery Service in your database VCN, use this procedure to register the subnet in Recovery Service.

Multiple protected databases can use the same Recovery Service subnet. In order to ensure that the required number of IP addresses are available to support the Recovery Service private endpoints, you can assign multiple subnets to a Recovery Service subnet that is used by more than one protected database.

> **Note**
>
> Select an IPv4-only subnet for Recovery Service in your database VCN. Do not select an IPv6-enabled subnet as Recovery Service does not support using an IPv6-enabled subnet.

Use the following steps to register Recovery Service subnet:

1. Open the **navigation menu**, select **Oracle AI Database**, and then select **Database Backups**.
2. Select **Recovery Service Subnets**.
   The Recovery Service subnets list page opens. All Recovery Service subnets in the selected compartment are displayed in a table.
3. Select **Register Recovery Service subnet**.
4. In the **Register Recovery Service subnet** panel, enter the following details:
5. In the **Name** field, enter a name for the Recovery Service subnet.

Copyright © 2022, 2026, Oracle and/or its affiliates.

6. In the **Compartment** field, select the compartment where you want to create the Recovery Service subnet.
7. In the **Virtual cloud network** field, select the database VCN. Select **Compartment** to select a VCN belonging to a different compartment.
8. In the **Subnet** field, select a private subnet that you have configured for Recovery Service operations in your database VCN. Select **Compartment** to select a private subnet from a different compartment.
9. (Optional) Select **+Another Subnet** to assign an additional subnet to the Recovery Service subnet. If a single subnet does not contain enough IP addresses to support the Recovery Service private endpoints, then you can assign multiple subnets.
10. Expand **Advanced options** to enter these additional features:
    - Network security groups
    - Tags

    If you have used a network security group (NSG) to implement security rules for Recovery Service in the database VCN, then you must add the Recovery Service NSG to the Recovery Service subnet. The Recovery Service NSG can reside in the same compartment or in a different compartment. However, the NSG must belong to the same VCN to which the specified subnet belongs.

    a. In the **Network security groups** section, select **Use network security groups to control traffic**.
    b. Select the Recovery Service NSG you have created in the database VCN.
    c. Select **+Another network security group** to associate multiple NSGs (maximum five).

    (Optional) In the **Tag Namespace** field, consider adding a tag namespace, or tagging the control with an existing tag namespace.
11. Select **Register**.

You can replace a subnet or add more subnets to support the required number of private endpoints.

Use these steps to update an existing Recovery Service subnet:

1. In the **Recovery Service subnet** details page, select the **Subnets** tab.
2. Select **Add subnets** and select the subnets you want to add.
3. To replace an existing subnet, Select the **Action** menu, and select **Delete**. You can then add another subnet.

> **Note**
>
> A Recovery Service subnet must be associated with at least one subnet belonging to your database VCN.

Use these steps to manage the network security groups (NSGs) for an existing Recovery Service subnet:

1. In the **Network security groups** tab, select **Add network security groups**.
2. Select and add the Recovery Service network security groups (maximum five).
3. To remove an NSG, select the resource and select **Remove**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Ensure that the Recovery Service Subnet Can Communicate with Oracle Services

The Recovery Service Subnet that you registered needs to communicate with the Recovery Service.

To access the service, the routing table for the private subnet needs to include **All IAD Services In Oracle Services Network**.

Ensure that Your Database Has TDE Fully Configured

When using the Recovery Service, you must have your database fully TDE encrypted.

For new databases that are born in the cloud, this should already be done. But if you creating a stub database in OCI and migrating a database to an Oracle Database Cloud Service from on-premises, or somewhere else, you might not meet all the criteria. For these databases you should verify that you meet the prerequisites for a backing up to the recovery service. I have a blog post you can find here that will outline what to check for with queries to execute.

You must meet these 3 criteria:

- You need to have `WALLET_ROOT` configured in the database. If you are still using `sqlnet.ora`, you need to use `dbaascli` to properly set `WALLET_ROOT` for all databases that will be utilizing the Recovery Service. To enable `wallet_root` SPFILE parameter for an existing database, run:

  ```
  dbaascli tde enableWalletRoot
  ```

  > **Note**
  >
  > Setting `ENCRYPTION_WALLET_LOCATION` in `sqlnet.ora` is not supported and will be deprecated.

- You need to have an encryption key set for the CDB and all PDBs in your database.
- ALL tablespaces must be TDE encrypted before executing your first backup.

Turn Off Any Manual Operational Backups

In some cases, OCI users perform manual operational backups. These backups are run outside the standard tooling and support point-in-time recovery (non-KEEP backups).

If you are running any of these types of operational backups, it is critical that you turn them off at this point. Running operational backups to two different locations can cause issues with both backups and can create data loss scenarios. Therefore, before you enable automatic backups, you must disable manual backup scripts and processes to other storage destinations.

> **Note**
>
> If you are using the tooling for object storage backups, and switch to the Recovery Service, the switchover will be automated by the tooling, and all of the previous backups will remain available.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Back Up a Database Using the Console

This article explains how to manage backups in the Recovery Service and the Object Storage using the Console.

For more information about backups, see *Back Up and Recovery in Base Database Service*.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be granted security access in a **policy** by an administrator. This access is required whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you get a message that you don't have permission or are unauthorized, verify with your administrator what type of access you have and which **compartment** to work in.

If you're new to policies, see Getting Started with Policies and Common Policies.

Navigate to the List of Standalone Backups for Your Current Compartment

Perform the following steps to navigate to the list of standalone backups in your compartment using the OCI Console:

1. Open the **navigation menu**, select **Oracle Database**, and then select **Oracle Base Database Service**.

   The **DB Systems** list page opens. All DB systems in the selected compartment are displayed in a table.
2. On the DB system list page, under **Resources**, select **Standalone Backups** to view the list of standalone backups.

Configure Automatic Backups for a Database

When you start a DB system, you can optionally enable automatic backups for the initial database.

Procedure

Perform the following steps to configure or disable automatic backups after the database is created using the OCI Console:

> **Note**
>
> Switching the backup destination will trigger an immediate full backup by default. This full backup cannot be rescheduled.

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, on the **Database information** tab, the **Backup** section indicate whether automatic backups are enabled. When backups are enabled, the details also indicate the selected backup retention period.

Copyright © 2022, 2026, Oracle and/or its affiliates.

4. On the database details page, from the **Actions** menu, select **Configure automatic backups**.
5. In the **Configure automatic backups** panel, enter the following details:

Switch on or off the **Enable automatic backups** toggle as per your requirement. If you are enabling automatic backups, you can select to configure the **Recovery Service** or **Object Storage** as the backup destination.

> **Note**
>
> To use Recovery Service as a backup destination, the Oracle Database version must be 19.18 or later.

Your choice to use Recovery Service as the backup destination depends on the available limits in your tenancy and the available capacity in the specific region. The following restrictions apply to all regions when you enable automatic backups and want to use Recovery Service as the backup destination:

- If you have available limits and if there is available capacity in the region, then your choices are Recovery Service (default) and Object Storage.
- If you have exhausted the default limits for the Recovery Service, then you can only use Object Storage. However, you can make an additional limits request and then use Recovery Service.
- If there is no available capacity in the region, then you can use only Object Storage. However, after the required capacity becomes available in the region, you can seamlessly switch from Object Storage to Recovery Storage.
- Ampere A1 shape-based DB systems can only be backed up in the Object Storage.

Effective June 23, 2026, when you enable automatic backups using the OCI Console, the Autonomous Recovery Service will be the only available backup destination under the following conditions:

- The tenancy was created on or after June 23, 2026.
- The database is deployed in one of the following OCI regions: Netherlands Northwest (Amsterdam) **AMS**, Colombia Central (Bogota) **BOG**, Brazil East (Sao Paulo) **GRU**, Indonesia North (Batam) **HSG**, India South (Hyderabad) **HYD**, US East (Ashburn) **IAD**, South Korea Central (Seoul) **ICN**, Malaysia West 2 (Kulai) **JBP**, Saudi Arabia West (Jeddah) **JED**, Morocco West (Casablanca) **LEJ**, UK South (London) **LHR**, Italy Northwest (Milan) **LIN**, Spain Central (Madrid) **MAD**, Israel Central (Jerusalem) **MTZ**, Italy North (Turin) **NRQ**, Spain Central (Madrid 3) **ORF**, Chile Central (Santiago) **SCL**, Chile West (Valparaiso) **VAP**, Brazil Southeast (Vinhedo) **VCP**, Canada Southeast (Montreal) **YUL**, and Canada Southeast (Toronto) **YYZ**.
- The Oracle Database version is later than 19.18 or 23.4.

If any of these conditions are not met, OCI Object Storage will be shown as a backup destination.

Effective May 15, 2026, when you enable automatic backups using the OCI Console, Autonomous Recovery Service will be the only available backup destination under the following conditions:

- The tenancy was created on or after May 15, 2026.
- The database is deployed in one of the following OCI regions:

Copyright © 2022, 2026, Oracle and/or its affiliates.

Sweden Central (Stockholm) **ARN**, UAE Central (Abu Dhabi) **AUH**, India West (Mumbai) **BOM**, France Central (Paris) **CDG**, UK West (Newport) **CWL**, UAE East (Dubai) **DXB**, South Africa Central (Johannesburg) **JNB**, Japan Central (Osaka) **KIX**, Australia Southeast (Melbourne) **MEL**, France South (Marseille) **MRS**, Mexico Northeast (Monterrey) **MTY**, US Midwest (Chicago) **ORD**, Mexico Central (Queretaro) **QRO**, Saudi Arabia Central (Riyadh) **RUH**, US West (San Jose) **SJC**, Singapore (Singapore) **SIN**, Australia East (Sydney) **SYD**, Singapore West (Singapore) **XSP**, South Korea North (Chuncheon) **YNY**, or Switzerland North (Zurich) **ZRH**.

- The Oracle Database version is later than 19.18 or 23.4.

If any of these conditions are not met, OCI Object Storage will be shown as a backup destination.

Effective August 06, 2025, when you enable automatic backups through the OCI Console, Autonomous Recovery Service will be the only available backup destination under the following conditions:

- The tenancy was created on or after August 06, 2025.
- The database is deployed in one of the following OCI regions:
  Germany Central (Frankfurt) **FRA**, US West (Phoenix) **PHX**, or Japan East (Tokyo) **NRT**.
- The Oracle Database version is later than 19.18 or 23.4.

If any of these conditions are not met, OCI Object Storage will be shown as a backup destination.

Recovery Service

If **Recovery Service** is selected as the **Backup destination**, you can configure the following options:

- **Protection policy**: You can select from one of the preset protection policies or a custom policy. The system automatically deletes your backups at the end of your chosen protection policy recovery window.

  The following retention periods are available for **Recovery Service**. The retention periods (in days) are defined in the Recovery Service protection policy.

  - Bronze (14 days)
  - **Silver (35 days)** (default)
  - Gold (65 days)
  - Platinum (95 days)
  - Custom (User defined protection policy)
- **Real-time data protection**: Real-time protection is the continuous transfer of redo changes from a protected database to **Recovery Service**. This reduces data loss and provides a recovery point objective (RPO) near 0. This is an extra cost option.
- **Deletion options after database termination**: You can use the following options to retain managed database backups after the database is terminated. These options can also help restore the database from backups in case of accidental or malicious damage to the database.
  - **Retain backups according to the retention period**: When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be removed at the end of the specified retention period.
  - **Retain backups for 72 hours, then delete**: When a database is terminated, the automatic database backups associated with the terminated database and all of its

Copyright © 2022, 2026, Oracle and/or its affiliates.

resources will be retained for 72 hours and then deleted. The backups are retained for 72 hours to safeguard against accidental deletion by the user.

- **Scheduled day for initial backup**: Select a day of the week for the initial backup to begin.
- **Scheduled time for initial backup (UTC)**: Select a time for the initial backup to begin. The initial backup could start at any time or within the chosen two-hour scheduling window.
- **Scheduled time for daily backup (UTC)**: Select a time for the daily backup to begin. The daily backup could start at any time or within the chosen two-hour scheduling window.
- **Take the first backup immediately**: A full backup is an operating system backup of all data files and the control file that constitute an Oracle Database. A full backup must also include the parameter files associated with the database. You can take a database backup when the database is shut down or while the database is open. You must typically not take a backup after an instance failure or other unusual circumstances. If you select to defer the initial backup, your database may not be recoverable in the event of a database failure.

Object Storage

If **Object Storage** is selected as the **Backup destination**, you can configure the following options:

- **Backup retention period**: If you select to enable automatic backups, you can select a policy with one of the preset retention periods. The system automatically deletes your incremental backups at the end of your chosen retention period. You can change the backup retention period after provisioning.

  The following retention periods are available for **Object Storage**.

  - 7 days
  - 15 days
  - **30 days** (default)
  - 45 days
  - 60 days
- **Deletion options after database termination**: You can use the following options to retain managed database backups after the database is terminated. These options can also help restore the database from backups in case of accidental or malicious damage to the database.
  - **Retain backups according to the retention period**: When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be removed at the end of the specified retention period.
  - **Retain backups for 72 hours, then delete**: When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be retained for 72 hours and then deleted. The backups are retained for 72 hours to safeguard against accidental deletion by the user.
- **Scheduled day for full backup**: Select a day of the week for the initial and future full backups to begin.
- **Scheduled time for full backup (UTC)**: Select a time for the full backup to begin. The full backup could start at any time or within the chosen two-hour scheduling window.
- **Scheduled time for incremental backup (UTC)**: Select a time for the incremental backup to begin. The incremental backup could start at any time or within the chosen two-hour scheduling window.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Take the first backup immediately**: A full backup is an operating system backup of all data files and the control file that constitute an Oracle Database. A full backup must also include the parameter files associated with the database. You can take a database backup when the database is shut down or while the database is open. You must typically not take a backup after an instance failure or other unusual circumstances. If you select to defer the initial backup, your database may not be recoverable in the event of a database failure.

Select **Save**.

The database details page displays the automatic backup status and details in the **Backup** section.

Configure Automatic Backups for a Standby Database

Procedure

Perform the following steps to configure automatic backups for a standby database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, select the **Data Guard associations** tab.
4. In the list of standby databases, select the database for which you want to configure automatic backups.
5. On the standby database details page, on the **Database information** tab, the **Backup** section indicate whether automatic backups are enabled. When backups are enabled, the details also indicate the selected backup retention period.
6. On the standby database details page, from the **Actions** menu, select **Configure automatic backups**.
7. In the **Configure automatic backups** panel, enter the following details:

> **Note**
>
> If the backup was already configured on the primary database, then the standby must use the same backup destination.

Switch on or off the **Enable automatic backups** toggle as per your requirement. If you are enabling automatic backups, you can select to configure the **Recovery Service** or **Object Storage** as the backup destination.

> **Note**
>
> To use Recovery Service as a backup destination, the Oracle Database version must be 19.18 or later.

Your choice to use Recovery Service as the backup destination depends on the available limits in your tenancy and the available capacity in the specific region. The following restrictions

Copyright © 2022, 2026, Oracle and/or its affiliates.

apply to all regions when you enable automatic backups and want to use Recovery Service as the backup destination:

- If you have available limits and if there is available capacity in the region, then your choices are Recovery Service (default) and Object Storage.
- If you have exhausted the default limits for the Recovery Service, then you can only use Object Storage. However, you can make an additional limits request and then use Recovery Service.
- If there is no available capacity in the region, then you can use only Object Storage. However, after the required capacity becomes available in the region, you can seamlessly switch from Object Storage to Recovery Storage.
- Ampere A1 shape-based DB systems can only be backed up in the Object Storage.

Effective June 23, 2026, when you enable automatic backups using the OCI Console, the Autonomous Recovery Service will be the only available backup destination under the following conditions:

- The tenancy was created on or after June 23, 2026.
- The database is deployed in one of the following OCI regions:
  Netherlands Northwest (Amsterdam) **AMS**, Colombia Central (Bogota) **BOG**, Brazil East (Sao Paulo) **GRU**, Indonesia North (Batam) **HSG**, India South (Hyderabad) **HYD**, US East (Ashburn) **IAD**, South Korea Central (Seoul) **ICN**, Malaysia West 2 (Kulai) **JBP**, Saudi Arabia West (Jeddah) **JED**, Morocco West (Casablanca) **LEJ**, UK South (London) **LHR**, Italy Northwest (Milan) **LIN**, Spain Central (Madrid) **MAD**, Israel Central (Jerusalem) **MTZ**, Italy North (Turin) **NRQ**, Spain Central (Madrid 3) **ORF**, Chile Central (Santiago) **SCL**, Chile West (Valparaiso) **VAP**, Brazil Southeast (Vinhedo) **VCP**, Canada Southeast (Montreal) **YUL**, and Canada Southeast (Toronto) **YYZ**.
- The Oracle Database version is later than 19.18 or 23.4.

If any of these conditions are not met, OCI Object Storage will be shown as a backup destination.

Effective May 15, 2026, when you enable automatic backups using the OCI Console, Autonomous Recovery Service will be the only available backup destination under the following conditions:

- The tenancy was created on or after May 15, 2026.
- The database is deployed in one of the following OCI regions:
  Sweden Central (Stockholm) **ARN**, UAE Central (Abu Dhabi) **AUH**, India West (Mumbai) **BOM**, France Central (Paris) **CDG**, UK West (Newport) **CWL**, UAE East (Dubai) **DXB**, South Africa Central (Johannesburg) **JNB**, Japan Central (Osaka) **KIX**, Australia Southeast (Melbourne) **MEL**, France South (Marseille) **MRS**, Mexico Northeast (Monterrey) **MTY**, US Midwest (Chicago) **ORD**, Mexico Central (Queretaro) **QRO**, Saudi Arabia Central (Riyadh) **RUH**, US West (San Jose) **SJC**, Singapore (Singapore) **SIN**, Australia East (Sydney) **SYD**, Singapore West (Singapore) **XSP**, South Korea North (Chuncheon) **YNY**, or Switzerland North (Zurich) **ZRH**.
- The Oracle Database version is later than 19.18 or 23.4.

If any of these conditions are not met, OCI Object Storage will be shown as a backup destination.

Effective August 06, 2025, when you enable automatic backups through the OCI Console, Autonomous Recovery Service will be the only available backup destination under the following conditions:

- The tenancy was created on or after August 06, 2025.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- The database is deployed in one of the following OCI regions:
  Germany Central (Frankfurt) **FRA**, US West (Phoenix) **PHX**, or Japan East (Tokyo) **NRT**.
- The Oracle Database version is later than 19.18 or 23.4.

If any of these conditions are not met, OCI Object Storage will be shown as a backup destination.

Recovery Service

If **Recovery Service** is selected as the **Backup destination**, you can configure the following options:

- **Protection policy**: You can select from one of the preset protection policies or a custom policy. The system automatically deletes your backups at the end of your chosen protection policy recovery window.

  The following retention periods are available for **Recovery Service**. The retention periods (in days) are defined in the Recovery Service protection policy.

  - Bronze (14 days)
  - **Silver (35 days)** (default)
  - Gold (65 days)
  - Platinum (95 days)
  - Custom (User defined protection policy)
- **Real-time data protection**: Real-time protection is the continuous transfer of redo changes from a protected database to **Recovery Service**. This reduces data loss and provides a recovery point objective (RPO) near 0. This is an extra cost option.
- **Deletion options after database termination**: You can use the following options to retain managed database backups after the database is terminated. These options can also help restore the database from backups in case of accidental or malicious damage to the database.
  - **Retain backups according to the retention period**: When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be removed at the end of the specified retention period.
  - **Retain backups for 72 hours, then delete**: When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be retained for 72 hours and then deleted. The backups are retained for 72 hours to safeguard against accidental deletion by the user.
- **Scheduled day for initial backup**: Select a day of the week for the initial backup to begin.
- **Scheduled time for initial backup (UTC)**: Select a time for the initial backup to begin. The initial backup could start at any time or within the chosen two-hour scheduling window.
- **Scheduled time for daily backup (UTC)**: Select a time for the daily backup to begin. The daily backup could start at any time or within the chosen two-hour scheduling window.
- **Take the first backup immediately**: A full backup is an operating system backup of all data files and the control file that constitute an Oracle Database. A full backup must also include the parameter files associated with the database. You can take a database backup when the database is shut down or while the database is open. You must typically not take a backup after an instance failure or other unusual circumstances. If you select to defer the initial backup, your database may not be recoverable in the event of a database failure.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Object Storage

If **Object Storage** is selected as the **Backup destination**, you can configure the following options:

- **Backup retention period**: If you select to enable automatic backups, you can select a policy with one of the preset retention periods. The system automatically deletes your incremental backups at the end of your chosen retention period. You can change the backup retention period after provisioning.

  The following retention periods are available for **Object Storage**.

  - 7 days
  - 15 days
  - **30 days** (default)
  - 45 days
  - 60 days
- **Deletion options after database termination**: You can use the following options to retain managed database backups after the database is terminated. These options can also help restore the database from backups in case of accidental or malicious damage to the database.
  - **Retain backups according to the retention period**: When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be removed at the end of the specified retention period.
  - **Retain backups for 72 hours, then delete**: When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be retained for 72 hours and then deleted. The backups are retained for 72 hours to safeguard against accidental deletion by the user.
- **Scheduled day for full backup**: Select a day of the week for the initial and future full backups to begin.
- **Scheduled time for full backup (UTC)**: Select a time for the full backup to begin. The full backup could start at any time or within the chosen two-hour scheduling window.
- **Scheduled time for incremental backup (UTC)**: Select a time for the incremental backup to begin. The incremental backup could start at any time or within the chosen two-hour scheduling window.
- **Take the first backup immediately**: A full backup is an operating system backup of all data files and the control file that constitute an Oracle Database. A full backup must also include the parameter files associated with the database. You can take a database backup when the database is shut down or while the database is open. You must typically not take a backup after an instance failure or other unusual circumstances. If you select to defer the initial backup, your database may not be recoverable in the event of a database failure.

Select **Save**.

The database details page displays the automatic backup status and details in the **Backup** section.

Create an On-Demand Backup of a Database

Perform the following steps to create an on-demand backup of a database using the OCI Console:

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> Object Storage creates a full backup, while the Recovery Service creates an incremental backup of the database.

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, select the **Backups** tab to view the list of backups.
4. Select **Create backup**.
5. In the **Create Backup** panel, enter the following details:
   - Provide a name for the backup database.
   - Select **Backup retention** from the following options:
     - **Retain backups per backup retention period**: Select this option to use the protection policy retention period for this backup.
     - **Specify long-term backup retention period**: Select this option to specify an LTR period with Recovery Service. The retention period must be in days (90 - 3650) or years (1 - 10) from when the backup was created.
6. Select **Create**.

Change the Retention Period of a Long-Term Retention Backup

Procedure

Perform the following steps to change the retention period of an LTR backup with Recovery Service using the OCI Console:

1. Open the **navigation menu**, select **Oracle Database**, and then select **Oracle Base Database Service**.

   The **DB Systems** list page opens. All DB systems in the selected compartment are displayed in a table.
2. If you have a standalone backup, under **Resources**, select **Standalone backups**, and skip to step **7** below.
3. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
4. On the **Databases** tab, select the database that you want to work with.
5. On the database details page, select the **Backups** tab to view the list of backups.
6. In the list of backups, select the **Actions** menu for the backup with type *Long-term backup* for which you want to change the retention period.
7. Select **Change retention period**.
8. Enter the **Backup retention period**. The retention period must be in days (90 - 3650) or years (1 - 10) from when the backup was created.
9. Select **Save**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

View Details of a Protected Database

Procedure

Perform the following steps to view the details of the protected database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, on the **Database information** tab, the **Backup** section indicate whether automatic backups are enabled. When backups are enabled, the details also indicate the chosen backup details.
4. Select the Recovery Service link in the **Backup destination** to view the details of the protected database.

For more information, see Managing Protected Databases.

View Status of a Backup

Procedure

Perform the following steps to view the backup status of a database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, select the **Backups** tab to view the list of backups.
4. The **State** column displays the status of the backup.
   - The following are the various states of the backup: **Active**, **Creating**, **Canceled**, **Canceling**, or **Failed**.

Cancel a Backup

Procedure

Perform the following steps to cancel a backup of a database using the OCI Console:

> **Note**
>
> Only backups in the **Creating** state can be canceled.

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, select the **Backups** tab to view the list of backups.

Copyright © 2022, 2026, Oracle and/or its affiliates.

4. The **State** column displays the status of the backup. The following are the various states of the backup: **Active**, **Creating**, **Canceled**, **Canceling**, or **Failed**.
5. Select the **Actions** menu for the backup you want to cancel.
6. Select **Cancel backup** and confirm.
7. The status of the backup changes to **Canceling**.

If the cancel backup fails, in the **Work requests** pane under **Resources,** you will see a line item called **Cancel database backup** in **Failed** state. There will also be a work request for the backup initiated by **Create database backup** that will reflect the state of the backup operation.

Delete Full Backups of a Database

Procedure

Perform the following steps to delete the full backup of a database using the OCI Console:

> **Note**
>
> You cannot explicitly delete automatic backups. Unless you terminate the database, automatic backups remain in the Recovery Service and Object Storage for the number of days specified by the user, after which time they are automatically deleted.

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, select the **Backups** tab to view the list of backups.
4. Select the **Actions** menu for the backup you want to delete.
5. Select **Delete** and confirm.

Delete Standalone Backups of a Database

Perform the following steps to delete a standalone backup of a database using the OCI Console:

1. Open the **navigation menu**, select **Oracle Database**, and then select **Oracle Base Database Service**.

   The **DB Systems** list page opens. All DB systems in the selected compartment are displayed in a table.
2. On the **DB Systems** list page, under **Resources**, select **Standalone Backups**.
   A list of standalone backups is displayed.
3. On the standalone backups list page, select the **Actions** menu for the backup you want to delete.
4. Select **Delete** and confirm.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Back Up a Database to Object Storage Using RMAN

This article explains how to use Recovery Manager (RMAN) to manage backups of your DB system database to your own Object Storage.

To back up to the service you'll need to create an Object Storage bucket for the backups, generate a password for the service, install the Oracle Database Cloud Backup Module, and then configure RMAN to send backups to the service. The backup module is a system backup to tape (SBT) interface that's tightly integrated with RMAN, so you can use familiar RMAN commands to perform backup and recovery operations.

You'll notice *Swift* mentioned in the Console and in the endpoint URL for the service. That's because the backup module is typically used to back up to the Oracle Database Backup Cloud Service, which is an OpenStack Swift object store.

> **Tip**
>
> On a DB system, you can use the DBCLI to back up to Object Storage. This is an alternative to installing the backup module and using RMAN for backups. For more information, see Objectstoreswift Commands.

Prerequisites

You'll need the following:

- A DB system and a database to back up.
- The DB system's cloud network (VCN) must be configured with access to Object Storage:
  - For Object Storage access in the same region as the DB system: Oracle recommends using a service gateway.
  - For Object Storage access in a different region than the DB system: Use an internet gateway. Note that the network traffic between the DB system and Object Storage does not leave the cloud and never reaches the public internet.

  For more information, see *VCN and Subnets*.
- An existing Object Storage bucket to use as the backup destination. You can use the Console or the Object Storage API to create the bucket.
  For more information, see Managing Buckets.
- An auth token generated by OCI. You can use the Console or the IAM API to generate the password.
  For more information, see *working with auth tokens* in Managing User Credentials.
- The user name (specified when you install and use the backup module) must have tenancy-level access to Object Storage. An easy way to do this is to add the user name to the Administrators group. However, that allows access to *all* of the cloud services. Instead, an administrator should create a policy like the following that limits access to only the required resources in Object Storage for backing up and restoring the database:

  ```
  Allow group <group_name> to manage objects in compartment
  <compartment_name> where target.bucket.name = '<bucket_name>'

  Allow group <group_name> to read buckets in compartment <compartment_name>
  ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

For more information about adding a user to a group, see Managing Groups. For more information about policies, see Getting Started with Policies.

Install the Backup Module On the DB System

1. SSH to the DB system.

   ```
   ssh -i <SSH_key_used_when_launching_the_DB_system>
   opc@<DB_system_IP_address_or_hostname>
   ```

2. Log in as `opc` user.

   ```
   login as: opc
   ```

3. `sudo` to the `oracle` user.

   ```
   sudo su - oracle
   ```

4. Change to the directory that contains the backup module `opc_install.jar` file.

   ```
   cd /opt/oracle/oak/pkgrepos/oss/odbcs
   ```

5. Use the following command syntax to install the backup module.

   ```
   java -jar opc_install.jar -opcId <user_id> -opcPass '<auth_token>' -
   container <bucket_name>;-walletDir ~/hsbtwallet/ -libDir ~/lib/ -
   configfile ~/config -host https://
   swiftobjectstorage.<region_name>.oraclecloud.com/v1/
   <object_storage_namespace>
   ```

   The parameters are:

Parameter	Description
`-opcId`	The user name for the Oracle Cloud Infrastructure user account, for example: `-opcId <username>@<example>.com`. This is the user name you use to sign in to the Console. The user name must be a member of the Administrators group, as described in the *Prerequisites* section. You can also specify the user name in single quotes. This might be necessary if the name contains special characters, for example: `-opcId 'j~smith@<example>.com'` Make sure to use straight single quotes and not slanted apostrophes.
`-opcPass`	The auth token generated by using the Console or IAM API, in single quotes, for example: `-opcPass <password>` Make sure to use straight single quotes and not slanted apostrophes. For more information, see Managing User Credentials. This is **not** the password for the Oracle Cloud Infrastructure user.
`-container`	The name of an existing bucket in Object Storage to use as the backup destination, for example: `-container DBBackups`

Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameter	Description
`-walletDir`	The directory where the install tool will create an Oracle Wallet containing the Oracle Cloud Infrastructure user name and auth token. `-walletDir ~/hsbtwallet` creates the wallet in the current user (oracle) home directory.
`-libDir`	The directory where the SBT library is stored. The directory must already exist before you run the command. This parameter causes the latest SBT library to be downloaded. `-libDir ~/lib/` downloads the `libopc.so` file to the current user's home directory, for example, `/home/oracle/lib/libopc.so`.
`-configfile`	The name of the initialization parameter file that will be created by the install tool. This file will be referenced by your RMAN jobs. `-configfile ~/config` creates the file in the current user's home directory, for example, `/home/oracle/config`.
`-host`	The endpoint URL to which backups are to be sent: `https://swiftobjectstorage.<region_name>.oraclecloud.com/v1/<object_storage_namespace>` where `object_storage_namespace` is your tenancy's Object Storage namespace. For more information, see Understanding Object Storage Namespaces. Do not add a slash after the Object Storage namespace. To look up the region name, see Regions and Availability Domains.

Configure RMAN

This topic describes how to configure RMAN to use the bucket as the default backup destination. The following assumes you are still logged in to the DB system.

1. On the DB system, set the `ORACLE_HOME` and `ORACLE_SID` environment variables using the `oraenv` utility.

   ```
   . oraenv
   ```

2. Connect to the database using RMAN.

   ```
   rman target /
   ```

3. Configure RMAN to use the SBT device and point to the `config` file that was created when you installed the backup module. A sample command for a version 12 database is shown here.

   ```
   CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' PARMS
       'SBT_LIBRARY=/home/oracle/lib/libopc.so,
       SBT_PARMS=(OPC_PFILE=/home/oracle/config)';
   ```

4. Configure RMAN to use SBT_TAPE by default. The following sample enables the controlfile and spfile autobackup to SBT_TAPE and configures encryption. There are other settings that may apply to your installation such as compression, number of backup and recovery channels to use, backup retention policy, archived log deletion policy, and more.

Copyright © 2022, 2026, Oracle and/or its affiliates.

See the Oracle Backup and Recovery documentation for your version of Oracle for more information on choosing the appropriate settings.

```
CONFIGURE DEFAULT DEVICE TYPE TO SBT_TAPE;
CONFIGURE BACKUP OPTIMIZATION ON;
CONFIGURE CONTROLFILE AUTOBACKUP ON;
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE SBT_TAPE TO '%F';
CONFIGURE ENCRYPTION FOR DATABASE ON;
```

> **Note**
>
> Backups must be encrypted. You will specify encryption when you perform a backup. You will get an error if a backup is not encrypted.

Once the RMAN configuration is complete, you can use the same RMAN commands that you regularly use for tape backups.

Back Up the Database

This topic provides examples of commonly used backup commands.

1. Set the database encryption:

   ```
   SET ENCRYPTION IDENTIFIED BY "password" ONLY;
   ```

   Note that this setting is not permanent; you must set it for each new RMAN session.

2. Back up the database and archivelogs. Below are some example commands. See the Oracle Backup and Recovery documentation for your version of Oracle for more information about choosing a back up procedure that meets your needs. Be sure to back up regularly to minimize potential data loss and always include a copy of the spfile and controlfile. Note that the example below uses multi-section incremental backups, which is a feature introduced in 12c. When using 11g, omit the `section size` clause.

   ```
   BACKUP INCREMENTAL LEVEL 0 SECTION SIZE 512M DATABASE PLUS ARCHIVELOG;


   BACKUP INCREMENTAL LEVEL 1 SECTION SIZE 512M DATABASE PLUS ARCHIVELOG;


   BACKUP INCREMENTAL LEVEL 1 CUMULATIVE SECTION SIZE 512M DATABASE PLUS
   ARCHIVELOG;
   ```

3. Backup archivelogs frequently to minimize potential data loss, and keep multiple backup copies as a precaution.

   ```
   BACKUP ARCHIVELOG ALL NOT BACKED UP 2 TIMES;
   ```

When the backup job completes, you can display the backup files in your bucket in the Console on the **Storage** page, by selecting **Object Storage**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Recover a Database Using the Console

This article explains how to recover a database from an automatic backup stored in the Recovery Service or Object Storage.

For more information about backups, see *Back Up and Recovery in Base Database Service*.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be granted security access in a **policy** by an administrator. This access is required whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you get a message that you don't have permission or are unauthorized, verify with your administrator what type of access you have and which **compartment** to work in.

If you're new to policies, see Getting Started with Policies and Common Policies.

Prerequisites

The DB system requires access to the Recovery Service or the Object Storage service, including connectivity to the applicable Swift endpoint for Object Storage. Oracle recommends using a service gateway with the VCN to enable this access. The Recovery Service requires a dedicated network path in each database VCN.

For more information about setting up your VCN for the DB system, including the service gateway, see *VCN and Subnets*.

For more information about the Swift endpoints to use, see Can I use Oracle Cloud Infrastructure Object Storage as a destination for my on-premises backups?.

General Information

- The default password profile will be reset when you restore a DB system from a backup.
- You can use the Console to restore the database from an automatic backup that was created by using the Console or the API. You can restore to the last known good state of the database, or you can specify a point in time or an existing System Change Number (SCN). You can also create a new database by using a standalone backup.
- The list of backups you see in the Console does not include any unmanaged backups (backups created directly by using `RMAN` or `dbcli`).
- If the database you are restoring was configured to use customer-managed encryption keys after the specified timestamp or SCN, the database will be restored without customer-managed keys enabled. You can change the encryption settings after restoring to use customer-managed keys.

For more information, see *Manage Encryption Keys*.

Procedure

Perform the following steps to restore a database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, from the **Actions** menu, select **Restore**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> You can also access the list of backups by selecting **Standalone backups** under **Resources**.

4. Select one of the following restore options:
   - **Restore to the latest**: Restores the database to the last known good state with the least possible data loss.
   - **Restore to a timestamp**: Restores the database to the timestamp specified.
   - **Restore to SCN**: Restores the database using the System Change Number (SCN) specified. This SCN must be valid.

   > **Note**
   >
   > You can determine the SCN number to use either by accessing and querying your database host or by accessing any online or archived logs.

5. Select **Restore** and confirm.

> **Note**
>
> If the restore operation fails, the database will be in a "Restore failed" state. You can try restoring again using a different restore option. However, Oracle recommends that you review the `RMAN` logs on the host and fix any issues before reattempting to restore the database.

Recover a Database from Object Storage Using RMAN Backup

This article explains how to recover a Recovery Manager (RMAN) backup stored in Object Storage.

Prerequisites

You'll need the following:

- A new DB system to restore the database to (see assumptions below). For more information, see *Overview of Creating a DB System*.
- The Oracle Database Cloud Backup Module must be installed on the DB system. For more information, see *Installing the Backup Module on the DB System* in *Back Up a Database to Object Storage Using RMAN*.

Assumptions

The procedures below assume the following:

- A new DB system has been created to host the restored database and no other database exists on the new DB system. It is possible to restore to a DB system that has existing databases, but that is beyond the scope of this topic.
- The original database is lost and all that remains is the latest RMAN backup. The procedure assumes the DB system (inclusive of the database) no longer exists.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> Any data not included in the most recent backup will be lost.

- The Oracle Wallet and/or encryption keys used by the original database at the time of the last backup is available.
- The RMAN backup contains a copy of the control file and spfile as of the most recent backup as well as all of the datafile and archivelog backups needed to perform a complete database recovery.
- An RMAN catalog will not be used during the restore.

Set Up Storage on the DB system

1. SSH to the DB System.

   ```
   ssh -i <private_key_path> opc@<db_system_ip_address>
   ```

2. Log in as opc and then sudo to the root user. Use `sudo su -` with a hyphen to invoke the root user's profile, which will set the PATH to the dbcli directory (`/opt/oracle/dcs/bin`).

   ```
   login as: opc
   sudo su -
   ```

3. You can use an existing empty database home or create a new one for the restore. Use the applicable commands to help you complete this step.
   If you will be using an existing database home:

   - Use the Dbhome Commands to list the database homes.

     ```
     dbcli list-dbhomes
     ```

     Output:

     ```
     ID                                       Name                 DB
     Version Home Location
     ---------------------------------------- --------------------
     ---------- ----------------------------------------------
     2e743050-b41d-4283-988f-f33d7b082bda     OraDB12102_home1
     12.1.0.2   /u01/app/oracle/product/12.1.0.2/dbhome_1
     ```

   - Use the Database Commands to ensure the database home is not associated with any database.

   If necessary, use the Dbhome Commands to create a database home for the restore.

4. Use the Dbstorage Commands to set up directories for DATA, RECO, and REDO storage. The following example creates 10GB of ACFS storage for the rectest database.

   ```
   dbcli create-dbstorage --dbname rectest --dataSize 10 --dbstorage ACFS
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> When restoring a version 11.2 database, ACFS storage must be specified.

Perform the Database Restore and Recovery

1. SSH to the DB system, log in as opc, and then become the oracle user.

   ```
   sudo su - oracle
   ```

2. Create an entry in `/etc/oratab` for the database. Use the same SID as the original database.

   ```
   db1:/u01/app/oracle/product/12.1.0.2/dbhome_1:N
   ```

3. Set the `ORACLE_HOME` and `ORACLE_SID` environment variables using the oraenv utility.

   ```
   . oraenv
   ```

4. Obtain the DBID of the original database. This can be obtained from the file name of the `controlfile` autobackup on the backup media. The file name will include a string that contains the DBID. The typical format of the string is `c-DDDDDDDDDDDD-YYYYMMDD-NN` where `DDDDDDDDDDDD` is the DBID, `YYYYMMDD` is the date the backup was created, and `NN` is a sequence number to make the file name unique. The DBID in the following examples is 1508405000. Your DBID will be different.
   Use the following curl syntax to perform a general query of Object Storage. The parameters in red are the same parameters you specified when installing the backup module as described in *Installing the Backup Module on the DB System* in *Back Up a Database to Object Storage Using RMAN*.

   ```
   curl -u '<user_ID>.com:<auth_token>' -v https://
   swiftobjectstorage.<region_name>.oraclecloud.com/v1/
   <object_storage_namespace>
   ```

   To look up the region name, see Regions and Availability Domains.

   For example:

   ```
   curl -u 'djones@mycompany.com:1cnk!d0++ptETd&C;tHR' -v https://
   swiftobjectstorage.<region_name>.oraclecloud.com/v1/
   myobjectstoragenamespace
   ```

   To get the DBID from the control file name, use the following syntax:

   ```
   curl -u '<user_id>.com:<auth_token>' -v https://
   swiftobjectstorage.<region_name>.oraclecloud.com/v1/
   <object_storage_namespace>/<bucket_name>?prefix=sbt_catalog/c-
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

For example:

```
curl -u 'djones@mycompany.com:1cnk!d0++ptETd&C;tHR' -v https://
swiftobjectstorage.<region_name>.oraclecloud.com/v1/
myobjectstoragenamespace/dbbackups/?prefix=sbt_catalog/c-
```

In the sample output below, 1508405000 is the DBID.

```
{
    "bytes": 1732,
    "content_type": "binary/octet-stream",
    "hash": "f1b61f08892734ed7af4f1ddaabae317",
    "last_modified": "2016-08-11T20:28:34.438000",
    "name": "sbt_catalog/c-1508405000-20160811-00/metadata.xml"
}
```

5. Run RMAN and connect to the target database. There is no need to create a `pfile` or `spfile` or use a backup `controlfile`. These will be restored in the following steps. Note that the target database is `(not started)`. This is normal and expected at this point.

```
rman target /
```

Output:

```
Recovery Manager: Release 12.1.0.2.0 - Production on Wed Jun 22 18:36:40
2016
Copyright (c) 1982, 2014, Oracle and/or its affiliates.  All rights
reserved.
connected to target database (not started)
```

6. Set the DBID using the value obtained above.

```
set dbid 1508405000;
```

7. Run the following command. If the server parameter file is not available, RMAN attempts to start the instance with a dummy server parameter file. The ORA-01078 and LRM-00109 errors are normal and can be ignored.

```
STARTUP NOMOUNT


startup failed: ORA-01078: failure in processing system parameters
LRM-00109: could not open parameter file '/u01/app/oracle/product/12.1.0.2/
dbhome_1/dbs/initdb1.ora'

starting Oracle instance without parameter file for retrieval of spfile
Oracle instance started

Total System Global Area    2147483648 bytes

Fixed Size                     2944952 bytes
Variable Size                847249480 bytes
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
Database Buffers           1254096896 bytes
Redo Buffers                 43192320 bytes
```

8. Restore the server parameter file from autobackup.
   The SBT_LIBRARY is the same library specified with the `-libDir` parameter when the Backup Module was installed, for example `/home/oracle/lib/`.

   The OPC_PFILE is the same file specified with the `-configfile` parameter when the Backup Module was installed, for example `/home/oracle/config`.

   ```
   set controlfile autobackup format for device type sbt to '%F';
   run {
     allocate channel c1 device type sbt PARMS  'SBT_LIBRARY=/home/oracle/lib/
   libopc.so, SBT_PARMS=(OPC_PFILE=/home/oracle/config)';
     restore spfile from autobackup;
   }
   ```

9. Create the directory for audit_file_dest. The default is `/u01/app/oracle/admin/$ORACLE_SID/adump`. You can see the setting used by the original database by searching the spfile for the string, `audit_file_dest`.

   ```
   strings ${ORACLE_HOME}/dbs/spfile${ORACLE_SID}.ora | grep audit_file_dest
   *.audit_file_dest='/u01/app/oracle/admin/db1/adump'

   mkdir -p /u01/app/oracle/admin/db1/adump
   ```

10. If block change tracking was enabled on the original database, create the directory for the block change tracking file. This will be a directory under `db_create_file_dest`. Search the `spfile` for the name of the directory.

    ```
    strings ${ORACLE_HOME}/dbs/spfile${ORACLE_SID}.ora | grep
    db_create_file_dest
    *.db_create_file_dest='/u02/app/oracle/oradata/db1'

    mkdir -p /u02/app/oracle/oradata/db1/<$ORA_UNQNAME if available or
    database name>/changetracking
    ```

11. Restart the instance with the restored server parameter file.

    ```
    STARTUP FORCE NOMOUNT;
    ```

12. Restore the controlfile from the RMAN autobackup and mount the database.

    ```
    set controlfile autobackup format for device type sbt to '%F';
    run {
      allocate channel c1 device type sbt PARMS 'SBT_LIBRARY=/home/oracle/lib/
    libopc.so, SBT_PARMS=(OPC_PFILE=/home/oracle/config)';
      restore controlfile from autobackup;
      alter database mount;
    }
    ```

13. Restore and recover the database.

    ```
    RESTORE DATABASE;
    RECOVER DATABASE;
    ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

14. RMAN will recover using archived redo logs until it can't find any more. It is normal for an error similar to the one below to occur when RMAN has applied the last archived redo log in the backup and can't find any more logs.

```
unable to find archived log
archived log thread=1 sequence=29
RMAN-00571: ===========================================================
RMAN-00569: =============== ERROR MESSAGE STACK FOLLOWS ===============
RMAN-00571: ===========================================================
RMAN-03002: failure of recover command at 06/28/2016 00:57:35
RMAN-06054: media recovery requesting unknown archived log for thread 1
with sequence 29 and starting SCN of 2349563
```

15. Open the database with resetlogs.

```
ALTER DATABASE OPEN RESETLOGS;
```

The recovery is complete. The database will have all of the committed transactions as of the last backed up archived redo log.

Recover a Database from the OCI Classic Object Store

This article explains how to recover a database using a backup created by the Oracle Database Backup Module and stored in Oracle Cloud Infrastructure Object Storage Classic.

The following terms are used throughout this topic:

- Source database: The database backup in Object Storage Classic.
- Target database: The new database on a DB system in Oracle Cloud Infrastructure.

Prerequisites

You'll need the following:

- The service name, identity name, container, user name, and password for Oracle Cloud Infrastructure Object Storage Classic.
- The backup password if password-based encryption was used when backing up to Object Storage Classic.
- The source database ID, database name, database unique name (required for setting up storage).
- If the source database is configured with Transparent Data Encryption (TDE), you'll need a backup of the wallet and the wallet password.
- Tnsnames to setup for any database links.
- The output of `Opatch lsinventory` for the source database Oracle_home, for reference.
- A copy of the sqlpatch directory from the source database home. This is required for rollback in case the target database does not include these patches.

Set Up Storage on the DB System

1. SSH to the DB System.

```
ssh -i <private_key_path> opc@<db_system_ip_address>
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

2. Log in as opc and then sudo to the root user. Use `sudo su -` with a hyphen to invoke the root user's profile, which will set the PATH to the dbcli directory (`/opt/oracle/dcs/bin`).

```
login as: opc
sudo su -
```

3. Use the Dbstorage Commands to set up directories for DATA, RECO, and REDO storage. The following example creates 10GB of ACFS storage for the tdetest database.

```
dbcli create-dbstorage --dbname tdetest --dataSize 10 --dbstorage ACFS
```

> **Note**
>
> When migrating a version 11.2 database, ACFS storage must be specified.

4. Use the Dbstorage Commands to list the storage ID. You'll need the ID for the next step.

```
dbcli list-dbstorages
```

Output:

```
ID                                       Type   DBUnique Name        Status
---------------------------------------- ------ --------------------
----------
9dcdfb8e-e589-4d5f-861a-e5ba981616ed     Acfs   tdetest
Configured
```

5. Use the Dbstorage Commands with the storage ID from the previous step to list the DATA, RECO and REDO locations.

```
dbcli describe-dbstorage --id 9dcdfb8e-e589-4d5f-861a-e5ba981616ed
```

Output:

```
DBStorage details
----------------------------------------------------------------
                     ID: 9dcdfb8e-e589-4d5f-861a-e5ba981616ed
                DB Name: tdetest
         DBUnique Name: tdetest
        DB Resource ID:
           Storage Type: Acfs
          DATA Location: /u02/app/oracle/oradata/tdetest
          RECO Location: /u03/app/oracle/fast_recovery_area/
          REDO Location: /u03/app/oracle/redo/
                  State: ResourceState(status=Configured)
                Created: August 24, 2016 5:25:38 PM UTC
            UpdatedTime: August 24, 2016 5:25:53 PM UTC
```

6. Note down the DATA, RECO and REDO locations. You'll need them later to set the `db_create_file_dest`, `db_create_online_log_dest`, and `db_recovery_file_dest` parameters for the database.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Choose an ORACLE_HOME

Decide which ORACLE_HOME to use for the database restore and then switch to that home with the correct ORACLE_BASE, ORACLE_HOME, and PATH settings. The ORACLE_HOME must not already be associated with a database.

To get a list of existing ORACLE_HOMEs and to ensure that the ORACLE_HOME is empty, use the Dbhome Commands and the Database Commands, respectively. To create a new ORACLE_HOME, use the Dbhome Commands.

Copy the Source Database Wallets

Skip this section if the source database is **not** configured with TDE.

1. On the DB system, become the oracle user:

   ```
   sudo su - oracle
   ```

2. Create the following directory, if it does not already exist:

   ```
   mkdir /opt/oracle/dcs/commonstore/wallets/tde/<db_unique_name>
   ```

3. Copy the ewallet.p12 file from the source database to the directory you created in the previous step.
4. On the target host, make sure that `$ORACLE_HOME/network/admin/sqlnet.ora` contains the following line:

   ```
   ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE)
   (METHOD_DATA=(DIRECTORY=/opt/oracle/dcs/commonstore/
   wallets/tde/$ORACLE_UNQNAME)))
   ```

   Add the line if it doesn't exist in the file. (The line might not be there if this is a new home and no database has been created yet on this host.)
5. Create the autologin wallet from the password-based wallet to allow auto-open of the wallet during restore and recovery operations.
   For a version 12.1 or later database, use the `ADMINISTER KEY MANAGEMENT` command:

   ```
   $cat create_autologin_12.sh

   #!/bin/sh
   if [ $# -lt 2 ]; then
           echo "Usage: $0 <dbuniquename><remotewalletlocation>"
           exit 1;
   fi

   mkdir /opt/oracle/dcs/commonstore/wallets/tde/$1
   cp $2/ewallet.p12* /opt/oracle/dcs/commonstore/wallets/tde/$1
   rm -f autokey.ora
   echo "db_name=$1"  > autokey.ora
   autokeystoreLog="autologinKeystore_`date +%Y%m%d_%H%M%S_%N`.log"
   echo "Enter Keystore Password:"
   read -s keystorePassword
   echo "Creating AutoLoginKeystore -> "
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
sqlplus "/as sysdba"  <<EOF
spool $autokeystoreLog
set echo on
startup nomount pfile=autokey.ora
ADMINISTER KEY MANAGEMENT CREATE AUTO_LOGIN KEYSTORE
FROM KEYSTORE '/opt/oracle/dcs/commonstore/wallets/tde/$1' -- Keystore
location
IDENTIFIED BY "$keystorePassword";
shutdown immediate;
EOF
```

Adjust the `cwallet.sso` permissions from `oracle:asmadmin` to `oracle:oinstall`.

```
ls -ltr /opt/oracle/dcs/commonstore/wallets/tde/<db_unique_name>
```

Output:

```
total 20
-rw-r--r-- 1 oracle oinstall 5680 Jul  6 11:39 ewallet.p12
-rw-r--r-- 1 oracle asmadmin 5725 Jul  6 11:39 cwallet.sso
```

For a version 11.2 database, use the `orapki` command:

```
orapki wallet create -wallet wallet_location -auto_login [-pwd password]
```

Install the Oracle Database Backup Module

The backup module JAR file is included on the DB system but you need to install it.

1. SSH to the DB system, log in as opc, and then become the oracle user.

   ```
   ssh -i <path to SSH key used when launching the DB System> opc@<DB System
   IP address or hostname>
   sudo su - oracle
   ```

2. Change to the directory that contains the backup module `opc_install.jar` file.

   ```
   cd /opt/oracle/oak/pkgrepos/orapkgs/oss/<version>/
   ```

3. To install the backup module, see the command syntax described in *Installing the Oracle Database Cloud Backup Module for OCI Classic* in Using Oracle Database Backup Cloud Service.

Set Environment Variables

Set the following environment variables for the RMAN and SQL*Plus sessions for the database:

```
ORACLE_HOME=<path of Oracle Home where the database is to be restored>
ORACLE_SID=<database instance name>
ORACLE_UNQNAME=<db_unique_name in lower case>
NLS_DATE_FORMAT="mm/dd/yyyy hh24:mi:ss"
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Allocate an RMAN SBT Channel

For each restore operation, allocate an SBT channel and set the SBT_LIBRARY parameter to the location of the `libopc.so` file and the OPC_FILE parameter to the location of the `opc_sbt.ora` file, for example:

```
ALLOCATE CHANNEL c1 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U' PARMS
'SBT_LIBRARY=/tmp/oss/libopc.so ENV=(OPC_PFILE=/<ORACLE_HOME>/dbs/
opc_sbt.ora)';
```

For more information about these files, see *Files Created When the Oracle Database Cloud Backup Module for OCI Classic is Installed* in Using Oracle Database Backup Cloud Service.

Ensure Decryption is Turned On

Make sure that decryption is turned on for all the RMAN restore sessions.

```
set decryption wallet open identified by <keystore password>;
```

For more information, see Providing the Password Required to Decrypt Encrypted Backups.

Restore Spfile

The following sample shell script restores the spfile. Set the `$dbID` variable to the dbid of the database being restored. By default, spfile is restored to `$ORACLE_HOME/dbs/spfile<sid>.ora`.

```
rman target / <<EOF

spool log to "`date +%Y%m%d_%H%M%S_%N`_dbid_${dbID}_restore_spfile.log"
startup nomount
set echo on
run {
ALLOCATE CHANNEL c1 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U' PARMS
'SBT_LIBRARY=/tmp/oss/libopc.so ENV=(OPC_PFILE=/tmp/oss/opc_sbt.ora)';
SET DBID=$dbID;
RESTORE SPFILE FROM AUTOBACKUP;
shutdown immediate;
EOF
```

Set the Database Parameters

1. Start the database in nomount mode.

   ```
   startup nomount
   ```

2. Update spfile and modify the following parameters.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- If the database storage type is ACFS, use the DATA, RECO, and REDO locations obtained from the `dbcli describe-dbstorage` command output, as described in Set Up Storage on the DB System:

```
alter system set db_create_file_dest='/u02/app/oracle/oradata/' scope =
spfile;
alter system set db_create_online_log_dest_1='/u03/app/oracle/redo'
scope = spfile;
alter system set db_recovery_file_dest='/u03/app/oracle/
fast_recovery_area' scope = spfile;
```

- If the database storage type is ASM:

```
alter system set db_create_file_dest='+DATA' scope = spfile;
alter system set db_create_online_log_dest_1='+RECO' scope = spfile;
alter system set db_recovery_file_dest='+RECO' scope = spfile;
```

- Set `db_recovery_file_dest_size` is not set or is set incorrectly:

```
alter system set db_recovery_file_dest_size=<sizeG> scope=spfile;
```

- Set `audit_file_dest` to the correct value:

```
alter system set audit_file_dest=/u01/app/oracle/admin/<db_unique_name
in lower case>/adump
```

3. Remove the `control_files` parameter. The Oracle Managed Files (OMF) parameters will be used to create the control file.

```
alter system reset control_files scope=spfile;
```

4. Restart the database in nomount mode using the newly added parameters.

```
shutdown immediate
startup nomount
```

Restore the Control File

Modify the following sample shell script for your environment to restore the control file. Set the `$dbID` variable to the dbid of the database being restored. Set SBT_LIBRARY to the location specified in the `-libDir` parameter when you installed the Backup Module. Set OPC-PFILE to the location specified in the `-configFile` parameter, which defaults to *ORACLE_HOME*`/dbs/opc`*SID*`.ora`.

```
rman target / <<EOF

spool log to "`date +%Y%m%d_%H%M%S_%N`_dbid_${dbID}_restore_controlfile.log"
set echo on
run {
ALLOCATE CHANNEL c1 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U' PARMS
'SBT_LIBRARY=/<Backup Module libDir>/libopc.so ENV=(OPC_PFILE=/<Backup Module
configFile>/opcSID.ora)';
SET DBID=$dbID;
RESTORE CONTROLFILE FROM AUTOBACKUP;
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
alter database mount;
}

exit;
EOF
```

Restore the Database

1. Preview and validate the backup. The database is now mounted and RMAN should be able to locate the backup from the restored controlfile. This step helps ensure that the list of archivelogs is present and that the backup components can be restored.
   In the following examples, modify SBT_LIBRARY and OPC_PFILE as needed for your environment.

   ```
   rman target / <<EOF

   spool log to "`date +%Y%m%d_%H%M%S_%N`_restore_database_preview.log"
   set echo on
   run {
       ALLOCATE CHANNEL c1 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U'
   PARMS 'SBT_LIBRARY=/tmp/oss/libopc.so ENV=(OPC_PFILE=/tmp/oss/
   opc_sbt.ora)';
       ALLOCATE CHANNEL c2 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U'
   PARMS 'SBT_LIBRARY=/tmp/oss/libopc.so ENV=(OPC_PFILE=/tmp/oss/
   opc_sbt.ora)';
       ALLOCATE CHANNEL c3 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U'
   PARMS 'SBT_LIBRARY=/tmp/oss/libopc.so ENV=(OPC_PFILE=/tmp/oss/
   opc_sbt.ora)';
       restore database validate header preview;
   }
   ```

   Review the output and if there are error messages, investigate the cause of the problem.

2. Redirect the restore using `set newname` to restore the data files in OMF format and use `switch datafile all` to allow the control file to update with the new data file copies.

   ```
   rman target / <<EOF

   spool log to "`date +%Y%m%d_%H%M%S_%N`_restore_database_preview.log"
   set echo on
   run {
       ALLOCATE CHANNEL c1 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U'
   PARMS 'SBT_LIBRARY=/tmp/oss/libopc.so ENV=(OPC_PFILE=/tmp/oss/
   opc_sbt.ora)';
       ALLOCATE CHANNEL c2 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U'
   PARMS 'SBT_LIBRARY=/tmp/oss/libopc.so ENV=(OPC_PFILE=/tmp/oss/
   opc_sbt.ora)';
       ALLOCATE CHANNEL c3 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U'
   PARMS 'SBT_LIBRARY=/tmp/oss/libopc.so ENV=(OPC_PFILE=/tmp/oss/
   opc_sbt.ora)';
       set newname for database to new;
       restore database;
       switch datafile all;
       switch tempfile all;
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
    recover database;
}
```

This recovery will attempt to use the last available archive log backup and then fail with an error, for example:

```
RMAN-00571: ===========================================================
RMAN-00569: =============== ERROR MESSAGE STACK FOLLOWS ===============
RMAN-00571: ===========================================================
RMAN-03002: failure of recover command at 07/20/2016 12:09:02
RMAN-06054: media recovery requesting unknown archived log for thread 1
with sequence 22 and starting SCN of 878327
```

3. To complete the incomplete recovery, run a recovery using the sequence number and thread number shown in the RMAN-06054 message, for example:

```
recover database until sequence 22 thread 1;
```

Reset the Logs

Reset the logs.

```
alter database open resetlogs;
```

Prepare to Register the Database

Before you register the database:

1. Make sure the database COMPATIBLE parameter value is acceptable. If the value is less than the minimum, the database cannot be registered until you upgrade the database compatibility.
2. Verify that the database has registered with the listener and the service name.

   ```
   lsnrctl services
   ```

3. Make sure the password file was restored or created for the new database.

   ```
   ls -ltr $ORACLE_HOME/dbs/orapw<oracle sid>
   ```

   If the file does not exist, create it using the orapwd utility.

   ```
   orapwd file=<$ORACLE_HOME/dbs/orapw<$ORACLE_SID>> password=<sys password>
   ```

4. Make sure the restored database if open in read write mode.

   ```
   select open_mode from v$database;
   ```

   The command output should indicate read write mode. The `dbcli register-database` command will attempt to run datapatch, which requires read write mode. If there are PDBs, they should also be in read write mode to ensure that datapatch runs on them.

Copyright © 2022, 2026, Oracle and/or its affiliates.

5. From oracle home on the restored database, use the following command verify the connection to SYS:

   ```
   conn sys/<password>@//<hostname>:1521/<database service name>
   ```

   This connection is required to register the database later. Fix any connection issues before continuing.

6. Make sure the database is running on spfile by using the SQL*Plus command.

   ```
   SHOW PARAMETERS SPFILE
   ```

7. (Optional) If you would like to manage the database backup with the dbcli command line interface, you can associate a new or existing backup configuration with the migrated database when you register it or after you register it. A backup configuration defines the backup destination and recovery window for the database. Use the Backupconfig Commands to create, list, and display backup configurations.

8. Copy the folder `$ORACLE_HOME/sqlpatch` from source database to the target database. This will enable the `dbcli register-database` command to roll back any conflicting patches.

> **Note**
>
> If you are migrating a version 11.2 database, additional steps are required after you register the database. For more information, see Roll Back Patches on a Version 11.2 Database.

Register the Database on the DB System

The Database Commands registers the restored database to the dcs-agent so it can be managed by the dcs-agent stack.

> **Note**
>
> The `dbcli register-database` command is not available on 2-node RAC DB systems.

As the root user, use the `dbcli register-database` command to register the database on the DB system, for example:

```
dbcli register-database --dbclass OLTP --dbshape odb1 --servicename tdetest --
syspassword
```

Output:

```
Password for SYS:
{
  "jobId" : "317b430f-ad5f-42ae-bb07-13f053d266e2",
  "status" : "Created",
  "message" : null,
  "reports" : [ ],
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
  "createTimestamp" : "August 08, 2016 05:55:49 AM EDT",
  "description" : "Database service registration with db service name:
tdetest",
  "updatedTime" : "August 08, 2016 05:55:49 AM EDT"
}
```

Update tnsnames.ora

Check the `tnsnames.ora` in the backup location, check the database links used in the cloned database, and then add any relevant connection strings to the cloned database file at `$ORACLE_HOME/network/admin/tnsnames.ora`.

Roll Back Patches on a Version 11.2 Database

For version 11.2 databases, the sqlpatch application is not automated, so any interim patches (previously known as a "one-off" patches) applied to the source database that are not part of the installed PSU must be rolled back manually in the target database. After registering the database, execute the `catbundle.sql` script and then the `postinstall.sql` script with the corresponding PSU patch (or the overlay patch on top of the PSU patch), as described below.

> **Tip**
>
> Some interim patches may include files written to the $ORACLE_HOME/rdbms/admin directory as well as the $ORACLE_HOME/sqlpatch directory. Oracle recommends that you roll back these patches in the source database using the instructions in the patch read-me prior to migrating the database to OCI environment. Contact Oracle Support if you need assistance with rolling back these patches.

1. On the DB System, use the `dbcli list-dbhomes` command to find the PSU patch number for the version 11.2 database home. In the following sample command output, the PSU patch number is the second number in the DB Version column:

   ```
   dbcli  list-dbhomes
   ```

   Output:

   ```
   ID                                   Name                DB
   Version                              Home
   Location                             Status
   ------------------------------------ -----------------
   ----------------------------------
   ---------------------------------------- ----------
   59d9bc6f-3880-4d4f-b5a6-c140f16f8c64 OraDB11204_home1    11.2.0.4.160719
   (23054319, 23054359)   /u01/app/oracle/product/11.2.0.4/dbhome_1 Configured
   ```

   (The first patch number, 23054319 in the example above, is for the OCW component in the database home.)

2. Find the overlay patch, if any, by using the `lsinventory` command. In the following example, patch number 24460960 is the overlay patch on top of the 23054359 PSU patch.

   ```
   $ORACLE_HOME/OPatch/opatch lsinventory
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Output:

```
...
Installed Top-level Products (1):

Oracle Database 11g
11.2.0.4.0
There are 1 products installed in this Oracle Home.


Interim patches (5) :

Patch  24460960     : applied on Fri Sep 02 15:28:17 UTC 2016
Unique Patch ID:  20539912
   Created on 31 Aug 2016, 02:46:31 hrs PST8PDT
   Bugs fixed:
     23513711, 23065323, 21281607, 24006821, 23315889, 22551446, 21174504
   This patch overlays patches:
     23054359
   This patch needs patches:
     23054359
   as prerequisites
```

3. Start SQL*Plus and execute the `catbundle.sql` script, for example:

```
startup
connect / as sysdba
@$ORACLE_HOME/rdbms/admin/catbundle.sql psu apply
exit
```

4. Apply the sqlpatch, using the overlay patch number from the previous step, for example:

```
connect / as sysdba
@$ORACLE_HOME/sqlpatch/24460960/postinstall.sql
exit
```

> **Note**
>
> If the source database has one-off patches installed and those patches are **not** part of the installed PSU in the cloud environment, then the SQL changes that correspond to those one-off patches need to be rolled back. To rollback the SQL changes, copy the `$ORACLE_HOME/sqlpatch/<patch#>/postdeinstall.sql` script from the source environment to the cloud environment and execute the `postdeinstall.sql` script.

Post Restore Checklist

After the database is restored and registered on the DB system, use the following checklist to verify the results and perform any post-restore customizations.

1. Make sure the database files were restored in OMF format.
2. Make sure the database is listed in the Database Commands output.
3. Check for the following external references in the database and update them if necessary:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- External tables: If the source database uses external tables, back up that data and migrate it to the target host.
- Directories: Customize the default directories as needed for the restored database.
- Database links: Make sure all the required TNS entries are updated in the `tnsnames.ora` file in ORACLE_HOME.
- Email and URLs: Make sure any email addresses and URLs used in the database are still accessible from the DB system.
- Scheduled jobs: Review the jobs scheduled in source database and schedule similar jobs as needed in the restored database.

4. If you associated a backup configuration when you registered the database, run a test back up using the Backup Commands.
5. If the restored database contains a CDB and PDBs, verify that patches have been applied to all PDBs.

Oracle Data Guard

Use Oracle Data Guard on a DB System

You can use Oracle Data Guard on a DB system. Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data.

For more information on Oracle Data Guard, see Introduction to Oracle Data Guard.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be granted security access in a **policy** by an administrator. This access is required whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you get a message that you don’t have permission or are unauthorized, verify with your administrator what type of access you have and which **compartment** to work in.

If you're new to policies, see Getting Started with Policies and Common Policies.

General Information

Oracle Data Guard configuration requires at least two databases: one in a primary role and one in a standby role. The primary and standby databases together constitute a Data Guard configuration. Most of the applications access the primary database. A standby database is a transactionally consistent copy of the primary database.

Data Guard Group provides the ability to create and manage multiple local and remote standby databases linked to a primary database, providing flexibility for both data protection and disaster recovery. In a Data Guard Group, a primary database can support up to a maximum of six standby databases.

Oracle Data Guard maintains the standby database by transmitting and applying redo data from the primary database. If the primary database becomes unavailable, then you can use Data Guard to switchover or failover the standby database to the primary role. This is true even if you have more than one standby database.

- **Switchover**: A switchover reverses the primary and standby database roles.
- **Failover**: A failover transitions the standby database into the primary role after the existing primary database fails or becomes unreachable.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Reinstate**: Reinstates a database into the standby role in a Data Guard configuration.

In a typical Data Guard configuration, two standby databases are commonly used:

- **Local Standby**: A standby database in the same region as the production database is ideal for failover scenarios, offering zero data loss for local failures (such as database, cluster, or availability domain failures). Application failover impact is reduced in this case, as applications continue operating without the performance overhead of communicating with a remote region.
- **Remote (Cross-Region) Standby**: A remote standby database, located in a different region, is typically used for disaster recovery or to offload read-only query processing. A remote standby database setup ensures data protection against regional failures.

Some enterprise customers aim for symmetry after a site switch. For example, they may prefer to have both the primary and local standby in Region 1, and a remote standby with its own local standby in Region 2. In this configuration, there will be three standby databases. After a site switch, you will still have a primary database and a local standby readily available in the new primary region.

Additionally, you can enhance their configurations by adding another standby database for testing purposes, leveraging our snapshot (read/write) standby capabilities.

A Data Guard configuration requires atleast two DB systems, one containing the primary database and one containing the standby database. When you enable Data Guard for a database, a new DB system with the standby database is created and associated with the primary database.

Additional requirement details are as follows:

- The database versions and editions must be identical. For 26ai databases, the primary and standby databases must be on the same major release version, while the standby database can be on a higher minor version.
- Data Guard does not support Oracle Database Standard Edition. (Active Data Guard requires Enterprise Edition Extreme Performance.)
- Each database in a Data Guard configuration must have a unique name (`DB_UNIQUE_NAME`) value that is not in use by other databases in the DB systems that house the Data Guard configuration. However, the primary and standby database can use the same database name `DB_NAME` value.
- The database edition determines whether Active Data Guard (ADG) can be used. ADG is only available with Enterprise Edition Extreme Performance. If you are using the BYOL licensing model and if your license does not include ADG, then you must ensure that ADG is not enabled when configuring Data Guard for Enterprise Edition Extreme Performance. Alternately, you can use Enterprise Edition or Enterprise Edition High Performance, which do not enable ADG by default. See *Use Oracle Data Guard with the Database CLI*.
- If your primary and standby databases are in the same region, then both must use the same virtual cloud network (VCN).
- If your primary and standby databases are in different regions, then you must peer the virtual cloud networks (VCNs) for each database. See Remote VCN Peering using an RPC.
- The primary and standby databases could be in a different region, but they must be in the same realm.
- The `SYS` password and the TDE wallet password of the primary and standby databases must all be the same.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- A cross-region Data Guard is also supported if the database is using OCI Vault for TDE keys.
- Configure the security list ingress and egress rules for the subnets of the DB systems in a Data Guard configuration to enable TCP traffic to move between the applicable ports. Ensure that the rules you create are stateful (the default).

  For example, if the subnet of the primary DB system uses the source CIDR 10.0.0.0/24 and the subnet of the standby DB system uses the source CIDR 10.0.1.0/24, then create rules as shown in the subsequent example.

  > **Note**
  >
  > The egress rules in the example show how to enable TCP traffic only for port 1521, which is a minimum requirement for the Data Guard to work. If TCP traffic is already enabled on all of your outgoing ports (0.0.0.0/0), then you do not need to explicitly add these specific egress rules.

- The standby databases in OCI are physical standbys.
- Creating a standby database associated with another standby database ("cascading standby") is not supported.
- You can create a Data Guard configuration when the database is encrypted using OCI Vault. However, OCI Vault currently supports key replication to only one remote region. As a result, cross-region Data Guard configurations for the databases using OCI Vault are limited to two regions: the primary region and one remote region. Multiple standby databases can be created within these regions, but configurations spanning three or more regions are not supported.

Known Issues

- In a cross-region Data Guard, after you perform a switchover, you cannot rotate the Database KMS key on the new primary database.
- In a cross-region Data Guard, after you perform a switchover, you cannot create a PDB on the new primary database.

Security List for Subnet on the Primary DB System

```
Ingress Rules:
    Stateless: No
    Source: 10.0.1.0/24
    IP Protocol: TCPSource Port Range: All
    Destination Port Range: 1521
    Allows: TCP traffic for ports: 1521

Egress Rules:
    Stateless: No
    Destination: 10.0.1.0/24
    IP Protocol: TCP
    Source Port Range: All
    Destination Port Range: 1521
    Allows: TCP traffic for ports: 1521
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Security List for Subnet on the Standby DB System

```
Ingress Rules:
    Stateless: No
    Source: 10.0.0.0/24
    IP Protocol: TCP
    Source Port Range: All
    Destination Port Range: 1521
    Allows: TCP traffic for ports: 1521

Egress Rules:
    Stateless: No
    Destination: 10.0.0.0/24
    IP Protocol: TCP
    Source Port Range: All
    Destination Port Range: 1521
    Allows: TCP traffic for ports: 1521
```

For information about creating and editing rules, see Security Lists.

Availability Domain and Fault Domain Considerations for Oracle Data Guard

Oracle recommends that the DB system that contains the standby database be in a different availability domain from that of the DB system containing the primary database to improve availability and disaster recovery. If you enable Oracle Data Guard for a database and your standby database is in the same availability domain as the primary database (either by choice, or because you are working in a single availability domain region), then Oracle recommends that you place the standby database in a different fault domain from that of the primary database.

> **Note**
>
> If your primary and standby databases are two-node Oracle RAC databases and both are in the same availability domain, then only one of the two nodes of the standby database can be in a fault domain that does not include any other nodes from either the primary or standby database. This is because each availability domain has only three fault domains, and the primary and standby databases have a combined total of four nodes. For more information on availability domains and fault domains, see Regions and Availability Domains.

Enable Oracle Data Guard on a DB System

This article provides the details and procedure to enable Data Guard on a DB system.

General Information

- When you enable Oracle Data Guard, a new Data Guard configuration will be created for the primary and the standby databases as a Data Guard Group resource with multiple standby databases or the older Data Guard association with a single standby database.
- If you have an existing Data Guard setup, you can continue to use current capabilities with no impact. However, if you wish to create multiple standby databases, you will need to switch to Data Guard Group, which can be done at any time.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- When you enable Oracle Data Guard, a new DB system must be created for the standby database.
- Ampere A1 shape-based DB systems do not support Data Guard with Intel or AMD shape-based DB systems.

Create the first standby

Perform the following steps to enable Data Guard on a DB system by creating the first standby for a database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, select the **Data Guard Associations** tab.
4. Select **Add standby**.
5. On the **Add standby** page, create a new peer DB system for the standby by entering the following information.

Data Guard resource type

You must select the Data Guard resource type when you are creating the first standby database.

- **Use the new Data Guard Group Resource**: With this option, your new Data Guard configuration will be created as a Data Guard Group resource. This option supports adding multiple standby databases and provides other enhancements. If you currently have automation that manages Data Guard operations using the existing Data Guard Association, you can update your applications to use the new Data Guard Group resource to take advantage of these new capabilities.
- **Use the existing Data Guard Association Resource**: Choose this option if your automation for managing Data Guard operations relies on the existing Data Guard Association resource. However, you will not be able to add multiple standby databases and will not get the enhancements provided by the new Data Guard Group resource.

> **Note**
>
> You can switch to the new Data Guard Group resource at any time without impact to your database.

Peer DB system

- **DB system name**: Enter a user-friendly name to help you easily identify the resource. Display name can be changed at any time.
- **Region**: Select the region of the new peer DB system. For more information on regions and availability domains, see About Regions and Availability Domains.
- **Availability domain**: Select the availability domain of the new peer DB system.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Shape

Shape: The shape determines the type of DB system created and the resources allocated to the system. For a complete list of shapes, see Available Shapes and How It Determines the Resources Allocated.

- By default, the **AMD VM.Standard.E4.Flex** shape with **4 OCPU**s is selected.
- The Enterprise Edition Developer is supported only on the Ampere A1 shape.
- If **Enterprise Edition Developer** is selected for the Oracle Database software edition, then by default, the **VM.Standard.A1.Flex** shape with **1 OCPU** is selected.
- The Standard Edition is not supported on the Ampere A1 shape.
- By default, the same shape as the primary is selected for standby.
- To specify a shape other than the default, select **Change shape**.

In the **Change shape** panel, provide the following details to select from available shapes:

Shape series

Select an AMD, Intel, or Ampere processor in the processor group.

- **AMD:** Shapes that use current-generation AMD processors. The AMD shapes are flexible.
- **Intel:** Standard and optimized shapes that use current-generation Intel processors. Both fixed and flexible Intel shapes are available.
- **Ampere**: Shapes that use Arm-based Ampere processors. The Ampere shapes are flexible.

> **Note**
>
> If you select Ampere A1, AMD E4, AMD E5, or Intel X9 flexible shapes, the memory, network bandwidth, and maximum theoretical IOPS scale proportionally.

Configure OCPU

Select the shape you want to use for this instance.

To change the number of OCPUs, perform the following steps:

- From the actions menu of the desired shape, select **Update OCPU count**.
- Select the desired number of OCPUs per node from the list.
- Select **Update**.

You can change the number of OCPUs for Ampere A1, AMD E4, AMD E5, and Intel X9 flexible shapes.

- For the Ampere A1 shape, a minimum of 1 OCPU and a maximum of 57 OCPUs can be selected.
- For AMD E4 and E5 shapes, a minimum of 1 OCPU and a maximum of 64 OCPUs can be selected.
- For the Intel X9 shape, a minimum of 1 OCPU and a maximum of 32 OCPUs can be selected.

The following resources scale proportionately to the number of **OCPUs** you selected.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Memory (GB)**: The amount of memory you want to allocate to this instance.
  For Ampere A1, AMD E4, AMD E5, and Intel X9 shapes, the memory will scale proportionally based on the number of OCPUs selected.
  - For the Ampere A1 shape, for each OCPU, 8 GB of memory is allocated. A minimum of 8 GB and a maximum of 456 GB of memory is allocated.
  - For AMD E4 and E5 shapes, for each OCPU, 16 GB of memory is allocated. A minimum of 16 GB and a maximum of 1024 GB of memory is allocated.
  - For the Intel X9 shape, for each OCPU, 16 GB of memory is allocated. A minimum of 16 GB and a maximum of 512 GB of memory is allocated.
- **Network bandwidth (Gbps)**: The amount of network bandwidth you want to allocate to this instance.
  For Ampere A1, AMD E4, AMD E5, and Intel X9 shapes, the bandwidth will scale proportionally based on the number of OCPUs selected. For each OCPU, 1 Gbps of network bandwidth is allocated.
  - For the Ampere A1 shape, a minimum of 1 Gbps and a maximum of 40 Gbps of network bandwidth is allocated.
  - For AMD E4 and E5 shapes, a minimum of 1 Gbps and a maximum of 40 Gbps of network bandwidth is allocated.
  - For the Intel X9 shape, a minimum of 1 Gbps and a maximum of 32 Gbps of network bandwidth is allocated.
- **Theoretical max IOPS**: The amount of input and output per second (IOPS) you want to allocate to this instance. Theoretical max IOPS is also dependent on the storage you select.
  For standard x86 shapes, the theoretical max IOPS will scale proportionally based on the number of ECPUs selected. For each ECPU, 8K theoretical max IOPS is allocated.

  For Ampere A1, AMD E4, AMD E5, and Intel X9 shapes, the theoretical max IOPS will scale proportionally based on the number of OCPUs selected. For each OCPU, 16K theoretical max IOPS is allocated.
  - For the Ampere A1 shape, a minimum of 16K and a maximum of 640K theoretical max IOPS is allocated.
  - For AMD E4 and E5 shapes, a minimum of 16K and a maximum of 640K theoretical max IOPS is allocated.
  - For the Intel X9 shape, a minimum of 16K to a maximum of 512K theoretical max IOPS is allocated.

DB system configuration

- **Total node count**: (Read-only) The number of nodes that is allocated to the standby instance. The node count will be the same as the primary node count.
- **License type**: The type of license you want to use for the DB system. Your choice affects metering for billing.
  - **License included** means the cost of this OCI Database service resource will include both the Oracle Database software licenses and the service.
  - **Bring Your Own License (BYOL)** means you will use your organization's Oracle Database software licenses for this OCI Database service resource. For more information, see Bring Your Own License.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Network information

- Select **Compartment** to select a VCN in a different compartment.
- **Virtual cloud network**: From the list, select a VCN in which to create the DB system. Select **Change compartment** to select a VCN in a different compartment.
- Select **Compartment** to select a subnet in a different compartment.
- **Client subnet**: The subnet to which the DB system attaches. For both single-node and multi-node RAC DB systems, do not use a subnet that overlaps with 192.168.16.16/28, which is used by the Oracle Clusterware private interconnect on the database instance. Specifying an overlapping subnet causes the private interconnect to malfunction.

  Choose a dual stack subnet if you want to configure the DB system with both IPv4 and IPv6 addresses.
- **Network security groups**: Optionally, you can specify one or more network security groups (NSGs) for your DB system. NSGs function as virtual firewalls, enabling you to apply a set of ingress and egress security rules to your DB system. A maximum of five NSGs can be specified.

  For more information, see Access and Security and *Security Rules for the DB System*.

  > **Note**
  >
  > If you select a subnet with a security list, the security rules for the DB system will be a union of the rules in the security list and the NSGs.

  To use network security groups:

  – Switch on the **Use network security groups to control traffic** toggle. Note that you must have a virtual cloud network selected to be able to assign NSGs to your DB system.

  – Specify the NSG to use with the DB system. You may need to use more than one NSG. If you're not sure, contact your network administrator.

  – To use additional NSGs, select **+ Another network security group**.
- **Hostname prefix**: Enter a hostname prefix for the DB system. The host name must begin with an alphabetic character and can contain only alphanumeric characters and hyphens (-). The maximum number of characters allowed is 16.

  > **Caution**
  >
  > The host name must be unique within the subnet. If it is not unique, the DB system will fail to provision.
- **Host domain name**: The domain name for the DB system. If the selected subnet uses the Oracle-provided Internet and VCN Resolver for DNS name resolution, then this field displays the domain name for the subnet, and it can't be changed. Otherwise, you can provide your choice of a domain name. Hyphens (-) are not permitted.
- **Host and domain URL**: Combines the host and domain names to display the fully qualified domain name (FQDN) for the database. The maximum length is 64 characters.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Private IP type**: Optionally, for non-RAC DB systems, you can define the IP address of the new DB system. This is useful in development contexts where you create and delete a DB system over and over, and you need each new iteration of the DB system to use the same IP address. If you specify an IP address that is currently in use within the subnet, the provisioning operation will fail with an error message regarding the invalid IP address.

  If a dual stack subnet is selected, then both IPv4 and IPv6 address options are displayed.

  - **IPv4 address**: You can either automatically assign an IPv4 address or enter it manually.
    * Select the **Automatically assign IPv4 addresses from subnet** option to assign an address automatically.
    * Select the **Manually assign IPv4 addresses** option to manually enter a private IP address. The IP address should be within the subnet CIDR range.
  - **IPv6 address**: You can either automatically assign an IPv6 address or enter it manually.
    * Select the **Automatically assign IPv6 addresses from subnet** option to assign an address automatically.
    * Select the **Manually assign IPv6 addresses** option to manually enter an IP address. The IP address should be within the subnet CIDR range.

Data Guard details

- **Data Guard type**: Select **Active Data Guard** or **Data Guard**. Active Data Guard provides additional features including: Real-Time Query and DML Offload, Automatic Block Repair, Standby Block Change Tracking, Far Sync, Global Data Services, and Application Continuity.
  - The Active Data Guard requires an Oracle Active Data Guard license.
  - For more information, see Introduction to Oracle Data Guard.
- **Protection mode**: (read-only) Default is **Max Performance**. The other possible protection modes are **Maximum Performance** and **Maximum Availability**. For more information, see Oracle Data Guard Protection Modes.
- **Transport type**: (read-only) Default is **Async**. The redo transport type used for this Data Guard. For more information, see Managing Redo Transport Services for Data Protection Modes.

> **Note**
>
> You cannot set the Protection Mode or Transport Type for the standby database during its creation. However, it is possible to edit it later. You can also edit other Data Guard details after provisioning if you need to. For more information, see *Edit the Oracle Data Guard Association*.

Diagnostics collection

The diagnostics collection and notifications feature enables Oracle Cloud Operations and you to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to events to get notified about resource state changes. You can enable or disable this feature at anytime.

Copyright © 2022, 2026, Oracle and/or its affiliates.

By default the options are selected for enabling. However, you can select to uncheck the diagnostic collection check boxes if you do not require the diagnostic feature.

- **Enable diagnostic events**: Enables and allows Oracle to collect and send fault notifications about critical, warning, and information events for you.
- **Enable health monitoring**: This diagnostics collection for Oracle Cloud operations viewing is not available for the Base Database Service.
- **Enable incident logs and trace collection**: Enables and allows Oracle to receive event notifications and collect incident logs and traces for fault diagnosis and issue resolution.

> **Note**
>
> You are opting in with the understanding that the list of events and log files can change in the future. You can opt out of this feature at any time.

Advanced options

Expand **Advanced options** to provide the advanced options for this resource.

Management

Expand **Management** to provide the following details:

- **Fault domain**: The fault domain(s) in which the DB system resides. You can select which fault domain to use for your DB system. For multi-node RAC DB systems, you can specify which two fault domains to use. Oracle recommends that you place each node of a multi-node RAC DB system in a different fault domain. For more information about fault domains, see About Regions and Availability Domains.
- **Time zone**: The default time zone for the DB system is UTC, but you can specify a different time zone. The time zone options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system. For more information, see DB System Time Zone. The following options are available:
  - **UTC:** configures your DB system to use coordinated universal time.
  - **Browser-detected:** The console displays the time zone detected by your browser for this option.
  - **Select another time zone:** To manually specify a time zone, first make a choice using the **Region or country selector** to select a geographic region, then use the **Time zone** selector to select your required time zone.

> **Tip**
>
> If you want to set a time zone other than UTC or the browser-detected time zone, and if you do not see the time zone you want, try selecting "Miscellaneous" in the **Region or country** list.

Security

Expand **Security** to provide security details. Optionally, you can specify one or more security attributes to configure Zero Trust Packet Routing (ZPR) for the DB system.

- Select the **Namespace** in which the required security attribute is available.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Select the **Key** and **Value** of the required security attribute.
- Select **Add security attribute**.

> **Note**
>
> - Administrators must set up security attribute namespaces and security attributes in a tenancy before users can apply security attributes to the DB systems.
> - A security attribute is effective only with appropriate policies. If a security attribute without any policies is added, all access will be denied by default, even if allowed in the Security List or NSGs.
> - If you use the security attribute, the security rules for the DB system will be a union of the rules in the security attributes along with any rules in the security list and the NSGs.
> - You may need to use more than one security attribute. If you're not sure, contact your network administrator.
> - A maximum of 3 security attributes can be specified for a DB system.

For more information about:

- Add, edit, or remove a security attribute after a DB system is created, see *Manage Security Attributes for the DB System*.
- Defining appropriate ZPR policies, see *Zero Trust Packet Routing*.
- Creating a security attribute, see Creating a Security Attribute.

Tags

Expand **Tags** to provide tag details for this resource.

You can add free-form tags or defined tags to this resource. You must have permission to use the tag namespace for defined tags. For information about using tags to manage your OCI resources, see Resource Tags.

> **Tip**
>
> Tags can be applied later or at any time to a resource.

Standby database

- **Database image:** *Optional*. By default, the latest database software image as the source database is used. You can specify what Oracle Database version is used for the database. You can mix database versions on the DB system, but not editions.
  - Select **Change database image** to choose a custom database software image that you or someone in your organization has created in your tenancy.
  - Use the **Compartment** filter to select the custom database software images from a specific compartment.
  - Select a database image from the table of available images for the Oracle Database version you selected.
  - Select the **Select** button.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Password**: Enter a password for the administrator. The password must meet the following criteria:
  - A strong password for SYS, SYSTEM, TDE wallet, and PDB administrator.
  - The password must be 9 to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters.
  - The special characters must be _, #, or -.
  - The password must not contain the user name (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reversed order and regardless of casing.

Enable

Select **Add**.

When you create Data Guard, the details for a database and its peer display their respective roles as **Primary** or **Standby**.

If setting up Data Guard fails for any reason, we recommend contacting Oracle Support to resolve the issue. The billing for the DB system starts as soon as it is created and is in **Available** status.

Create multiple standby databases

In a Data Guard group, a primary database can support up to a maximum of six standby databases.

Perform the following steps to create multiple standby databases in a Data Guard group using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, select the **Data Guard Group** tab.
4. Select **Add standby**.
5. On the **Add standby** page, create a new peer DB system for the standby by entering the following information.

Peer DB system

- **DB system name**: Enter a user-friendly name to help you easily identify the resource. Display name can be changed at any time.
- **Region**: Select the region of the new peer DB system. For more information on regions and availability domains, see About Regions and Availability Domains.
- **Availability domain**: Select the availability domain of the new peer DB system.

Shape

Shape: The shape determines the type of DB system created and the resources allocated to the system. For a complete list of shapes, see Available Shapes and How It Determines the Resources Allocated.

- By default, the **AMD VM.Standard.E4.Flex** shape with **4 OCPU**s is selected.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- The Enterprise Edition Developer is supported only on the Ampere A1 shape.
- If **Enterprise Edition Developer** is selected for the Oracle Database software edition, then by default, the **VM.Standard.A1.Flex** shape with **1 OCPU** is selected.
- The Standard Edition is not supported on the Ampere A1 shape.
- By default, the same shape as the primary is selected for standby.
- To specify a shape other than the default, select **Change shape**.

In the **Change shape** panel, provide the following details to select from available shapes:

Shape series

Select an AMD, Intel, or Ampere processor in the processor group.

- **AMD:** Shapes that use current-generation AMD processors. The AMD shapes are flexible.
- **Intel:** Standard and optimized shapes that use current-generation Intel processors. Both fixed and flexible Intel shapes are available.
- **Ampere**: Shapes that use Arm-based Ampere processors. The Ampere shapes are flexible.

> **Note**
>
> If you select Ampere A1, AMD E4, AMD E5, or Intel X9 flexible shapes, the memory, network bandwidth, and maximum theoretical IOPS scale proportionally.

Configure OCPU

Select the shape you want to use for this instance.

To change the number of OCPUs, perform the following steps:

- From the actions menu of the desired shape, select **Update OCPU count**.
- Select the desired number of OCPUs per node from the list.
- Select **Update**.

You can change the number of OCPUs for Ampere A1, AMD E4, AMD E5, and Intel X9 flexible shapes.

- For the Ampere A1 shape, a minimum of 1 OCPU and a maximum of 57 OCPUs can be selected.
- For AMD E4 and E5 shapes, a minimum of 1 OCPU and a maximum of 64 OCPUs can be selected.
- For the Intel X9 shape, a minimum of 1 OCPU and a maximum of 32 OCPUs can be selected.

The following resources scale proportionately to the number of **OCPUs** you selected.

- **Memory (GB)**: The amount of memory you want to allocate to this instance.
  For Ampere A1, AMD E4, AMD E5, and Intel X9 shapes, the memory will scale proportionally based on the number of OCPUs selected.
  - For the Ampere A1 shape, for each OCPU, 8 GB of memory is allocated. A minimum of 8 GB and a maximum of 456 GB of memory is allocated.

Copyright © 2022, 2026, Oracle and/or its affiliates.

  - For AMD E4 and E5 shapes, for each OCPU, 16 GB of memory is allocated. A minimum of 16 GB and a maximum of 1024 GB of memory is allocated.
  - For the Intel X9 shape, for each OCPU, 16 GB of memory is allocated. A minimum of 16 GB and a maximum of 512 GB of memory is allocated.
- **Network bandwidth (Gbps)**: The amount of network bandwidth you want to allocate to this instance.
  For Ampere A1, AMD E4, AMD E5, and Intel X9 shapes, the bandwidth will scale proportionally based on the number of OCPUs selected. For each OCPU, 1 Gbps of network bandwidth is allocated.
  - For the Ampere A1 shape, a minimum of 1 Gbps and a maximum of 40 Gbps of network bandwidth is allocated.
  - For AMD E4 and E5 shapes, a minimum of 1 Gbps and a maximum of 40 Gbps of network bandwidth is allocated.
  - For the Intel X9 shape, a minimum of 1 Gbps and a maximum of 32 Gbps of network bandwidth is allocated.
- **Theoretical max IOPS**: The amount of input and output per second (IOPS) you want to allocate to this instance. Theoretical max IOPS is also dependent on the storage you select.
  For standard x86 shapes, the theoretical max IOPS will scale proportionally based on the number of ECPUs selected. For each ECPU, 8K theoretical max IOPS is allocated.

  For Ampere A1, AMD E4, AMD E5, and Intel X9 shapes, the theoretical max IOPS will scale proportionally based on the number of OCPUs selected. For each OCPU, 16K theoretical max IOPS is allocated.
  - For the Ampere A1 shape, a minimum of 16K and a maximum of 640K theoretical max IOPS is allocated.
  - For AMD E4 and E5 shapes, a minimum of 16K and a maximum of 640K theoretical max IOPS is allocated.
  - For the Intel X9 shape, a minimum of 16K to a maximum of 512K theoretical max IOPS is allocated.

DB system configuration

- **Total node count**: (Read-only) The number of nodes that is allocated to the standby instance. The node count will be the same as the primary node count.
- **License type**: The type of license you want to use for the DB system. Your choice affects metering for billing.
  - **License included** means the cost of this OCI Database service resource will include both the Oracle Database software licenses and the service.
  - **Bring Your Own License (BYOL)** means you will use your organization's Oracle Database software licenses for this OCI Database service resource. For more information, see Bring Your Own License.

Network information

- Select **Compartment** to select a VCN in a different compartment.
- **Virtual cloud network**: From the list, select a VCN in which to create the DB system. Select **Change compartment** to select a VCN in a different compartment.
- Select **Compartment** to select a subnet in a different compartment.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Client subnet**: The subnet to which the DB system attaches. For both single-node and multi-node RAC DB systems, do not use a subnet that overlaps with 192.168.16.16/28, which is used by the Oracle Clusterware private interconnect on the database instance. Specifying an overlapping subnet causes the private interconnect to malfunction.

  Choose a dual stack subnet if you want to configure the DB system with both IPv4 and IPv6 addresses.

- **Network security groups**: Optionally, you can specify one or more network security groups (NSGs) for your DB system. NSGs function as virtual firewalls, enabling you to apply a set of ingress and egress security rules to your DB system. A maximum of five NSGs can be specified.

  For more information, see Access and Security and *Security Rules for the DB System*.

  > **Note**
  >
  > If you select a subnet with a security list, the security rules for the DB system will be a union of the rules in the security list and the NSGs.

  To use network security groups:

  – Switch on the **Use network security groups to control traffic** toggle. Note that you must have a virtual cloud network selected to be able to assign NSGs to your DB system.

  – Specify the NSG to use with the DB system. You may need to use more than one NSG. If you're not sure, contact your network administrator.

  – To use additional NSGs, select **+ Another network security group**.

- **Hostname prefix**: Enter a hostname prefix for the DB system. The host name must begin with an alphabetic character and can contain only alphanumeric characters and hyphens (-). The maximum number of characters allowed is 16.

  > **Caution**
  >
  > The host name must be unique within the subnet. If it is not unique, the DB system will fail to provision.

- **Host domain name**: The domain name for the DB system. If the selected subnet uses the Oracle-provided Internet and VCN Resolver for DNS name resolution, then this field displays the domain name for the subnet, and it can't be changed. Otherwise, you can provide your choice of a domain name. Hyphens (-) are not permitted.
- **Host and domain URL**: Combines the host and domain names to display the fully qualified domain name (FQDN) for the database. The maximum length is 64 characters.
- **Private IP type**: Optionally, for non-RAC DB systems, you can define the IP address of the new DB system. This is useful in development contexts where you create and delete a DB system over and over, and you need each new iteration of the DB system to use the same IP address. If you specify an IP address that is currently in use within the subnet, the provisioning operation will fail with an error message regarding the invalid IP address.

  If a dual stack subnet is selected, then both IPv4 and IPv6 address options are displayed.

  – **IPv4 address**: You can either automatically assign an IPv4 address or enter it manually.

Copyright © 2022, 2026, Oracle and/or its affiliates.

        * Select the **Automatically assign IPv4 addresses from subnet** option to assign an address automatically.
        * Select the **Manually assign IPv4 addresses** option to manually enter a private IP address. The IP address should be within the subnet CIDR range.
    - **IPv6 address**: You can either automatically assign an IPv6 address or enter it manually.
        * Select the **Automatically assign IPv6 addresses from subnet** option to assign an address automatically.
        * Select the **Manually assign IPv6 addresses** option to manually enter an IP address. The IP address should be within the subnet CIDR range.

Data Guard details

- **Data Guard type**: Select **Active Data Guard** or **Data Guard**. Active Data Guard provides additional features including: Real-Time Query and DML Offload, Automatic Block Repair, Standby Block Change Tracking, Far Sync, Global Data Services, and Application Continuity.
    - The Active Data Guard requires an Oracle Active Data Guard license.
    - For more information, see Introduction to Oracle Data Guard.
- **Protection mode**: (read-only) Default is inherited from the first standby. The other possible protection modes are **Maximum Performance** and **Maximum Availability**. For more information, see Oracle Data Guard Protection Modes.
- **Transport type**: (read-only) Default is **Async**. The redo transport type used for this Data Guard. For more information, see Managing Redo Transport Services for Data Protection Modes.

> **Note**
>
> You cannot set the Protection Mode or Transport Type for the standby database during its creation. However, it is possible to edit it later. You can also edit other Data Guard details after provisioning if you need to. For more information, see *Edit the Oracle Data Guard Association*.

Diagnostics collection

The diagnostics collection and notifications feature enables Oracle Cloud Operations and you to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to events to get notified about resource state changes. You can enable or disable this feature at anytime.

By default the options are selected for enabling. However, you can select to uncheck the diagnostic collection check boxes if you do not require the diagnostic feature.

- **Enable diagnostic events**: Enables and allows Oracle to collect and send fault notifications about critical, warning, and information events for you.
- **Enable health monitoring**: This diagnostics collection for Oracle Cloud operations viewing is not available for the Base Database Service.
- **Enable incident logs and trace collection**: Enables and allows Oracle to receive event notifications and collect incident logs and traces for fault diagnosis and issue resolution.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> You are opting in with the understanding that the list of events and log files can change in the future. You can opt out of this feature at any time.

Advanced options

Expand **Advanced options** to provide the advanced options for this resource.

Management

Expand **Management** to provide the following details:

- **Fault domain**: The fault domain(s) in which the DB system resides. You can select which fault domain to use for your DB system. For multi-node RAC DB systems, you can specify which two fault domains to use. Oracle recommends that you place each node of a multi-node RAC DB system in a different fault domain. For more information about fault domains, see About Regions and Availability Domains.
- **Time zone**: The default time zone for the DB system is UTC, but you can specify a different time zone. The time zone options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system. For more information, see DB System Time Zone. The following options are available:
  - **UTC:** configures your DB system to use coordinated universal time.
  - **Browser-detected:** The console displays the time zone detected by your browser for this option.
  - **Select another time zone:** To manually specify a time zone, first make a choice using the **Region or country selector** to select a geographic region, then use the **Time zone** selector to select your required time zone.

> **Tip**
>
> If you want to set a time zone other than UTC or the browser-detected time zone, and if you do not see the time zone you want, try selecting "Miscellaneous" in the **Region or country** list.

Security

Expand **Security** to provide security details. Optionally, you can specify one or more security attributes to configure Zero Trust Packet Routing (ZPR) for the DB system.

- Select the **Namespace** in which the required security attribute is available.
- Select the **Key** and **Value** of the required security attribute.
- Select **Add security attribute**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> - Administrators must set up security attribute namespaces and security attributes in a tenancy before users can apply security attributes to the DB systems.
> - A security attribute is effective only with appropriate policies. If a security attribute without any policies is added, all access will be denied by default, even if allowed in the Security List or NSGs.
> - If you use the security attribute, the security rules for the DB system will be a union of the rules in the security attributes along with any rules in the security list and the NSGs.
> - You may need to use more than one security attribute. If you're not sure, contact your network administrator.
> - A maximum of 3 security attributes can be specified for a DB system.

For more information about:

- Add, edit, or remove a security attribute after a DB system is created, see *Manage Security Attributes for the DB System*.
- Defining appropriate ZPR policies, see *Zero Trust Packet Routing*.
- Creating a security attribute, see Creating a Security Attribute.

Tags

Expand **Tags** to provide tag details for this resource.

You can add free-form tags or defined tags to this resource. You must have permission to use the tag namespace for defined tags. For information about using tags to manage your OCI resources, see Resource Tags.

> **Tip**
>
> Tags can be applied later or at any time to a resource.

Standby database

- **Database image:** *Optional*. By default, the latest database software image as the source database is used. You can specify what Oracle Database version is used for the database. You can mix database versions on the DB system, but not editions.
  - Select **Change database image** to choose a custom database software image that you or someone in your organization has created in your tenancy.
  - Use the **Compartment** filter to select the custom database software images from a specific compartment.
  - Select a database image from the table of available images for the Oracle Database version you selected.
  - Select the **Select** button.
- **Password**: Enter a password for the administrator. The password must meet the following criteria:
  - A strong password for SYS, SYSTEM, TDE wallet, and PDB administrator.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- The password must be 9 to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters.
- The special characters must be _, #, or -.
- The password must not contain the user name (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reversed order and regardless of casing.

Enable

Select **Add**.

When you create Data Guard, the details for a database and its peer display their respective roles as **Primary** or **Standby**.

If setting up Data Guard fails for any reason, we recommend contacting Oracle Support to resolve the issue. The billing for the DB system starts as soon as it is created and is in **Available** status.

Perform Database Switchover and Failover

You initiate a switchover or failover operation on a standby database that is a member of the Data Guard configuration.

Database Switchover

A switchover reverses the primary and standby database roles. Each database continues to participate in Oracle Data Guard in its new role. A switchover ensures no data loss. You can use a switchover before you perform planned maintenance on the primary database. Performing planned maintenance on a DB system with Data Guard is typically done by switching the primary database to the standby role, performing maintenance on the standby database, and then switching it back to the primary role.

Procedure

Perform the following steps to perform a switchover using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, select the **Data Guard Group** tab.
4. For the standby database on which you want to perform a switchover, from the **Actions** menu, select **Switchover**.
5. In the **Switchover database** panel, enter the database admin password, and then select **Switchover**.

This database should now assume the role of the primary, and the old primary should assume the role of the standby.

Database Failover

A failover transitions the standby database into the primary role after the existing primary database fails or becomes unreachable. A failover might result in some data loss when you use Maximum Performance protection mode.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Procedure

Perform the following steps to perform a failover using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, select the **Data Guard Group** tab.
4. For the standby database on which you want to perform a failover, from the **Actions** menu, select **Failover**.
5. In the **Failover database** panel, enter the database admin password, and then select **Failover**.

This database should now assume the role of the primary, and the old primary's role should display as **Disabled standby**.

Edit the Oracle Data Guard Association

This article provides the details and procedure to edit Oracle Data Guard.

Procedure

Perform the following steps on the primary database to edit Oracle Data Guard configuration using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, select the **Data Guard Group** tab.
4. For the database you want to edit, from the **Actions** menu, select **Edit**.
5. In the **Edit Data Guard Group** panel, configure the following:
   - **Data Guard type**: Select **Active Data Guard** or **Data Guard**. Active Data Guard provides additional features including: Real-Time Query and DML Offload, Automatic Block Repair, Standby Block Change Tracking, Far Sync, Global Data Services, and Application Continuity.
     - The Active Data Guard requires an Oracle Active Data Guard license.
     - For more information, see Introduction to Oracle Data Guard.
   - **Protection mode**: The protection mode can be **Maximum Performance** or **Maximum Availability**. For more information, see Oracle Data Guard Protection Modes.
   - **Transport type**: The redo transport type used for this Data Guard configuration. For more information, see Managing Redo Transport Services for Data Protection Modes.
   - **Database admin password**: Enter the ADMIN password for the database.
6. Select **Edit**.

Reinstate a Database

This article provides the details and procedure to reinstate a database.

Copyright © 2022, 2026, Oracle and/or its affiliates.

The reinstate moves a database into the standby role in an Oracle Data Guard configuration. You can use the `reinstate` command to return a failed database into service after correcting the cause of failure.

After you fail over a primary database to its standby, the standby assumes the primary role, and the old primary is identified as a disabled standby. After you correct the cause of failure, you can reinstate the failed database as a functioning standby for the current primary database.

> **Note**
>
> Before you can reinstate a 12.2 database, you must perform some steps on the database host to stop the database or start it in `MOUNT` mode.
> Set your `ORACLE_UNQNAME` environment variable to the value of the Database Unique Name (as seen in the Console), and then run these commands:
>
> ```
> srvctl stop database -d db-unique-name -o abort
> srvctl start database -d db-unique-name -o mount
> ```

Procedure

Perform the following steps on the current primary database from where we can reinstate the failed database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, select the **Data Guard Group** tab.
4. For the database you want to reinstate, from the **Actions** menu, select **Reinstate**.
5. In the **Reinstate database** panel, enter the database admin password, and then select **Ok**.

This database should now be reinstated as the standby.

Terminate an Oracle Data Guard Association on a DB System

This article provides the details and procedure to terminate Data Guard configuration on a DB system.

You can remove a Data Guard configuration by terminating all the standby databases.

You must terminate the DB system with the standby database before terminating the DB system with the primary database. Terminating the DB system with the standby database automatically terminates the Data Guard configuration. If you try terminating a DB system that has the primary before terminating the standby, the terminate operation will not complete.

Alternatively, you can switch over the primary database to the standby role and then terminate it.

For more instructions on terminating a DB system, see *Terminate a DB System*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Procedure

Perform the following steps on the standby database to remove a Data Guard configuration by terminating the standby database using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, from the **More actions** menu, select **Terminate** and confirm.

The DB system's icon indicates '**Terminating**'.

Use Oracle Data Guard with the Database CLI

This article explains how to use the database CLI to set up Data Guard with Fast-Start Failover (FSFO) in Oracle Cloud Infrastructure. The topics in this article explain how to prepare the primary and standby databases, and then configure Data Guard to transmit redo data from the primary database and apply it to the standby database.

Oracle recommends that you use the Console instead of the database CLI to set up and work with Data Guard in Oracle Cloud Infrastructure.

> **Note**
>
> This article assumes that you are familiar with Data Guard and FSFO. To learn more about them, see *Use Oracle Data Guard on a DB System*.

Prerequisites

To perform the procedures in this topic, you'll need the following information for the primary and standby databases.

- db_name (or oracle_sid)
- db_unique_name
- oracle home directory (or database home)

Finding the Database Information

After you've launched the primary and standby DB systems and created databases as described later in this topic, you can use the CLI on those systems to find the needed database information.

1. SSH to the DB System.

   ```
   ssh -i <private_key_path> opc@<db_system_ip_address>
   ```

2. Log in as `opc` and then `sudo` to the `root` user. Use `sudo su -` with a hyphen to invoke the root user's profile, which will set the `PATH` to the dbcli directory (`/opt/oracle/dcs/bin`).

   ```
   sudo su -
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

3. To find the db_name (or oracle_sid) and db_uniqueName, run the `dbcli list-databases -j` command.

   ```
   dbcli list-databases -j
   ```

   Output:

   ```
   [ {
     "id" : "80ad855a-5145-4f8f-a08f-406c5e4684ff",
     "name" : "dbtst",
     "dbName" : "dbtst",
     "databaseUniqueName" : "dbtst_phx1cs",
     "dbVersion" : "12.1.0.2",
     "dbHomeId" : "2efe7af7-0b70-4e9b-ba8b-71f11c6fe287",
     "instanceOnly" : false,
      .
      .
      .
   ```

4. To find the oracle home directory (or database home), run the `dbcli list-dbhomes` command. If there are multiple database homes on the DB system, use the one that matches the "dbHomeId" in the `dbcli list-databases -j` command output shown above.

   ```
   dbcli list-dbhomes
   ```

   Output:

   ```
   ID                                       Name                 DB
   Version                                  Home
   Location                                 Status
   ---------------------------------------- --------------------
   ----------------------------------------
   ---------------------------------------------- ----------
   2efe7af7-0b70-4e9b-ba8b-71f11c6fe287     OraDB12102_home1
   12.1.0.2.160719 (23739960, 23144544)     /u01/app/oracle/product/12.1.0.2/
   dbhome_1     Configured
   33ae99fe-5413-4392-88da-997f3cd24c0f     OraDB11204_home1
   11.2.0.4.160719 (23054319, 23054359)     /u01/app/oracle/product/11.2.0.4/
   dbhome_1     Configured
   ```

Create a Primary DB System

If you don't already have a primary DB system, create one as described in *Overview of Creating a DB System*. The DB system will include an initial database. You can create additional databases by using the Database Commands available on the DB system.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Create a Standby DB System

> **Note**
>
> The standby database must have the same db_name as the primary database, but it must have a different db_unique_name. If you use the same database name for the standby and primary, you will have to delete the database from the standby DB system by using the `dbcli delete-database` command before you can run the `dbcli create-database` command described below. Deleting and creating the database will take several minutes to complete. The `dbcli` commands must be run as the root user.

1. Create a standby DB system as described in *Overview of Creating a DB System* and wait for the DB system to finish provisioning and become available.
   You can create the standby DB system in a different availability domain from the primary DB system for availability and disaster recovery purposes (this is strongly recommended). You can create the standby DB system in the primary DB system's cloud network so that both systems are in a single, routable network.
2. SSH to the DB System.

   ```
   ssh -i <private_key_path> opc@<db_system_ip_address>
   ```

3. Log in as `opc` and then `sudo` to the `root` user. Use `sudo su -` with a hyphen to invoke the root user's profile, which will set the `PATH` to the dbcli directory (`/opt/oracle/dcs/bin`).

   ```
   sudo su -
   ```

4. The DB system will include an initial database, but you'll need to create a standby database by using the `dbcli create-database` command with the `--instanceonly` parameter. This parameter creates only the database storage structure and starts the database in nomount mode (no other database files are created).
   When using `--instanceonly`, both the `--dbname` and `--adminpassword` parameters are required and they should match the dbname and admin password of the primary database to avoid confusion.

   The following sample command prompts for the admin password and then creates a storage structure for a database named dbname.

   ```
   dbcli create-database --dbname <same as primary dbname>;--
   databaseUniqueName <different from primary uniquename>;--instanceonly --
   adminpassword
   ```

   If you are using pluggable databases, also specify the `--cdb` parameter.

   For complete command syntax, see Database Commands.
5. Wait a few minutes for the `dbcli create-database` command to create the standby database.
   You can use the `dbcli list-jobs` command to verify that the creation job ran successfully, and then the `dbcli list-databases` command verify that the database is configured.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Prepare the Primary DB System

To prepare the primary DB system, you'll need to configure static listeners, update tnsnames.ora, and configure some database settings and parameters.

Configuring the Static Listeners

Create static listeners to be used by RMAN and Data Guard Broker.

1. SSH to the primary DB system, log in as the `opc` or `root` user, and `sudo` to the grid OS user.

   ```
   sudo su - grid
   ```

2. Edit `/u01/app/<version>/grid/network/admin/listener.ora` and add the following content to it. The first static listener shown here is optional. The second `DGMGRL` static listener is optional for version 12.1 or later databases, but required for version 11.2 databases.

   ```
   SID_LIST_LISTENER=
     (SID_LIST=
       (SID_DESC=
       (SDU=65535)
       (GLOBAL_DBNAME = <primary_db_unique_name>.<primary_db_domain>)
       (SID_NAME = <primary_oracle_sid>)
       (ORACLE_HOME=<oracle_home_directory>)
       (ENVS="TNS_ADMIN=<oracle_home_directory>/network/admin")
       )
       (SID_DESC=
       (SDU=65535)
       (GLOBAL_DBNAME = <primary_db_unique_name>_DGMGRL.<primary_db_domain>)
       (SID_NAME = <primary_oracle_sid>)
       (ORACLE_HOME=<oracle_home_directory>)
       (ENVS="TNS_ADMIN=<oracle_home_directory>/network/admin")
       )
     )
   ```

3. Save your changes and then restart the listener.

   ```
   srvctl stop listener
   srvctl start listener
   ```

Adding Net Service Names to tnsnames.ora

As the oracle user, edit `$ORACLE_HOME/network/admin/tnsnames.ora` and add the standby database net service name to it.

```
<standby db_unique_name> =
  (DESCRIPTION =
    (SDU=65535)
    (ADDRESS = (PROTOCOL = TCP)(HOST = <standby_server>.<domain>) (PORT =
1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = <standby db_unique_name>.<standby db_domain>)
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
    )
  )
```

The sample above assumes that name resolution is working and that the <*standby_server*>.<*domain*> is resolvable at the primary database. You can also use the private IP address of the standby server if the IP addresses are routable within a single cloud network (VCN).

Configuring Primary Database Parameters

> **Tip**
>
> If the primary and standby hosts have different directory structures, you might need to set additional parameters that are not discussed here, such as the `log_file_name_convert` parameter. See the RMAN documentation for more information about how to create standbys for hosts with different directory structures.

1. As the oracle user, enable automatic standby file management.

   ```
   alter system set standby_file_management=AUTO;
   ```

2. Identify the Broker configuration file names and locations. The commands used for this depend on the type of database storage. If you're not sure of the database storage type, use the Database Commands on the DB system.
   For ACFS database storage, use the following commands to set the Broker configuration files.

   ```
   alter system set dg_broker_config_file1='/u02/app/oracle/oradata/<Primary
   db_unique_name>/dbs/dr1<Primary db_unique_name>.dat';
   alter system set dg_broker_config_file2='/u02/app/oracle/oradata/<Primary
   db_unique_name>/dbs/dr2<Primary db_unique_name>.dat';
   ```

   For ASM database storage, use the following commands to set the Broker configuration files.

   ```
   alter system set dg_broker_config_file1='+DATA/<Primary db_unique_name>/
   dr1<db_unique_name>.dat';
   alter system set dg_broker_config_file2='+DATA/<Primary db_unique_name>/
   dr2<db_unique_name>.dat';
   ```

3. Enable Broker DMON process for the database.

   ```
   alter system set dg_broker_start=true;
   ```

4. Force database logging for all database transactions.

   ```
   alter database force logging ;
   ```

5. Add Standby Redo Logs (SRLs), based on the Online Redo Logs (ORLs). On a newly launched DB system, there will be three ORLs of size 1073741824, so create four SRLs of the same size.

Copyright © 2022, 2026, Oracle and/or its affiliates.

You can use the query below to determine the number and size (in bytes) of the ORLs.

```
select group#, bytes from v$log;
```

Output:

```
    GROUP# BYTES
---------- ----------
         1 1073741824
         2 1073741824
         3 1073741824
```

All of the ORLs must be the same size.

The SRLs must be the same size as the ORLs, but there must be at least one more SRL than the ORLs. In the example above, there are three ORLs, so four SRLs are required. So specify the current redo logs plus one, and use the same size as the redo logs.

```
alter database add standby logfile thread 1 size <size>;
```

There should be only one member in the SRL group (by default, a DB system is created with only one member per SRL group). To ensure this, you can name the file with the following syntax.

```
alter database add standby logfile thread 1 group 4 (<logfile name with
full path>) size 1073741824, group 5(<logfile name with full path>) size
1073741824 ...
```

For ASM/OMF configurations, the above command uses the diskgroup instead of <*logfile name with full path*>.

```
alter database add standby logfile thread 1 group 4 (+RECO) size
1073741824, group 5(+RECO) size 1073741824 ...
```

> **Tip**
>
> ORLs and SRLs should be sized so that log switches do not occur more frequently than every 10 minutes. This requires knowledge of the application and may need to be adjusted after deployment. For more information, see Use Standby Redo Logs and Configure Size Appropriately.

6. Verify that you created the correct number of SRLs.

```
select group#, bytes from v$standby_log;
```

7. Make sure the database is in ARCHIVELOG mode.

```
archive log list
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

8. Enable database FLASHBACK. The minimum recommended value for db_flashback_retention_target is 120 minutes.

```
alter database flashback on ;
alter system set db_flashback_retention_target=120;
```

9. Perform a single switch redo log to activate archiving if database is newly created. (At least one log must be archived prior to running the RMAN duplicate.)

```
alter system switch logfile;
```

Prepare the Standby Database

Before you prepare the standby database, make sure the database home on the standby is the same version as on the primary. (If the primary and standby databases are both newly created with the same database version, the database homes will be the same.) If it is not, create a database home that is the same version. You can use the Dbhome Commands to verify the versions and create a new database home as needed.

To prepare the standby DB system, you'll need to configure static listeners, update tnsnames.ora, configure TDE Wallet, create a temporary password file, verify connectivity, run RMAN DUPLICATE, enable FLASHBACK, and then create the database service.

Configuring the Static Listeners

Create static listeners to be used by RMAN and Data Guard Broker.

1. SSH to the standby DB system, log in as the `opc` or `root` user, and `sudo` to the grid OS user.

```
sudo su - grid
```

2. Append the following content to `/u01/app/<db_version>/grid/network/admin/listener.ora`.
The first static listener shown below is required for RMAN DUPLICATE. The second DGMGRL static listener is optional for database versions 12.2.0.1 and 12.1.0.2, but required for database version 11.2.0.4.

```
SID_LIST_LISTENER=
  (SID_LIST=
    (SID_DESC=
    (SDU=65535)
    (GLOBAL_DBNAME = <standby db_unique_name>.<standby db_domain>)
    (SID_NAME = <standby oracle_sid>)
    (ORACLE_HOME=<oracle home directory>)
    (ENVS="TNS_ADMIN=<oracle home directory>/network/admin")
    )
    (SID_DESC=
    (SDU=65535)
    (GLOBAL_DBNAME = <standby db_unique_name>_DGMGRL.<standby db_domain>)
    (SID_NAME = <standby oracle_sid>)
    (ORACLE_HOME=<oracle home directory>)
    (ENVS="TNS_ADMIN=<oracle home directory>/network/admin")
    )
  )
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

3. Restart the listener.

```
srvctl stop listener
srvctl start listener
```

4. Verify that the static listeners are available. The sample output below is for database version 12.1.0.2. Note that the `...status UNKNOWN` messages are expected at this point.

```
lsnrctl status
```

Output:

```
LSNRCTL for Linux: Version 12.1.0.2.0 - Production on 29-SEP-2016 21:09:25

Copyright (c) 1991, 2014, Oracle.  All rights reserved.

Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=IPC)(KEY=LISTENER)))
STATUS of the LISTENER
------------------------
Alias                     LISTENER
Version                   TNSLSNR for Linux: Version 12.1.0.2.0 -
Production
Start Date                29-SEP-2016 21:09:19
Uptime                    0 days 0 hr. 0 min. 5 sec
Trace Level               off
Security                  ON: Local OS Authentication
SNMP                      OFF
Listener Parameter File   /u01/app/12.1.0.2/grid/network/admin/listener.ora
Listener Log File         /u01/app/grid/diag/tnslsnr/dg2/listener/alert/
log.xml
Listening Endpoints Summary...
  (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=LISTENER)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=10.0.1.24)(PORT=1521)))
Services Summary...
Service "dg2_phx2hx.oratst.org" has 1 instance(s).
  Instance "dg2", status UNKNOWN, has 1 handler(s) for this service...
Service "dg2_phx2hx_DGMGRL.oratst.org" has 1 instance(s).
  Instance "dg2", status UNKNOWN, has 1 handler(s) for this service...
The command completed successfully
```

Adding Net Service Names to tnsnames.ora

As the oracle user, add the standby database net service name to `$ORACLE_HOME/network/admin/tnsnames.ora`. $ORACLE_HOME is the database home where the standby database is running.

```
<Primary db_unique_name> =
  (DESCRIPTION =
    (SDU=65535)
    (ADDRESS = (PROTOCOL = TCP)(HOST = <primary_server>.<domain>) (PORT =
1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = <primary db_unique_name).<primary db_domain>)
    )
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
  )

<Standby db_unique_name> =
  (DESCRIPTION =
    (SDU=65535)
    (ADDRESS = (PROTOCOL = TCP)(HOST = <standby_server>.<domain>) (PORT =
1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = <standby db_unique_name>.<db_domain>)
    )
  )
```

Copying the TDE Wallets to the Standby System

Copy the TDE wallet files from the primary DB system to standby DB system using SCP. The following sample command assumes the SCP command is being run by the oracle OS user and that the private key for oracle has been created and exists on the host where SCP is being run.

```
scp -i <private key> primary_server:/opt/oracle/dcs/commonstore/wallets/tde/
<primary db_unique_name>/* standby_server:/opt/oracle/dcs/commonstore/
wallets/tde/<standby db_unique_name>
```

Setting Up the Standby System Configuration

As the oracle user, create the following directory for database version 11.2.0.4. This step is optional for version 12.2.0.1 and version 12.1.0.2.

```
mkdir  -pv /u03/app/oracle/redo/<standby db_unique_name uppercase>/controlfile
```

Creating the Audit File Destination

As the oracle user, create the following directory to use as the audit file destination.

```
mkdir -p /u01/app/oracle/admin/<db_name>/adump
```

Otherwise, the RMAN duplicate command used later will fail.

Creating a Temporary Password File

As the oracle user, create a temporary password file.

```
orapwd file=$ORACLE_HOME/dbs/orapw<standby oracle_sid>    password=<admin
password for primary> entries=5
```

The password **must** be the same as the admin password of the primary database. Otherwise, the `RMAN` duplicate step below will fail with: `RMAN-05614: Passwords for target and auxiliary connections must be the same when using active duplicate.`

Copyright © 2022, 2026, Oracle and/or its affiliates.

Verifying the Standby Database is Available

1. As the oracle user, set the environment variables.

   ```
   . oraenv
   ```

2. Replace `$ORACLE_HOME/dbs/init<standby sid_name>.ora` with the following content:

   ```
   db_name=<Primary db_name>
   db_unique_name=<standby db_unique_name>
   db_domain=<standby db_domain>
   ```

3. Remove the spfile from the standby.

   ```
   /u02/app/oracle/oradata/<standby db_unique_name>/dbs/spfile$ORACLE_SID.ora
   ```

   The database needs to be started in nomount mode with no spfile specified, but the original init file contains an spfile parameter which will prevent the `RMAN` duplicate step from working.

4. Set the `ORACLE_UNQNAME` environment variable to point to your `DB_UNIQUE_NAME`.

   ```
   export ORACLE_UNQNAME =db_unique_name
   ```

   > **Note**
   >
   > If you do not perform this step, the wallet will not be opened, and running the `RMAN DUPLICATE` command in the subsequent step will fail.

5. The `dbcli create-database --instanceonly` command used earlier opens the standby database as a primary in read/write mode, so the database needs to be brought down before proceeding to the nomount step below.

   ```
   sqlplus / as sysdba

   shutdown immediate
   ```

6. Start the database in nomount mode.

   ```
   startup nomount
   ```

Verifying the Database Connections

Verify the connection between the primary and standby databases.

1. Make sure that the listener port 1521 is open in the security list(s) used for the primary and standby DB systems. For more information, see *Update the Security List for the DB System*.
2. From the primary database, connect to standby database.

   ```
   sqlplus sys/<password>@<standby net service name> as sysdba
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

3. From standby database, connect to primary database.

```
sqlplus sys/<password>@<primary net service name> as sysdba
```

Running the RMAN DUPLICATE Command

Run the RMAN DUPLICATE command on the standby DB system, as the oracle user.

If the primary database is large, you can allocate additional channels to improve performance. For a newly installed database, one channel typically runs the database duplication in a couple of minutes.

Make sure that there are no errors generated by the RMAN DUPLICATE command. If errors occur, restart the database using the `init.ora` file (not spfile) in case it is generated under `$ORACLE_HOME/dbs` as part of RMAN DUPLICATE.

In the following examples, use lowercase for the <*Standby db_unique_name*> unless otherwise specified.

For ACFS storage layout, run the following commands.

```
rman target sys/<password>@<primary alias>  auxiliary sys/<password>@<standby
alias> log=rman.out

run { allocate channel prim1 type disk;
    allocate auxiliary channel sby type disk;
    duplicate target database for standby from active database
    dorecover
    spfile
    parameter_value_convert '/<Primary db_unique_name>/','/<Standby
db_unique_name>/','/<Primary db_unique_name uppercase>/','/<Standby
db_unique_name uppercase >/'
    set db_unique_name='<Standby db_unique_name>'
    set db_create_file_dest='/u02/app/oracle/oradata/<Standby db_unique_name>'
    set dg_broker_config_file1='/u02/app/oracle/oradata/<Standby
db_unique_name>/dbs/dr1<Standby db_unique_name>.dat'
    set dg_broker_config_file2='/u02/app/oracle/oradata/<Standby
db_unique_name>/dbs/dr2<Standby db_unique_name>.dat'
    set dispatchers ='(PROTOCOL=TCP) (SERVICE=<Standby db_unique_name>XDB)'
    set instance_name='<Standby db_unique_name>'
    ;
}
```

For ASM storage layout, run the following commands.

```
rman target sys/<password>@<primary alias>  auxiliary sys/<password>@<standby
alias> log=rman.out

run {
    allocate channel prim1 type disk;
    allocate auxiliary channel sby type disk;
    duplicate target database for standby from active database
    dorecover
    spfile
    parameter_value_convert '/<Primary db_unique_name>/','/<Standby
db_unique_name>/','/<Primary db_unique_name uppercase>/','/<Standby
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
db_unique_name uppercase>/'
    set db_unique_name='<Standby db_unique_name>'
    set dg_broker_config_file1='+DATA/<Standby db_unique_name>/dr1<Standby
db_unique_name>.dat'
    set dg_broker_config_file2='+DATA/<Standby db_unique_name>/dr2<Standby
db_unique_name>.dat'
    set dispatchers ='(PROTOCOL=TCP) (SERVICE=<Standby db_unique_name>XDB)'
    set instance_name='<Standby db_unique_name>'
    ;
}
```

Enabling Database FLASHBACK

1. As a Data Guard best practice, enable flashback and set `db_flashback_retention_target` to at least 120 minutes on both the primary and standby databases.

   ```
   alter database flashback on;
   alter system set db_flashback_retention_target=120;
   ```

2. Verify that the standby database is created properly.

   ```
   select FORCE_LOGGING, FLASHBACK_ON, OPEN_MODE,
   DATABASE_ROLE,SWITCHOVER_STATUS, DATAGUARD_BROKER, PROTECTION_MODE from
   v$database ;
   ```

Creating a Database Service

Oracle recommends creating a database service for the standby database by using srvctl.

For ACFS storage layout.

1. Create a shared directory and copy the spfile file to it.

   ```
   mkdir -pv /u02/app/oracle/oradata/<Standby db_unique_name>/dbs
   cp $ORACLE_HOME/dbs/spfile<standby oracle_sid>.ora /u02/app/oracle/oradata/
   <Standby db_unique_name>/dbs
   ```

2. Stop and remove the existing database service.

   ```
   srvctl stop database
       -d <standby db_unique_name>

   srvctl remove database
       -d <standby db_unique_name>
   ```

3. Create the database service.

   ```
   srvctl add database
       -d <standby db_unique_name>
       -n <standby db_name>
       -o $ORACLE_HOME
       -c SINGLE
       -p '/u02/app/oracle/oradata/<standby db_unique_name>/dbs/
   spfile<standby db_name>.ora'
       -x <standby hostname>
       -s "READ ONLY"
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
    -r PHYSICAL_STANDBY
    -i <db_name>

srvctl setenv database
    -d <standby db_unique_name>
    -t "ORACLE_UNQNAME=<standby db_unique_name>"

srvctl config database
    -d <standby db_unique_name>
```

4. Start the database service.

```
srvctl start database
    -d <standby db_unique_name>
```

5. Clean up the files from `$ORACLE_HOME/dbs`.

```
rm $ORACLE_HOME/dbs/spfile<standby oracle_sid>.ora
rm $ORACLE_HOME/dbs/init<standby oracle_sid>.ora
```

6. Create the `$ORACLE_HOME/dbs/init<*standby oracle_sid*>.ora` file to reference the new location of the spfile file.

```
SPFILE='/u02/app/oracle/oradata/<standby db_unique_name>/dbs/
spfile<standby db_name>.ora'
```

7. Stop the standby database and then start it by using `srvctl`.

```
srvctl stop database
    -d <standby db_unique_name>

srvctl start database
    -d <standby db_unique_name>
```

For ASM storage layout.

1. Consider generating the spfile file under +DATA.

```
create pfile='init<standby oracle_sid>.ora' from spfile ;
create spfile='+DATA' from pfile='init<standby oracle_sid>.ora' ;
```

2. Stop and remove the existing database service.

```
srvctl stop database
    -d <standby db_unique_name>

srvctl remove database
    -d <standby db_unique_name>
```

3. Create the database service.

```
srvctl add database
    -d <standby db_unique_name>
    -n <standby db_name>
    -o $ORACLE_HOME
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
    -c SINGLE
    -p '+DATA/<standby db_unique_name>/PARAMETERFILE/spfile.xxx.xxxxxx'
    -x <standby hostname>
    -s "READ ONLY"
    -r PHYSICAL_STANDBY
    -i <db_name>

srvctl setenv database
    -d <standby db_unique_name>
    -t "ORACLE_UNQNAME=<standby db_unique_name>"

srvctl config database
    -d <standby db_unique_name>
```

4. Start the database service.

```
srvctl start database
    -d <standby db_unique_name>
```

5. Clean up the files from `$ORACLE_HOME/dbs`.

```
rm $ORACLE_HOME/dbs/init<standby oracle_sid>.ora
rm $ORACLE_HOME/dbs/spfile<standby oracle_sid>.ora
```

6. Create `$ORACLE_HOME/dbs/init<standby oracle_sid>.ora` file to reference the new location of the spfile file.

```
SPFILE='+DATA/<standby db_unique_name>/PARAMETERFILE/spfile.xxx.xxxxxx'
```

7. Stop the database and start the standby database by using srvctl.

```
srvctl start database -d <standby db_unique_name>
```

Configure Data Guard

Perform the following steps to complete the configuration of Data Guard and enable redo transport from the primary database and redo apply in the standby database.

1. Run the `dgmgrl` command line utility from either the primary or standby DB system and connect to the primary database using sys credentials.

```
connect sys/<sys password>@<primary tns alias>
```

2. Create the Data Guard configuration and identify for the primary and standby databases in the `dgmgrl` command line utility.

```
create configuration mystby as primary database is <primary
db_unique_name> connect identifier is <primary tns alias>;
add database  <standby db_unique_name> as connect identifier is <standby
tns alias>  maintained as physical;
```

3. Enable Data Guard configuration the `dgmgrl` command line utility.

```
enable configuration;
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

4. Verify that Data Guard setup was done properly. Run the following SQL in **both** the primary and standby databases in the SQL prompt.

```
select FORCE_LOGGING, FLASHBACK_ON, OPEN_MODE, DATABASE_ROLE,
SWITCHOVER_STATUS, DATAGUARD_BROKER, PROTECTION_MODE from v$database;
```

5. Verify that Data Guard processes are initiated in the standby database.

```
select PROCESS,PID,DELAY_MINS from V$MANAGED_STANDBY;
```

6. Verify parameter configuration on primary and standby.

```
show parameter log_archive_dest_
show parameter log_archive_config
show parameter fal_server
show parameter log_archive_format
```

7. Verify that the Data Guard configuration is working in the `dgmgrl` command line utility. Specifically, make sure redo shipping and redo apply are working and that the standby is not unreasonably lagging behind the primary.

```
show configuration verbose
show database verbose <standby db_unique_name>
show database verbose <primary db_unique_name>
```

Any discrepancies, errors, or warnings should be resolved. You can also run a transaction on the primary and verify that it's visible in the standby.

8. Verify that the Data Guard configuration is functioning as expected by performing switchover and failover in both directions. Run `show configuration` after each operation and make sure there are no errors or warnings in the `dgmgrl` command line utility.

> ⚠ **Caution**
>
> This step is optional, based on your discretion. If for any reason the configuration is not valid, the switchover and/or failover will fail and it might be difficult or impossible to start the primary database. A recovery of the primary might be required, which will affect availability.

```
switchover to <standby db_unique_name>
switchover to <primary db_unique_name>
#connect to standby before failover:

connect sys/<sys password>@<standby db_unique_name>
failover to <standby db_unique_name>
reinstate database <primary db_unique_name>
#connect to primary before failover:

connect sys/<sys password>@<primary db_unique_name>
failover to <primary db_unique_name>
reinstate database <standby db_unique_name>
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Configure Observer (Optional)

The best practice for high availability and durability is to run the primary, standby, and observer in separate availaility domains. The observer determines whether or not to failover to a specific target standby database. The server used for observer requires the Oracle Client Administrator software, which includes the Oracle SQL NET and Broker. Execute following commands in `dgmgrl` command line utility.

1. Configure TNS alias names for both the primary and standby databases as described previously, and verify the connection to both databases.
2. Change protection mode to either maxavailability or maxperformance (maxprotection is not supported for FSFO).
   To enable maxavailability:

   ```
   edit database <standby db_unique_name> set property 'logXptMode'='SYNC';
   edit database <primary db_unique_name> set property 'logXptMode'='SYNC';
   edit configuration set protection mode as maxavailability;
   ```

   To enable maxperformance:

   ```
   edit configuration set protection mode as maxperformance;
   edit database <standby db_unique_name> set property 'logXptMode'='ASYNC';
   edit database <primary db_unique_name> set property 'logXptMode'='ASYNC';
   ```

   For maxperformance, the `FastStartFailoverLaglimit` property limits the maximum amount of permitted data loss to 30 seconds by default.
3. The following properties should also be considered. Run `show configuration verbose` to see their current values.
   - FastStartFailoverPmyShutdown
   - FastStartFailoverThreshold
   - FastStartFailoverTarget
   - FastStartFailoverAutoReinstate

   (Running `show configuration` will result in the following error until the observer is started: `Warning : ORA-16819: fast-start failover observer not started`.)
4. Enable fast-start failover from Broker:

   ```
   Enable fast_start failover
   ```

5. Verify the fast-start failover and associated settings.

   ```
   show fast_start failover
   ```

6. Start the observer from Broker (it will run in the foreground, but can also be run in the background).

   ```
   start observer
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

7. Verify fast-start failover is enabled and without errors or warnings.

```
show configuration verbose
```

8. Always test failover in both directions to ensure that everything is working as expected. Verify that FSFO is running properly by performing a shutdown abort of the primary database.
The observer should start the failover to the standby database. If protection mode is set to maxprotection, some loss of data can occur, based on the FastStartFailoverLaglimit value.

Encryption Keys

Manage Encryption Keys

This article provides the details about encryption and encryption keys.

Oracle Base Database Service encrypts data stored in tables and tablespaces using Transparent Data Encryption (TDE).

Transparent Data Encryption
The Base Database Service uses TDE to encrypt and decrypt all user-created tablespaces.

Encryption Keys
You can choose to encrypt the database using your own encryption keys ("customer-managed keys") or use Oracle-managed keys. By default, the Base Database Service uses Oracle-managed keys. The customer-managed key is stored in the OCI Vault, which is external to the database host.

OCI Vault Key
In the OCI Vault, the Encryption keys are logical entities that contain one or more key versions that are used for encryption and decryption. These key versions can be auto-generated by OCI Vault or imported from an external source (Bring-Your-Own-Key).

For more information, see Introduction to Transparent Data Encryption and OCI Vault Key Management.

Required IAM Policy

If you want to use your own encryption keys to encrypt a database, then you must create a dynamic group and assign specific policies to the group for customer-managed encryption keys. See Managing Dynamic Groups and *Let security admins manage vaults, keys, and secrets* topic in Common Policies.

General Information

While creating a new DB system, a key will be assigned to both the container database and pluggable database.

The key version, if provided, will only be used for the container database and not for its pluggable database. The pluggable database will be assigned an automatically generated new key version. Specific key versions cannot be assigned to the pluggable databases during creation.

The pluggable database will always use the same key as the container database, but with the same or a different key version.

Copyright © 2022, 2026, Oracle and/or its affiliates.

You can specify any key version, including the latest version of the selected key.

By default, the database is configured using Oracle-managed keys. However, you can choose to configure it using customer-managed keys.

Rotate Encryption Key

The rotate encryption key operation generates a new key version for the same key.

You can perform any number of key rotations. Periodically rotating keys limits the amount of data encrypted or signed by one key version. The history of retired keys is also maintained, which enables you to rotate the key and still be able to decrypt data that was encrypted by an earlier key.

The rotate key at container database and pluggable database levels works independently of each other. The rotate key operation on a container database will not rotate keys in the pluggable databases. Similarly, rotating keys in one pluggable database will not rotate keys in other pluggable databases or its container database.

To ensure you are using the latest version, rotate keys from the database details page on the OCI Console instead of the Vault service's Console page.

> **Note**
>
> Rotating the encryption keys is not available for the databases that use Oracle-managed encryption.

To rotate an encryption key using the OCI Console, see *Rotate Encryption Key for a Database* and *Rotate Encryption Key for a Pluggable Database*.

Assign Key Version

You can create and assign new key versions for both container databases and pluggable databases. Only the key version can be changed. The key cannot be changed.

To assign the key version using the OCI Console, see *Assign a New Key Version for a Database* and *Assign a New Key Version for a Pluggable Database*.

Change Key Management

You can switch from Oracle-managed keys to customer-managed keys on existing databases. However, switching from customer-managed keys to Oracle-managed keys is not supported.

When a key is changed for a container database, it is also automatically applied to a pluggable database. The key of a pluggable database cannot be changed independently. The pluggable database will always use the same key as that of the container database, but they can use the same or a different key version.

When switching to customer-managed keys, the container database and all its pluggable databases must be open, and all tablespaces must be in read/write mode.

To change the key management type using the OCI Console, see *Change Key Management Type for a Database*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Clone, Remote Clone, and Relocate Pluggable Database

The cloned database will use the same key version as the source database when cloning a DB system that uses customer-managed encryption keys.

The source and target databases must use the same key but can have a different key version. The remote cloning or relocating operation fails if the source and target databases use different keys.

The keys are rotated in the target key vault after remote cloning and relocation operations. So new key versions will be generated for the remote cloned or relocated pluggable database in the target database.

Rotate Encryption Key for a Database

This article provides the details and procedure to rotate encryption key for a database.

The rotate encryption key operation generates a new key version for the same key.

You can perform any number of key rotations. Periodically rotating keys limits the amount of data encrypted or signed by one key version. The history of retired keys is also maintained, which enables you to rotate the key and still be able to decrypt data that was encrypted by an earlier key.

The rotate key at container database and pluggable database levels works independently of each other. The rotate key operation on a container database will not rotate keys in the pluggable databases. Similarly, rotating keys in one pluggable database will not rotate keys in other pluggable databases or its container database.

To ensure you are using the latest version, rotate keys from the database details page on the OCI Console instead of the Vault service's Console page.

> **Note**
>
> Rotating the encryption keys is not available for the databases that use Oracle-managed encryption.

Procedure

Perform the following steps to rotate the encryption key for a database in a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, from the **Actions** menu, select **Manage encryption key**.
4. In the **Manage encryption key** panel, enter the following values:
   - Select **Rotate Encryption Key**.
   - Select **Update**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Assign a New Key Version for a Database

This article provides the details and procedure to assign a new key version for a database.

You can create and assign new key versions for both container databases and pluggable databases. Only the key version can be changed. The key cannot be changed.

Procedure

Perform the following steps to assign a new key version for a database in a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, from the **Actions** menu, select **Manage encryption key**.
4. In the **Manage encryption key** panel, enter the following values:
   - Select **Assign a new key version**.
   - Enter the OCID of the key version that you want to use.
   - Select **Update**.

Change Key Management Type for a Database

This article provides the details and procedure to change the key management type for a database.

You can switch from Oracle-managed keys to customer-managed keys on existing databases. However, switching from customer-managed keys to Oracle-managed keys is not supported.

When a key is changed for a container database, it is also automatically applied to a pluggable database. The key of a pluggable database cannot be changed independently. The pluggable database will always use the same key as that of the container database, but they can use the same or a different key version.

When switching to customer-managed keys, the container database and all its pluggable databases must be open, and all tablespaces must be in read/write mode.

Procedure

Perform the following steps to change key management for a database in a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, from the **Actions** menu, select **Manage encryption key**.
4. In the **Manage encryption key** panel, enter the following values:
   - Select **Change Key Management Type**.
   - Select a compartment where the key is located from the list.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Select a **Vault** from the list.
- Select a **Master encryption key** from the list.
- Optionally, to specify a key version other than the latest version of the selected key, switch on the **Choose the key version** toggle and enter the OCID of the key you want to use in the **Key version OCID** field.

> **Note**
>
> If you do not choose a version, the latest version of the key is used.

- Enter the **TDE wallet password** to access the current Oracle-managed key.
- Select **Update**.

Rotate Encryption Key for a Pluggable Database

This article provides the details and procedure to rotate encryption key for a pluggable database.

The rotate encryption key operation generates a new key version for the same key.

You can perform any number of key rotations. Periodically rotating keys limits the amount of data encrypted or signed by one key version. The history of retired keys is also maintained, which enables you to rotate the key and still be able to decrypt data that was encrypted by an earlier key.

The rotate key at container database and pluggable database levels works independently of each other. The rotate key operation on a container database will not rotate keys in the pluggable databases. Similarly, rotating keys in one pluggable database will not rotate keys in other pluggable databases or its container database.

To ensure you are using the latest version, rotate keys from the database details page on the OCI Console instead of the Vault service's Console page.

> **Note**
>
> Rotating the encryption keys is not available for the databases that use Oracle-managed encryption.

Procedure

Perform the following steps to rotate the encryption key for a pluggable database in a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the pluggable database details page, from the **Actions** menu, select **Manage encryption key**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

5. In the **Manage encryption key** panel, enter the following values:
   - Select **Rotate Encryption Key**.
   - Select **Update**.

Assign a New Key Version for a Pluggable Database

This article provides the details and procedure to assign a new key version for a pluggable database.

You can create and assign new key versions for both container databases and pluggable databases. Only the key version can be changed. The key cannot be changed.

Procedure

Perform the following steps to assign a new key version for a pluggable database in a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that contains the pluggable database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that contains the pluggable database that you want to work with.
3. On the **Pluggable Databases** tab, select the pluggable database that you want to work with.
4. On the pluggable database details page, from the **Actions** menu, select **Manage encryption key**.
5. In the **Manage encryption key** panel, enter the following values:
   - Select **Assign a new key version**.
   - Select **Update**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

6

Secure

Security Guide for Base Database Service

This article describes security for Base Database Service.

Security Overview

This topic provides an overview of the security in the Base Database Service. Oracle manages security for most components, while users are responsible for the security of some components.

The cloud service components are classified into user-accessible services and Oracle-managed infrastructure. User-accessible service refers to the components that users can access as part of their subscription to the Base Database Service. These are virtual machines and database services commonly called as DB systems and databases respectively. Oracle-managed infrastructure refers to the hardware that Oracle owns and operates to support user-accessible services. It consists of AMD or Intel-based database computing shapes.

Oracle will manage the security and access to the Oracle-managed infrastructure components. Users will manage the security and access to the user-accessible services that include access to DB system and database services, network access to the DB system, authentication to access the DB system, and authentication to access databases running on the DB systems. Oracle staff are not authorized to access user-accessible services.

Users access Oracle Databases running on DB systems via a layer 2 (tagged VLAN) connection from user equipment using standard Oracle Database connection methods, such as Oracle Net on port 1521. Users can use the standard Oracle Linux methods to connect to the DB system running the Oracle Databases, such as token-based SSH on port 22.

The Base Database Service employs multiple, independent, and mutually-reinforcing security controls to help organizations create a secure operating environment for their workloads and data. The Base Database Service provides the following security controls:

- Defense in Depth to Secure the Operating Environment
- Least Privilege for Services and Users
- Audit and Accountability of Events and Actions
- Automating Cloud Operations

Defense in Depth to Secure the Operating Environment

The Base Database Service provides several controls to maintain confidentiality, integrity, and accountability across the service. The Base Database Service promotes the principle of defense-in-depth as follows:

- The virtual machines for DB systems are built from the hardened operating system image based on Oracle Linux 7. It secures the core operating environment by restricting the installation image to only the required software packages, disabling unnecessary services, and implementing secure configuration parameters throughout the system.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Additional secure default configuration choices are implemented in the service instances in addition to inheriting all the strengths of the mature Oracle Linux platform. For example, all database tablespaces require transparent data encryption (TDE), strong password enforcement for initial database users and superusers, and enhanced audit and event rules.
- The Base Database Service also constitutes a complete deployment and service and is subject to industry-standard external audits such as PCI, HIPPA, and ISO27001. These external audit requirements impose additional value-added service features such as antivirus scanning, automated alerting for unexpected changes to the system, and vulnerability scans for all Oracle-managed infrastructure systems in the fleet.

Least Privilege for Services and Users

Oracle secure coding standards require the paradigm of least privilege. Ensuring that applications, services, and users have access to the capabilities that they need to perform their tasks is only one side of the least-privilege principle. It is equally important to ensure that access to unnecessary capabilities, services, and interfaces are limited. Base Database Service promotes the principle of least-privilege as follows:

- Each process and daemon must run as a normal, unprivileged user unless it can prove a requirement for a higher level of privilege. This helps contain any unforeseen issues or vulnerabilities to unprivileged user space and not compromise an entire system.
- This principle also applies to Oracle operations team members who use individual named accounts to access the infrastructure for maintenance or troubleshooting. Only when necessary will they use the audited access to higher levels of privilege to solve or resolve an issue. Most issues are resolved through automation, so we also employ least privilege by not permitting human operators to access a system unless the automation is unable to resolve the issue.

Audit and Accountability of Events and Actions

A system must be able to recognize and notify incidents as they occur. Similarly, when an incident cannot be averted, an organization must be able to identify its occurrence in order to take the appropriate actions. Base Database Service encourages audit and accountability in the following ways:

- Auditing and accountability ensure that both Oracle and users are aware of the activity done on the system and its time. These details not only ensure that we remain compliant with reporting requirements for external audits, but they can also assist in identifying the activity that led to unexpected behavior.
- Auditing capabilities are provided for all infrastructure components to ensure all actions are captured. Users can also configure auditing for their database and user domain (domU) configuration and may choose to integrate those with other enterprise auditing systems.
- Oracle does not access the user domU.

Automating Cloud Operations

By eliminating manual operations required to provision, patch, maintain, troubleshoot, and configure systems, the possibility for error is reduced and a secure configuration is ensured.

The Base Database Service is designed to be secure by automating all provisioning, configuration, and the majority of other operational tasks. By automating, it is possible to avoid missed configurations and ensure all necessary paths into the system are properly configured.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Security Features

This topic describes the security features available in the Base Database Service.

The Base Database Service provides the following security features:

- Hardened OS Image
- Minimized Attack Surface
- Additional Security Features Enabled
- Secure Access Methods
- Auditing and Logging

Hardened OS Image

- **Minimal package installation**: Only the necessary packages required to run an efficient system are installed. By installing a smaller set of packages, the attack surface of the operating system is reduced and the system remains more secure.
- **Secure configuration**: Many non-default configuration parameters are set during installation to enhance the security posture of the system and its content. For example, SSH is configured to only listen on certain network interfaces, sendmail is configured to only accept local host connections, and many other similar restrictions are implemented during installation.
- **Run only necessary services**: Any services that may be installed on the system but are not required for normal operation are disabled by default. For example, while NFS is a service often configured by users for various application purposes, it is disabled by default as it is not required for normal database operations. Users may choose to optionally configure services as per their requirements.

Minimized Attack Surface

As part of the hardened image, the attack surface is reduced by installing and running only the software required to deliver the service.

Additional Security Features Enabled

- Base Database Service is designed to be secure by default and provides a complete security stack, from network firewall control to access control security policies.
- FIPS, SE Linux, and STIG can be enabled additionally to improve security on systems using the `dbcli secure-dbsystem` CLI.
- The STIG tool is provided to increase compliance with DISA's Oracle Linux 7 STIG on each system node in provisioned systems.

Secure Access Methods

- Access database servers via SSH using strong cryptographic ciphers. Weak ciphers are disabled by default.
- Access databases via encrypted Oracle Net connections. By default, our services are available using encrypted channels, and a default configured Oracle Net client will use encrypted sessions.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Auditing and Logging

By default, auditing and logging do not add any additional configuration for commercial deployments from what the operating system provides, but it can be improved by adding additional security settings by enabling STIG.

For more information, see:

- *Enable FIPS, SE Linux, and STIG on the DB System Components*
- *Security Technical Implementation Guide (STIG) Tool for the DB System*

User Security

This topic describes the user security available in the Base Database Service. The Base Database Service components are regularly managed by several user accounts. Oracle uses and recommends token-based SSH login only. Oracle users or processes do not use password-based authentication.

The following kinds of users are created by default:

- Default Users: No Logon Privileges
- Default Users: With Login Privileges

Default Users: No Logon Privileges

This user list consists of default operating system users. These users should not be altered. These users cannot login to the system.

```
bin:x:1:1:bin:/bin:/sbin/nologin
daemon:x:2:2:daemon:/sbin:/sbin/nologin
adm:x:3:4:adm:/var/adm:/sbin/nologin
lp:x:4:7:lp:/var/spool/lpd:/sbin/nologin
sync:x:5:0:sync:/sbin:/bin/sync
shutdown:x:6:0:shutdown:/sbin:/sbin/shutdown
halt:x:7:0:halt:/sbin:/sbin/halt
mail:x:8:12:mail:/var/spool/mail:/sbin/nologin
operator:x:11:0:operator:/root:/sbin/nologin
games:x:12:100:games:/usr/games:/sbin/nologin
ftp:x:14:50:FTP User:/var/ftp:/sbin/nologin
nobody:x:99:99:Nobody:/:/sbin/nologin
systemd-network:x:192:192:systemd Network Management:/:/sbin/nologin
dbus:x:81:81:System message bus:/:/sbin/nologin
rpc:x:32:32:Rpcbind Daemon:/var/lib/rpcbind:/sbin/nologin
polkitd:x:999:996:User for polkitd:/:/sbin/nologin
rpcuser:x:29:29:RPC Service User:/var/lib/nfs:/sbin/nologin
nfsnobody:x:65534:65534:Anonymous NFS User:/var/lib/nfs:/sbin/nologin
ntp:x:38:38::/etc/ntp:/sbin/nologin
tss:x:59:59:Account used by the trousers package to sandbox the
tcsd daemon:/dev/null:/sbin/nologin
sssd:x:998:994:User for sssd:/:/sbin/nologin
named:x:25:25:Named:/var/named:/sbin/nologin
sshd:x:74:74:Privilege-separated SSH:/var/empty/sshd:/sbin/nologin
dhcpd:x:177:177:DHCP server:/:/sbin/nologin
saslauth:x:997:76:Saslauthd user:/run/saslauthd:/sbin/nologin
tcpdump:x:72:72::/:/sbin/nologin
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Default Users: With Login Privileges

These privileged users are responsible for accomplishing most of the tasks in the system. These users should never be altered or deleted as it would have a significant impact on the running system. SSH keys are used for logging in.

The following is the list of default users with login privileges.

- `root` is a Linux requirement. It is used sparingly to run local privileged commands. The root is also used for some processes like TFA Agent. It runs the local agent (aka "DCS Agent") that performs lifecycle operations for RDBMS software (patching, create database, etc.)
- `oracle` owns the Oracle Database software installation and runs RDBMS processes.
- `grid` owns the Oracle Grid Infrastructure software installation and runs GI processes.
- `opc` is used by Oracle Cloud Automation for automation tasks. User has the ability to run certain privileged commands without further authentication (to support automation functions).
- `mysql` is a critical user, and it needs to be up and running for successful operation of DCS Agent as it owns the metastore of DCS Agent.

```
root:x:0:0:root:/root:/bin/bash
opc:x:54322:54323::/home/opc:/bin/bash
mysql:x:54323:54331::/home/mysql:/bin/bash
grid:x:102:1001::/home/grid:/bin/bash
oracle:x:101:1001::/home/oracle:/bin/bash
```

Security Settings

This topic describes the security settings available in the Base Database Service. The following are the default security settings provided in the system.

Table 6-1 Security Settings and Default Values

Security Settings	Default Values
Password complexity	• Password minimum length: 15 • Password to maximum of consecutive repeating characters from the same character class: 4 • Password maximum consecutive repeating characters: 3 • Password strength minimum digit characters: 1 • Password strength minimum different categories: 4 • Password strength minimum different characters: 8 • Password strength minimum special characters: 1 • Password strength minimum lowercase characters: 1 • Password strength minimum uppercase characters: 1
User account configuration	• Maximum number of days a password may be used: 60 • Minimum number of days allowed between password changes: 1 • Encryption hash algorithm: SHA512 • Logon failure delay: 4 seconds
Disabled options	• Disabled Ctrl-Alt-Del Reboot • DCCP support is disabled • USB storage device is disabled • X Windows Package Group is removed

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 6-1 (Cont.) Security Settings and Default Values

Security Settings	Default Values
SSH Configurations	• Only SSH Protocol 2 is allowed • Enabled use of privilege separation • SSH idle timeout interval: 600 seconds • GSSAPI Authentication disabled • Compression set to delayed • Non-certificate trusted allowed to SSH logon to the system • SSH daemon does not allow authentication using known hosts authentication
Packages	• Remove all software components after updated versions have been installed • The system prevents the installation of local packages for unverified software, patches, service packs, device drivers, or operating system components
Logging	• System and Kernel messages are sent to remote host (rsyslog) • Cron configured to log to rsyslog • Configure AIDE for periodic execution
Others	• Authentication required upon booting into single-user and maintenance modes • Interactive session timeout: 600 seconds • Configured PAM in SSSD devices • PAM system service configured to store only encrypted representations of passwords • Configured SSSD LDAP backend client CA certificate location • Configured SSSD LDAP backend to use TLS for all transactions • Disable account after password expires • Group account administration utilities are configured to store only encrypted representations of passwords

Additionally, by default, ONSR regions enable FIPS, SE Linux, and STIG to comply with the requirements standards. You can improve the system security by enabling additional configurations. The configuration standard (STIG) can be set to follow the most restrictive standards and increase security compliance with DISA's Oracle Linux 7 STIG. A tool is provided as a part of the image to enable FIPS, SE Linux, and STIG.

For more information, see:

- *Enable FIPS, SE Linux, and STIG on the DB System Components*
- *Security Technical Implementation Guide (STIG) Tool for the DB System*

Security Processes

This topic describes the default security processes available in the Base Database Service. The following are the list of processes that are run by default on the user virtual machine (DB system) also called the domU.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 6-2 Security Processes

Processes	Description
domU agent	It is a cloud agent for handling database lifecycle operations. • Runs as `root` user • process table shows it running as a java process with the following jar names: – `dcs-agent-VersionNumber-SNAPSHOT.jar` – `dcs-admin-VersionNumber-SNAPSHOT.jar`
TFA Agent	The Oracle Trace File Analyzer (TFA) provides several diagnostic tools in a single bundle, making it easy to gather diagnostic information about the Oracle Database and Clusterware, which in turn helps with problem resolution when dealing with Oracle Support. • Runs as `root` user • runs as `initd` demon (`/etc/init.d/init.tfa`) • process tables show a java application (`oracle.rat.tfa.TFAMain`)
Database and GI (clusterware)	• Runs as `oracle` and `grid` users • some of CRS/clusterware daemon process runs as `root` user • process table shows following applications: – `ora_*`, `apx_*`, `ams_*`, and `oracle+ASM*` – `mysqld` and `zookeeper` – some of other process from `/u01/<version>/grid/*`

Network Security

This topic describes the network security in the Base Database Service. The following are the list of default ports, processes, and iptables rules that are run by default on the user virtual machine (DB system), also called the domU.

Ports for domU Service

The following table provides a list of default ports for domU services.

Table 6-3 Default port matrix for domU services

Type of interface	Name of interface	Port	Process running
Listen on all interfaces	0.0.0.0	22	SSH
		1522	RDBMS: TNS listener
		7060	DCS Admin
		7070	DCS Agent
		2181	Zookeeper
		8888, 8895	RAC: Quality of Management Service (QOMS) Server
		9000	RAC: Oracle Clusterware
		68	DHCP
		123	NTP
		5353	Multicast DNS
Client Interface	ens3	1521	RDBMS: TNS listener
		5000	RDBMS: Autonomous Health Framework (AHF) (includes TFA)

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 6-3 (Cont.) Default port matrix for domU services

Type of interface	Name of interface	Port	Process running
Cluster Interconnect	ens3:1	1521	RDBMS: TNS listener
	ens3:2	1521	RDBMS: TNS listener
	ens3:3	1521	RDBMS: TNS listener
	ens4	1525	RDBMS: TNS listener
		2888	Zookeeper
		3888	Zookeeper
		6000	RAC: Grid inter-process communication
		7000	RAC: High availability service

iptables Rules for domU

The default iptables is set up to `ACCEPT` connections on input, forward, and output chains.

The following are the default iptables rules for domU services:

- `CHAIN INPUT`
- `CHAIN FORWARD`
- `CHAIN OUTPUT`

Example 6-1 iptables rules

The following example provides the default iptables rules for domU services.

```
iptables -L -n -v
```

Output:

```
Chain INPUT (policy ACCEPT 0 packets, 0 bytes)
 pkts bytes target     prot opt in     out     source
destination
  43M  110G ACCEPT     all  --  *      *       0.0.0.0/0
0.0.0.0/0           state RELATED,ESTABLISHED
 2664  224K ACCEPT     icmp --  *      *       0.0.0.0/0            0.0.0.0/0
40793 2441K ACCEPT     all  --  lo     *       0.0.0.0/0            0.0.0.0/0
    0     0 ACCEPT     all  --  ens4   *       0.0.0.0/0            0.0.0.0/0
    3   192 ACCEPT     tcp  --  *      *       0.0.0.0/0
0.0.0.0/0           state NEW tcp dpt:22
   40  2400 ACCEPT     tcp  --  *      *       0.0.0.0/0
0.0.0.0/0           state NEW tcp dpt:1521 /* Required for access to
Database Listener, Do not remove or modify.  */
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0
0.0.0.0/0           state NEW tcp dpt:5000 /* Required for TFA traffic.  */
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0
0.0.0.0/0           state NEW tcp dpt:6200 /* This rule is recommended and
enables the Oracle Notification Services (ONS) to communicate about Fast
Application Notification (FAN) events.  */
  343 20580 ACCEPT     tcp  --  *      *       169.254.0.0/16
0.0.0.0/0           state NEW tcp dpt:7070 /* Required for instance
management by the Database Service, Do not remove or modify.  */
  132  7920 ACCEPT     tcp  --  *      *       169.254.0.0/16
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
0.0.0.0/0            state NEW tcp dpt:7060 /* Required for instance
management by the Database Service, Do not remove or modify.  */
    0     0 ACCEPT     tcp  --  *      *       169.254.0.0/16
0.0.0.0/0            state NEW tcp dpt:22 /* Required for instance management
by the Database Service, Do not remove or modify.  */
    3   424 REJECT     all  --  *      *       0.0.0.0/0
0.0.0.0/0            reject-with icmp-host-prohibited

Chain FORWARD (policy ACCEPT 0 packets, 0 bytes)
 pkts bytes target     prot opt in     out     source
destination
    0     0 REJECT     all  --  *      *       0.0.0.0/0
0.0.0.0/0            reject-with icmp-host-prohibited

Chain OUTPUT (policy ACCEPT 51078 packets, 3218K bytes)
 pkts bytes target     prot opt in     out     source
destination
    0     0 ACCEPT     all  --  *      ens4    0.0.0.0/0            0.0.0.0/0
  52M  170G ACCEPT     all  --  *      *       0.0.0.0/0
0.0.0.0/0            state RELATED,ESTABLISHED
 8003  548K InstanceServices  all  --  *      *       0.0.0.0/0
169.254.0.0/16

Chain InstanceServices (1 references)
 pkts bytes target     prot opt in     out     source
destination
   11   660 ACCEPT     tcp  --  *      *       0.0.0.0/0
169.254.2.0/24       owner UID match 0 tcp dpt:3260 /* See the Oracle-
Provided Images section in the Oracle documentation for security impact of
modifying or removing this rule */
    1    60 ACCEPT     tcp  --  *      *       0.0.0.0/0
169.254.0.2          owner UID match 0 tcp dpt:3260 /* See the Oracle-
Provided Images section in the Oracle documentation for security impact of
modifying or removing this rule */
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0
169.254.0.2          tcp dpt:80 /* See the Oracle-Provided Images section in
the Oracle documentation for security impact of modifying or removing this
rule */
  678 63323 ACCEPT     udp  --  *      *       0.0.0.0/0
169.254.169.254      udp dpt:53 /* See the Oracle-Provided Images section in
the Oracle documentation for security impact of modifying or removing this
rule */
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0
169.254.169.254      tcp dpt:53 /* See the Oracle-Provided Images section in
the Oracle documentation for security impact of modifying or removing this
rule */
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0
169.254.0.3          owner UID match 0 tcp dpt:80 /* See the Oracle-Provided
Images section in the Oracle documentation for security impact of modifying
or removing this rule */
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0
169.254.0.4          tcp dpt:80 /* See the Oracle-Provided Images section in
the Oracle documentation for security impact of modifying or removing this
rule */
 2569  195K ACCEPT     udp  --  *      *       0.0.0.0/0
169.254.169.254      udp dpt:123 /* Allow access to OCI local NTP service */
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
 4727  284K ACCEPT     tcp  --  *      *       0.0.0.0/0
169.254.169.254      tcp dpt:80 /* See the Oracle-Provided Images section in
the Oracle documentation for security impact of modifying or removing this
rule */
   15  4920 ACCEPT     udp  --  *      *       0.0.0.0/0
169.254.169.254      udp dpt:67 /* See the Oracle-Provided Images section in
the Oracle documentation for security impact of modifying or removing this
rule */
    0     0 ACCEPT     udp  --  *      *       0.0.0.0/0
169.254.169.254      udp dpt:69 /* See the Oracle-Provided Images section in
the Oracle documentation for security impact of modifying or removing this
rule */
    0     0 REJECT     tcp  --  *      *       0.0.0.0/0
169.254.0.0/16       tcp /* See the Oracle-Provided Images section in the
Oracle documentation for security impact of modifying or removing this rule
*/ reject-with tcp-reset
    0     0 REJECT     udp  --  *      *       0.0.0.0/0
169.254.0.0/16       udp /* See the Oracle-Provided Images section in the
Oracle documentation for security impact of modifying or removing this rule
*/ reject-with icmp-port-unreachable
```

User Responsibilities for Security Settings

This topic describes the Oracle Cloud Operations responsibilities and user responsibilities for security settings in the Base Database Service. The following table provides a list of security settings that the Oracle Cloud Operations and user need to perform.

Table 6-4 Oracle Cloud Operations and User Responsibilities for Various Operations

Operation	Oracle Cloud Platform		User / Tenant Instances	
	Oracle Cloud Responsibility	**User Responsibility**	**Oracle Cloud Responsibility**	**User Responsibility**
DATABASE DEPLOYMENT	Software infrastructure and guidance for Base Database Service deployment	**Network Admin:** Configure cloud network infrastructure (VCN and subnets, gateway, etc). **Database Admin:**Setup database requirements (memory, storage, computation, database version, database type, etc).	Install operating system, database and Grid Infrastructure system if selected	**Database Admin:** Update Oracle Database software version, shape of virtual machine requirements (CPU / memory), data storage and recovery storage configuration size resources based on workloads if required (upgrade/ downgrade resources).
MONITORING	Physical security, infrastructure, control plane, hardware faults, availability, capacity	Nothing required	Infrastructure availability to support user monitoring of user services.	**Database Admin:** Monitoring of user operating system, databases, apps and Grid Infrastructure

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 6-4 (Cont.) Oracle Cloud Operations and User Responsibilities for Various Operations

Operation	Oracle Cloud Platform		User / Tenant Instances	
	Oracle Cloud Responsibility	User Responsibility	Oracle Cloud Responsibility	User Responsibility
INCIDENT MANAGEMENT AND RESOLUTION	Incident management and remediationspare parts and field dispatch	Nothing required	Support for any incidents related to the underlying platform	**Database Admin:** Incident Management and resolution for user apps
PATCH MANAGEMENT	Proactive patching of hardware, IaaS/ PaaS control stack	Nothing required	Staging of available patches, for example, Oracle Database patch set	**Database Admin:** Patching of tenant instances, testing **OS Admin:** OS patching
BACKUP AND RESTORATION	Infrastructure and control plane backup and recovery, recreate user virtual machines	Nothing required	Provide running and user accessible virtual machines	**Database Admin:** Snapshots / backup and recovery of user IaaS and PaaS data using Oracle native or third-party capability

Enable Additional Security Capabilities

The Base Database Service provides the following additional security capabilities:

- dbcli NetSecurity
- OCI Vault Integration
- CLI to Enable FIPS

dbcli NetSecurity

The dbcli NetSecurity deals with the encryption of data as it travels through the network. When the data moves from Oracle Database to a third party or from a server to client, it has to be encrypted at the sender's end and decrypted at the receiver's end. In NetSecurity, rules are configured with default values for both client and server during provisioning and database home creation operations. The `dcs-agent` CLI interface provides commands to update these NetSecurity rules and enhance security for encryption algorithms, integrity algorithms, and connection types.

By default, `dcs-agent` configures the following default rules for the database home:

- `SQLNET.ENCRYPTION_SERVER=REQUIRED`
- `SQLNET.CRYPTO_CHECKSUM_SERVER=REQUIRED`
- `SQLNET.ENCRYPTION_TYPES_SERVER=(AES256,AES192,AES128)`
- `SQLNET.CRYPTO_CHECKSUM_TYPES_SERVER=(SHA1)`
- `SQLNET.ENCRYPTION_CLIENT=REQUIRED`
- `SQLNET.CRYPTO_CHECKSUM_CLIENT=REQUIRED`

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `SQLNET.ENCRYPTION_TYPES_CLIENT=(AES256,AES192,AES128)`
- `SQLNET.CRYPTO_CHECKSUM_TYPES_CLIENT=(SHA1)`

For more information on updating the settings, see *Oracle Database CLI Reference*.

OCI Vault Integration

The Base Database Service now has integration with the OCI Vault service in all OCI commercial regions. You can now create and manage TDE master keys within the OCI Vault that protect your databases. With this feature, you have the option to start using the OCI Vault service to store and manage the master encryption keys. The OCI Vault keys used for protecting databases are stored in a highly available, durable, and managed service.

> **Note**
>
> The OCI Vault integration is only available for Oracle Database versions 19.13 and later.

With OCI Vault integration with Base Database Service, you can:

- Centrally control and manage TDE master keys by enabling OCI Vault-based key encryption while provisioning Oracle Databases on the Base Database Service.
- Have your TDE master keys stored in a highly available, durable, and managed service wherein the keys are protected by hardware security modules (HSM) that meet Federal Information Processing Standards (FIPS) 140-2 Security Level 3 security certification.
- Rotate your encryption keys periodically to maintain security compliance and, in cases of personnel changes, disable access to a database.
- Migrate from Oracle-managed keys to user-managed keys for your existing databases.
- Bring in your own keys—that's BYOK (Bring Your Own Key)—and use them while creating databases with user-managed encryption.

> **Note**
>
> - BYOK is applicable to the container database (CDB) only. The pluggable database (PDB) will be assigned an automatically generated new key version.
> - Oracle Databases that use user-managed encryption support DB system cloning, in-place restore, out-of-place restore, intra-region Data Guard configuration, and PDB-specific operations like PDB creation and local cloning.

CLI to Enable FIPS

Oracle provides a tool for commercial users to improve security by default. This tool is used to enable FIPS, SE Linux, and STIG to follow the most rigorous standards.

For more information, see *Enable FIPS, SE Linux, and STIG on the DB System Components*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Use Identity and Access Management Authentication with Base Database Service

You can configure the Oracle Database in the Base Database Service to use Oracle Cloud Infrastructure Identity and Access Management (OCI IAM) authentication and authorization to allow IAM users to access the database with IAM credentials.

> **Note**
>
> Oracle Database supports the Base Database Service integration for OCI IAM with identity domains as well as the legacy IAM, which does not include identity domains. Both default and non-default domain users and groups are supported when using IAM with identity domains.

Support for non-default custom domains is only available with Oracle Database Release 19c, Version 19.21 and higher (but not Oracle Database Release 21c).

About IAM Authentication

OCI IAM integration with Base Database Service supports both database password verifier authentication and token-based authentication. For more information on the architecture for using IAM users on Base Database Service, see Authenticating and Authorizing IAM Users for Oracle DBaaS Databases.

IAM Database Password Authentication

You can enable an Oracle Database instance to allow user access with an OCI IAM database password (using a password verifier).

For password verifier database access, you create the mappings for IAM users and OCI applications to the Oracle 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 a custom, non-default domain.

> **Note**
>
> Any supported 12c and above database client can be used for IAM database password access to the Oracle Database.

An OCI IAM database password allows an IAM user to log in to a 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 OCI Console password. Using an IAM user with a password verifier you can login to the database with any supported database client as long as the database client supports Oracle Database 12c password verifiers.

IAM SSO Token Based Authentication

You can enable an Oracle Database instance to use OCI IAM SSO tokens.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For token database access, you create the mappings for IAM users and OCI applications to the Oracle 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 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. IAM database token usage requires the Oracle Database client 19.16 and above (not 21c). Limited (not full) IAM database token capabilities are available with some Oracle Database clients 21.5 and above.
- If the application or tool does not support requesting an IAM database token and sending it to the database through the client API, the IAM user can first use OCI 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 OCI 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 OCI IAM instance principal or resource principal to get an IAM database token, and use the IAM database token to authenticate itself to a database instance.
- IAM users and OCI applications can request a database token from IAM with several methods, including using an API-key.
  For more information about configuring client connection, see Configure Client Connection for SQL*Plus that Uses an IAM Token. For more information about other methods such as using a delegation token within an OCI cloud shell, see Authenticating and Authorizing IAM Users for Oracle DBaaS Databases.

If a user enters a username and password to login, then the database driver uses the password verifier method to access the database as the default method.

Prerequisites

The following prerequisites are required for IAM authentication on Base Database Service.

Network Settings

Before using IAM authentication on databases, you must use the Networking service to add a service gateway, a route rule, and an egress security rule to the Virtual Cloud Network (VCN) and subnets where your database resources reside.

1. Create a service gateway in the VCN where your database resources reside by following the instructions in Create the service gateway.
2. After creating the service gateway, add a route rule and an egress security rule to each subnet (in the VCN) where the database resources reside so that these resources can use the gateway to use IAM authentication:
   a. Go to the **Subnet Details** page for the subnet.
   b. In the **Subnet Information** tab, click the name of the subnet's **Route Table** to display its **Route Table Details** page.
   c. In the table of existing **Route Rules**, check whether there is already a rule with the following characteristics:
      - **Destination**: All IAD Services In Oracle Services Network

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Target Type**: Service Gateway
- **Target**: The name of the service gateway you just created in the VCN

If such a rule does not exist, click **Add Route Rules** and add a route rule with these characteristics.

d. Return to the **Subnet Details** page for the subnet.

e. In the subnet's **Security Lists** table, click the name of the subnet's security list to display its **Security List Details** page.

f. In the side menu, under **Resources**, click **Egress Rules**.

g. In the table of existing **Egress Rules**, check whether there is already a rule with the following characteristics:

- **Stateless**: No
- **Destination**: All IAD Services In Oracle Services Network
- **IP Protocol**: TCP
- **Source Port Range**: All
- **Destination Port Range**: 443

h. If such a rule does not exist, click **Add Egress Rules** and add an egress rule with these characteristics.

Environment Settings

Check if WALLET_ROOT is configured or not:

```
show parameters wallet_root;


NAME                 TYPE        VALUE
-------------------- ----------- --------
wallet_root          string
```

If a directory location does not show up for WALLET_ROOT, you will not be able to configure this database with IAM. WALLET_ROOT should be set the next time your database is patched. New databases will come with WALLET_ROOT set.

TLS Configuration

When sending IAM tokens from the database client to the database server, a TLS connection must be established. The TLS wallet with the database certificate for the Base DB Service instance must be stored under the WALLET_ROOT location. Create a tls directory so it looks like: `WALLET_ROOT/<PDB GUID>/tls`

When configuring TLS between the database client and server there are several options to consider.

- Using a self-signed database server certificate vs a database server certificate signed by a commonly known certificate authority.
- One-way TLS (TLS) vs Mutual or two-way TLS (mTLS).
- Client with or without a wallet.

Self-signed certificate: Using a self-signed certificate is a common practice for internally facing IT resources since you can create these yourself and it's free. The resource (in our

Copyright © 2022, 2026, Oracle and/or its affiliates.

case, the database server) will have a self-signed certificate to authenticate itself to the database client. The self-signed certificate and root certificate will be stored in the database server wallet. For the database client to be able to recognize the database server certificate, a copy of the root certificate will also be needed on the client. This self-created root certificate can be stored in a client-side wallet or installed in the client system default certificate store (Windows and Linux only). When the session is established, the database client will check to see that the certificate sent over by the database server has been signed by the same root certificate.

A well-known certificate authority: Using a commonly known root certificate authority has some advantages in that the root certificate is most likely already stored in the client system default certificate store. There is no extra step for the client to store the root certificate if it is a common root certificate. The disadvantage is that this normally has a cost associated with it.

One-way TLS: In the standard TLS session, only the server provides a certificate to the client to authenticate itself. The client doesn't need to have a separate client certificate to authenticate itself to the server (similar to how HTTPS sessions are established). While the database requires a wallet to store the server certificate, the only thing the client needs to have is the root certificate used to sign the server certificate.

Two-way TLS (also called Mutual TLS, mTLS): In mTLS, both the client and server have identity certificates that are presented to each other. In most cases, the same root certificate will have signed both of these certificates so the same root certificate can be used with the database server and client to authenticate the other certificate. mTLS is sometimes used to authenticate the user since the user identity is authenticated by the database server through the certificate. This is not necessary for passing IAM tokens but can be used when passing IAM tokens.

Client with a wallet: A client wallet is mandatory when using mTLS to store the client certificate. However, the root certificate can be stored either in the same wallet or in the system default certificate store.

A client without a wallet: Clients can be configured without a wallet when using TLS under these conditions:

1. One-way TLS is being configured where the client does not have its own certificate, and
2. the root certificate that signed the database server certificate is stored in the system default certificate store. The root certificate would most likely already be there if the server certificate is signed by a common certificate authority. If it's a self-signed certificate, then the root certificate would need to be installed in the system default certificate store to avoid using a client wallet.

For details on how to configure TLS between the database client and database server including the options described above, see Configuring Transport Layer Security Authentication.

If you choose to use self-signed certificates and for additional wallet related tasks, refer to the `orapki` command line interface (CLI) reference guide in the Database Security Guide. See Managing Public Key Infrastructure (PKI) Elements.

Change External Identity Providers

This topic describes the steps to change the external identity provider from Centrally Managed Users (CMU) to OCI IAM authentication and authorization and vice-versa on Base Database Service.

OCI IAM authentication and authorization for users is not enabled for newly provisioned databases, by default. Another option for external authentication is to use Centrally Managed

Copyright © 2022, 2026, Oracle and/or its affiliates.

Users with Active Directory (CMU-AD). There can only be one external authentication scheme enabled at any given time.

Enable OCI IAM Authentication and Authorization

Perform the following steps to enable OCI IAM authentication and authorization.

1. Enable OCI IAM authentication and authorization using the `ALTER SYSTEM` command.

   ```
   ALTER SYSTEM SET IDENTITY_PROVIDER_TYPE=OCI_IAM SCOPE=BOTH;
   ```

2. Verify the value of `IDENTITY_PROVIDER_TYPE` system parameter.

   ```
   SELECT NAME, VALUE FROM V$PARAMETER WHERE NAME='identity_provider_type';


   NAME                    VALUE
   ---------------------- -------
   identity_provider_type OCI_IAM
   ```

3. Check if the `IDENTITY_PROVIDER_CONFIG` parameter has been set.

   ```
   SELECT NAME, VALUE FROM V$PARAMETER WHERE NAME='identity_provider_config';
   ```

4. If the `IDENTITY_PROVIDER_CONFIG` parameter has been set, then reset this parameter.

   ```
   ALTER SYSTEM RESET IDENTITY_PROVIDER_CONFIG SCOPE=BOTH;
   ```

Disable OCI IAM Authentication and Authorization

Perform the following step to disable OCI IAM authentication and authorization.

1. Disable OCI IAM integration using the `ALTER SYSTEM` command.

   ```
   ALTER SYSTEM RESET IDENTITY_PROVIDER_TYPE SCOPE=BOTH;
   ```

Enable CMU-AD

Perform the following steps to enable Active Directory (AD) users to connect to the database using CMU:

1. Disable IAM integration as described in Disable OCI IAM Authentication and Authorization.
2. Configure CMU-AD as described in Configuring Centrally Managed Users with Microsoft Active Directory.

Disable CMU-AD

Perform the following step to disable CMU-AD:

1. Disable CMU-AD using the `ALTER SYSTEM` command.

   ```
   ALTER SYSTEM SET LDAP_DIRECTORY_ACCESS = 'NONE';
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Re-enable OCI IAM Authentication and Authorization

Perform the following steps to re-enable IAM users to connect to the database using OCI IAM authentication and authorization:

1. Disable CMU-AD as described in Disable CMU-AD.
2. Enable OCI IAM authentication and authorization as described in Enable OCI IAM Authentication and Authorization.

Create IAM Groups and Policies for IAM Users

This topic describes the steps to write policy statements for an IAM group to enable IAM user access to OCI resources, specifically the 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 the database. A policy is not required when using IAM database passwords to access the database.

To enable the database to allow IAM users to connect to the database using IAM tokens:

1. Perform OCI IAM prerequisites by creating a group and adding users to the group.

   For example, create the group `sales_dbusers`.

   For more information, see Managing Groups.
2. Write policy statements to enable access to OCI resources.
   a. In the OCI Console, click **Identity and Security** and click **Policies**.
   b. To a write policy, 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 users in IAM group `DBUsers` to access any database in their tenancy:

      ```
      Allow group DBUsers to use database-connections in tenancy
      ```

      For example to create a policy that limits members of `DBUsers` group to access the databases in compartment `testing_compartment` only:

      ```
      allow group DBUsers to use database-connections in compartment
      testing_compartment
      ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

For example to create a policy that limits group access to a single database in a compartment:

```
allow group DBUsers to use database-connections in compartment
testing_compartment
    where target.database.id = 'ocid1.database.oc1.iad.aabbcc'
```

d. Click **Create**.
For more information on policies, see Managing Policies.

> **Note**
>
> The following is required for creating policies for use with IAM users on database in the Base Database Service.
>
> - Policies can allow IAM users to access database instances across the entire tenancy, in a compartment, or can limit access to a single database instance.
> - You can use either instance principal or resource principal to retrieve database tokens to establish a connection from your application to an database instance. If you are using an instance pricipal 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 OCI.
>
>   For more information, see Managing Dynamic Groups.

Add IAM Users

To add IAM users to allow access to the 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 a database instance works by mapping database global users (schemas) to IAM users (exclusive mapping) or IAM groups (shared schema mapping).

Authorize IAM Users on a Database Instance:

Perform the following steps to authorize IAM users on a 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 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';
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

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, `sales_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 database is enabled for IAM authentication.

Exclusively Map a Local IAM User to an Oracle Database Global User:

Perform the following steps to exlcusively 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 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

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.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Creating global roles is optional for an IAM user with an exclusive IAM mapping to a database user (schema). When the IAM mapping is to a shared schema, creating a global role is also optional. 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.

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 works by mapping the database global roles to IAM groups.

Map the Database Global Roles to IAM Groups:

Perform the following steps to map the 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 the 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;
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

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 the above 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 the database by supplying a username and password, you must create an IAM database password. An IAM username and IAM database password can be used in one of two ways.

1. The IAM user can enter the IAM username and IAM database password when accessing the database. By default, the database client will follow the normal password authentication mechanism with the database and the database will retrieve the IAM database password verifier from IAM.
2. The database client can be configured to get an IAM database token using the IAM username and IAM database password. The database client will send this database token to the database for user access.

For more information on getting an IAM database token using the IAM username and IAM database password, see Configuring IAM for Oracle DBaaS.

For more information on IAM database passwords, see *Working with IAM Database User Names and Passwords* in Managing User Credentials.

Connect to Database with IAM Authentication

After the database ADMIN user maps global users and global roles to the IAM users and IAM groups, users log in to the database instance using their OCI IAM credentials or access the database through an OCI IAM database token.

You can still log in to the database using your local database account username and password (non-global database user account).

You can use a database client to access a database instance as an OCI IAM user. To use a client with OCI IAM username and password credentials and a password verifier, the database client must be 12c or newer.

IAM database token usage requires the Oracle Database client 19.16 and above (not 21c). Limited (not full) IAM database token capabilities are available with some Oracle Database clients 21.5 and above.

> **Note**
>
> If your database instance is in Restricted Mode, only the users with the `RESTRICTED SESSION` privilege such as ADMIN can connect to the database.

Copyright © 2022, 2026, Oracle and/or its affiliates.

About Connecting to a Database Instance Using IAM

IAM users can connect to the database instance by using either an IAM database password verifier or an IAM token.

Using the IAM database password verifier is similar to the database password authentication process. However, instead of the password verifier (encrypted hash of the password) being stored in the database, the verifier is instead stored as part of the OCI IAM user profile.

The second connection method makes use of an IAM token for the database. The use of token-based access is a better fit for Cloud resources such as Oracle Databases in the Base Database Service. 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.

The IAM user enters the IAM user name and IAM database password (not the 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 database.

The database client is also able to request a token directly from IAM using the IAM username and IAM database password. For more information about configuring the client to get a token, see Client Connections That Use a Token Requested by an IAM User Name and Database Password.

Client Connections that Use a Token

For IAM token access to the 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 OCI 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.
- IAM username and IAM database password, which can be used by the database client to retrieve an IAM database token directly when configured to do so.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- 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 OCI services.
- Using an IAM username and IAM database password (which can only be requested by the 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, unless the database client is configured to retrieve a database token from IAM with the IAM username and IAM database password. The instructions in this topic 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: SQL*Plus on-premises, SQLcl on-premises, SQL*Plus in Cloud Shell, or applications that use SEP wallets.

The following topics explain how to:

- Configure a Client Connection for SQL*Plus that Uses an IAM Database Password
- Configure Client Connection for SQL*Plus that Uses an IAM Token
- Use Instance Principal to Access Database with IAM Authentication

Configure 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 database 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 a 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"
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Configure Client Connection for SQL*Plus that Uses an IAM Token

Perform the following steps to 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 the 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. For more information about installing and configuring the OCI CLI, see Quickstart.
4. 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 OCI CLI:

        ```
        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, execute the following command in the OCI CLI:

        ```
        oci iam db-token get
        ```

      - Retrieving an instance token by using the OCI CLI:

        ```
        oci iam db-token get --auth instance_principal
        ```

   For more information,see Required Keys and OCIDs.
5. This configuration only works with the Oracle Database 19c client. Ensure that you are using the latest release updates for this client.

   > **Note**
   >
   > Oracle Database client release 21c offers limited IAM token features.

6. Follow the existing process to download the wallet from the 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`.

Copyright © 2022, 2026, Oracle and/or its affiliates.

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_cert_dn="CN=example.uscom-east-1.oraclecloud.com,
     OU=Oracle BMCS US, O=Example Corporation,
     L=Redwood City, ST=California, C=US")
  (TOKEN_AUTH=OCI_TOKEN)))
```

After the connect string is updated with the `TOKEN_AUTH` parameter, the IAM user can log in to the 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,
     OU=Oracle BMCS US, 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 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` unless another parameter is specified: `PASSWORD_AUTH = OCI_TOKEN`. This directs the database client to get the token from IAM using the IAM username and IAM database password. In addition to setting `PASSWORD_AUTH`, you will also need to set `OCI_IAM_URL`, `OCI_TENANCY` and optionally `OCI_COMPARTMENT` and `OCI_DATABASE`.

Use Instance Principal to Access Database with IAM Authentication

After the ADMIN user enables OCI IAM on the database, an application can access the database through an OCI IAM database token using an instance principal.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For more information, see Accessing the Oracle Cloud Infrastructure API Using Instance Principals.

Configure Proxy Authentication

Proxy authentication allows an IAM user to proxy to a database schema for tasks such as application maintenance.

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.

Configuring 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 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
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Validating 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 database instance as a user who has the `CREATE USER` and `ALTER USER` system privileges.
2. Connect at the IAM user and execute 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`. You will need to connect to the database using the different types of authentication methods shown here, but the output of the commands that you execute will be the same for all types.
   - 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"
     ```

Use Database Link with IAM Authenticated Users

You can use a database link to connect from one database instance to another as an OCI IAM user.

You can use either connected user or fixed user database link to connect to a database as an OCI IAM user.

> **Note**
>
> Current user database link is not supported for connecting to a database in Base Database Service as an OCI IAM user.

- **Connected User Database Link**: For a connected user database link, an IAM user must be mapped to a schema in both the source and target databases connected by a database

Copyright © 2022, 2026, Oracle and/or its affiliates.

link. You can use a database password verifier or an IAM database token to use a connected user database link.

- **Fixed User Database Link**: A fixed user database link can be created using a database user or an IAM user. When using an IAM user as a fixed user database link, the IAM user must have a schema mapping in the target database. The IAM user for a database link can be configured with a password verifier only.

Disable IAM Authentication

You can disable IAM user access on your database instance using the `ALTER SYSTEM` command as shown below:

```
ALTER SYSTEM RESET IDENTITY_PROVIDER_TYPE SCOPE=BOTH;
```

If you also want to update access to IAM from the resource, you may need to remove or modify the IAM group and the policies you set up to allow access to IAM from those resources.

Use Azure Active Directory Authentication with Base Database Service

You can configure the Oracle Database in the Base Database Service to use Microsoft Azure Active Database authentication and authorization to allow Azure AD users to access the database with Azure AD credentials.

About Integrating Azure AD with Base Database Service

An Oracle Database in a Base Database Service instance can be configured for Microsoft Azure Active Directory (Azure AD) users to connect using Azure `OAuth2` access tokens.

For more information about authorizing Azure AD users, architecture, user mappings, use cases, and the integration process, see Introduction to Authorizing Microsoft Azure AD Users for an Oracle Database section in the Oracle Database Security Guide.

Prerequisites

The following prerequisites are required for Azure AD authentication on Base Database Service.

- Network Settings
- TLS Configuration

Each of these are described in detail in the following topics.

Network Settings

Before using Azure AD authentication on databases, you must use the Networking service to add a service gateway, a route rule, and an egress security rule to the Virtual Cloud Network (VCN) and subnets where your database resources reside. Perform the following steps to configure outbound connectivity to Azure AD using a NAT gateway.

1. Create a NAT gateway in the VCN where your database resources reside by following the instructions in Create the service gateway.

Copyright © 2022, 2026, Oracle and/or its affiliates.

2. After creating the service gateway, add a route rule and an egress security rule **to each subnet** (in the VCN) where the database resources reside so that these resources can use the gateway to obtain a public key from your Azure AD instance to use Azure AD authentication:
   a. Go to the **Subnet Details** page for the subnet.
   b. In the **Subnet Information** tab, click the name of the subnet's **Route Table** to display its **Route Table Details** page.
   c. In the table of existing **Route Rules**, check whether there is already a rule with the following characteristics:
      - **Destination**: 0.0.0.0/0
      - **Target Type**: NAT Gateway
      - **Target**: The name of the service gateway you just created in the VCN

      If such a rule does not exist, click **Add Route Rules** and add a route rule with these characteristics.
   d. Return to the **Subnet Details** page for the subnet.
   e. In the subnet's **Security Lists** table, click the name of the subnet's security list to display its **Security List Details** page.
   f. In the side menu, under **Resources**, click **Egress Rules**.
   g. In the table of existing **Egress Rules**, check whether there is already a rule with the following characteristics:
      - **Destination Type**: CIDR
      - **Destination**: 0.0.0.0/0
      - **IP Protocol**: TCP
      - **Source Port Range**: 443
      - **Destination Port Range**: All
   h. If such a rule does not exist, click **Add Egress Rules** and add an egress rule with these characteristics.

TLS Configuration

When sending Azure AD tokens from the database client to the database server, a TLS connection must be established. The TLS wallet with the database certificate for the Base Database Service instance must be stored under the WALLET_ROOT location. Create a tls directory so it looks like: `WALLET_ROOT/<PDB GUID>/tls`

When configuring TLS between the database client and server there are several options to consider.

- Using a self-signed database server certificate vs a database server certificate signed by a commonly known certificate authority.
- One-way TLS (TLS) vs Mutual or two-way TLS (mTLS).
- Client with or without a wallet.

Self-signed certificate: Using a self-signed certificate is a common practice for internally facing IT resources since you can create these yourself and it's free. The resource (in our case, the database server) will have a self-signed certificate to authenticate itself to the database client. The self-signed certificate and root certificate will be stored in the database server wallet. For the database client to be able to recognize the database server certificate, a

Copyright © 2022, 2026, Oracle and/or its affiliates.

copy of the root certificate will also be needed on the client. This self-created root certificate can be stored in a client-side wallet or installed in the client system default certificate store (Windows and Linux only). When the session is established, the database client will check to see that the certificate sent over by the database server has been signed by the same root certificate.

A well-known certificate authority: Using a commonly known root certificate authority has some advantages in that the root certificate is most likely already stored in the client system default certificate store. There is no extra step for the client to store the root certificate if it is a common root certificate. The disadvantage is that this normally has a cost associated with it.

One-way TLS: In the standard TLS session, only the server provides a certificate to the client to authenticate itself. The client doesn't need to have a separate client certificate to authenticate itself to the server (similar to how HTTPS sessions are established). While the database requires a wallet to store the server certificate, the only thing the client needs to have is the root certificate used to sign the server certificate.

Two-way TLS (also called Mutual TLS, mTLS): In mTLS, both the client and server have identity certificates that are presented to each other. In most cases, the same root certificate will have signed both of these certificates so the same root certificate can be used with the database server and client to authenticate the other certificate. mTLS is sometimes used to authenticate the user since the user identity is authenticated by the database server through the certificate. This is not necessary for passing Azure AD tokens but can be used when passing Azure AD tokens.

Client with a wallet: A client wallet is mandatory when using mTLS to store the client certificate. However, the root certificate can be stored either in the same wallet or in the system default certificate store.

A client without a wallet: Clients can be configured without a wallet when using TLS under these conditions:

1. One-way TLS is being configured where the client does not have its own certificate, and
2. the root certificate that signed the database server certificate is stored in the system default certificate store. The root certificate would most likely already be there if the server certificate is signed by a common certificate authority. If it's a self-signed certificate, then the root certificate would need to be installed in the system default certificate store to avoid using a client wallet.

For details on how to configure TLS between the database client and database server including the options described above, see Configuring Transport Layer Security Authentication.

If you choose to use self-signed certificates and for additional wallet related tasks, refer to the `orapki` command line interface (CLI) reference guide in the Database Security Guide. See Managing Public Key Infrastructure (PKI) Elements.

Configure Base Database Service for Integration with Azure AD

The Base Database Service integration with the Azure AD requires the database to be registered with Azure AD so that the database can request the Azure AD public key.

To configure the Base Database Service for integration with the Azure AD, you must first complete the prerequisites in the Prerequisites section and then follow the instructions in the Configuring the Oracle Database for Microsoft Azure AD Integration section of the Oracle Database Security Guide.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Map Oracle Database Schemas and Roles

Azure AD users will be mapped to one database schema and optionally to one or more database roles.

For more information on the options for mapping Oracle Database schemas and roles to Azure AD users, see Mapping Oracle Database Schemas and Roles section in the Oracle Database Security Guide.

Configure Client Connections to Azure ADs

There are numerous ways that you can configure a client to connect with a Base Database Service instance using Azure AD tokens.

For more information about configuring Azure AD client connections, see Configuring Azure AD Client Connections to the Oracle Database section in the Oracle Database Security Guide.

Testing the Accessibility of the Azure Endpoint

You can ensure that your Base Database instance can access the Azure AD endpoint by running a connection test.

For more information about testing the connection, see Testing the Accessibility of the Azure Endpoint section in the Oracle Database Security Guide.

Trace Files Used for Troubleshooting Connections

You can use trace files to troubleshoot Base Database Service client connections with Azure AD connections.

For more information about trace files and setting client tracing for token authentication, see Trace Files for Troubleshooting Oracle Database Client Connections with Azure AD section in the Oracle Database Security Guide.

Add SSH Keys to a DB System

This article provides the procedure to add SSH keys to a DB system.

Procedure

Perform the following steps to add SSH keys to a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the DB system details page, from the **Actions** menu, select **Add SSH keys**.
3. In the **Add SSH keys** panel, enter the following details:
   - **SSH keys**: Add the public key portion of each key pair you want to use for SSH access. Select one of the following options:
     - **Generate SSH key pair**: Use this option to create a new SSH key pair. Select both **Save private key** and **Save public key** when using this option. The private key is downloaded to your local system and must be stored in a safe location. You cannot download another copy of the private key generated during this operation after completing the operation.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- **Upload SSH key files**: Select this option to browse or drag and drop your existing public key (.pub) files.
- **Paste SSH keys**: Select this option to paste in individual public keys. To paste multiple keys, select **+ Another SSH key** and supply a single key for each entry.

4. Select **Save**.

Open Ports on the DB System

You can open ports on the DB systems using the following steps.

Open the following ports as needed on the DB system:

- 6200 - For Oracle Notification Service (ONS).
- 1158 - For Enterprise Manager Cloud Control. 1158 is the default port, but each additional console enabled on the DB system will have a different port. For more information, see *Monitor Using Oracle Enterprise Manager*.

For important information about critical firewall rules, see essential firewall rules in *Security Rules for the DB System*.

Procedure

1. SSH to the DB System.

   ```
   ssh -i <private_key_path> opc@<db_system_ip_address>
   ```

2. Log in as opc and then sudo to the root user.

   ```
   login as: opc
   sudo su -
   ```

3. Save a copy of iptables as a backup.

   ```
   iptables-save > /tmp/iptables.orig
   ```

   (If necessary, you can restore the original file by using the command `iptables-restore < /tmp/iptables.orig`.)

4. Dynamically add a rule to iptables to allow inbound traffic on the console port, as shown in the following sample. Change the port number and comment as needed.

   ```
   iptables -I INPUT 8 -p tcp -m state --state NEW -m tcp --dport 1158 -j
   ACCEPT -m comment
       --comment "Required for Enterprise Manager Database Control."
   ```

5. Make sure the rule was added.

   ```
   service iptables status
   ```

6. Save the updated file to `/etc/sysconfig/iptables`.

   ```
   /sbin/service iptables save
   ```

   The change takes effect immediately and will remain in effect when the node is rebooted.

Copyright © 2022, 2026, Oracle and/or its affiliates.

7. Update the DB system's security list as described in *Update the Security List for the DB System*.

Manage Administrator and TDE Wallet Passwords

This article provides the administrative tasks for updating the administrator and TDE wallet passwords of a database in the DB system.

Procedure

Perform the following steps to update administrator and TDE wallet passwords in your DB system using the OCI Console:

> **Note**
>
> Changing the TDE wallet password does NOT rotate the TDE key.

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, from the **Actions** menu, select **Manage passwords**.
4. In the **Manage passwords** panel, select **Update administrator password** or **Update TDE wallet password**, depending on which password you want to update.
5. Enter the new password:
   - For the administrator password, enter the new password in both the **New administrator password** and **Confirm administrator password** fields.
   - For the TDE wallet password, enter the current wallet password in the **Enter existing TDE wallet password** field, then enter the new password in both the **New TDE wallet password** and the **Confirm TDE wallet password** fields.
6. Select **Apply** to update your chosen password.

Enable FIPS, SE Linux, and STIG on the DB System Components

This article describes the procedure to add Federal Information Processing Standards (FIPS), Security Enhanced (SE) Linux, and Security Technical Implementation Guide (STIG) standards security enhancements to the DB system.

For more information, see:

- Oracle Database FIPS 140-2 Settings.
- Use SELinux on Oracle Linux.
- *Security Technical Implementation Guide (STIG) Tool for the DB System*.

Enable FIPS, SE Linux, and STIG

Perform the following steps on each system node.

Copyright © 2022, 2026, Oracle and/or its affiliates.

1. Open an SSH session to the DB system node and switch to the root user, then navigate to `/opt/oracle/dcs/bin`.

```
sudo -s
cd /opt/oracle/dcs/bin
```

2. Run the following command.

```
dbcli secure-dbsystem -se -sd -fo -fd
```

Output:

```
Job details
----------------------------------------------------------------
ID: <job_ID_number>
Description: Secure DB System
Status: Created
Created: November 8, 2020 4:12:29 PM UTC
Progress: 0%
Message:

Task Name Start Time End Time Status
```

3. Verify the job details.

```
dbcli describe-job -i <job_ID_number>
```

The output provides information about the progress, status, and details of the job.

```
Job details
----------------------------------------------------------------
ID: <job_ID_number>
Description: Secure DB System
Status: Success
Created: November 8, 2020 4:12:29 PM UTC
Progress: 100%
Message:

Task Name Start Time End Time Status
------------------------------------------------------------------------
---------------------------------- -------
Enable SE Linux [<name>] November 8, 2020 4:12:31 PM UTC November 8, 2020
4:12:31 PM UTC Success
Enable STIG for DOD [<name>] November 8, 2020 4:12:31 PM UTC November 8,
2020 4:12:49 PM UTC Success
Enable FIPS for OS [<name>] November 8, 2020 4:12:49 PM UTC November 8,
2020 4:14:43 PM UTC Success
Enable FIPS for DB Home [<DB_home_name_1>] November 8, 2020 4:14:43 PM UTC
November 8, 2020 4:14:43 PM UTC Success
Enable FIPS for DB[<DB_name_1>] November 8, 2020 4:14:43 PM UTC November
8, 2020 4:14:46 PM UTC Success
Enable FIPS for DB Home [<DB_home_name_2>] November 8, 2020 4:14:46 PM UTC
November 8, 2020 4:14:46 PM UTC Success
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
Enable FIPS for DB[<DB_name_2>] November 8, 2020 4:14:46 PM UTC November
8, 2020 4:14:49 PM UTC Success
```

4. After the job details output shows the **Status** as "Success", you must restart your DB system node using the Console. This is required because enabling FIPS and SE Linux updates the OS kernel. For instructions, see *Reboot a DB System*.

Checking a DB System Node for FIPS and SE Linux Configurations

To confirm that FIPS and SE Linux are enabled on your DB system node, use the following `dbcli` command.

```
dbcli get-dbsystemsecurestatus
```

The system returns details as shown in the following example.

```
{
  "isSELinuxEnabledForOS" : true,
  "isFipsEnabledForOS" : true,
  "fipsStatusForDBs" : [ {
    "databaseResId" : "<DB_ID_number>",
    "status" : true
  } ]
}
```

Security Technical Implementation Guide (STIG) Tool for the DB System

This article describes the STIG tool, a Python script, for DB Systems provisioned using Oracle Linux 7.

A Security Technical Implementation Guide (STIG) is a document written by the Defense Information Systems Agency (DISA) that provides guidance on configuring a system to meet cybersecurity standards for deployment within the Department of Defense's (DoD) IT network systems. STIG requirements help secure the network against cybersecurity threats by focusing on infrastructure and network security to mitigate vulnerabilities.

The STIG tool, a Python script, is used to ensure security compliance with DISA's Oracle Linux 7 STIG. This tool:

- makes the base image of the DB System compliant with the Oracle Linux 7 STIG,
- embeds certain STIG rules into the system that can be activated after provisioning when required to address security compliance requirements,
- categorizes the embedded rules, enabling you to view and monitor the rules in the following categories:
  - **Static** rules that are included in the base image,
  - **DoD** rules that are optionally activated after provisioning when needed to meet U.S. Department of Defense compliance standards, and
  - **Runtime** rules that are activated after provisioning when needed and are intended for use by all users needing to harden security for DB Systems (including users outside of the U.S. Department of Defense),

Copyright © 2022, 2026, Oracle and/or its affiliates.

- provides a rollback capability, enabling you to roll back a DB System to a state with no configuration modifications made by the script, and
- provides a compliance check capability, enabling you to see how many of the rules are successfully passed by the DB System.

Acquire the STIG Tool

The STIG tool is provided for all newly provisioned DB Systems. The STIG tool is provided in the following OS directory location on DB System nodes: `/opt/oracle/dcs/bin/dbcsstig`

Updated versions of the STIG tool will be available for download from the Oracle Technology Network (OTN). Updated versions of the STIG tool are also provided when you update the DB System agent.

Use the STIG Tool

Use the following syntax for the STIG tool:

```
dbcsstig --<operation><category>
```

For example:

```
dbcsstig --fix dod
```

Command Reference

Operations

Table 6-5 Operations

Operation Parameter	Definition
`--check, -c`	Checks for compliance with rules included in the specified category.
`--fix, -f`	Applies fixes for rules included in the specified category.
`--rollback, -rb`	Rolls back system configuration changes implemented by the STIG tool.
`--version, -v`	Provides version information for the STIG tool script.
`--help, -h`	Provides command-line help information.

Rule Categories

Table 6-6 Rule Categories

Category Parameter	Definition
`static`	To specify rules included in the base image of the DB System.
`dod`	To specify rules required for compliance with DISA's Oracle Linux 7 STIG.
`runtime`	To specify rules activated after provisioning for general security hardening.
`all`	To specify all rules.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Manage Database Security with Oracle Data Safe

Oracle Data Safe, which is included with the Base Database Service, provides a unified control center that helps you manage the day-to-day security and compliance requirements of Oracle Databases.

About Oracle Data Safe

Your corporate policy requires that you monitor your databases and retain audit records. Your developers are asking for copies of production data for that new application, and you're wondering what kinds of sensitive information it will contain. Meanwhile, you need to make sure that recent maintenance activities haven't left critical security configuration gaps on your production databases and that staff changes haven't left dormant user accounts on the databases. Oracle Data Safe assists you with these tasks and is included with your Base Database Service.

Data Safe supports you to evaluate security controls, assess user security, monitor user activity, and address data security compliance requirements for your database by evaluating the sensitivity of your data as well as masking sensitive data for non-production databases.

Data Safe provides the following features:

- **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 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 - and take proactive steps to secure them. User Assessment Baselines make it easy to know when new accounts are added, or an account's privileges are modified. Use OCI events to receive proactive notifications when a database deviates from its baseline.
- **Sensitive Data Discovery:** Knowing what sensitive data is managed in your applications is critical for security and privacy. Data discovery scans your database for over 180 different types of sensitive data, helping you understand what types and how much sensitive data you are storing. Use these reports to formulate audit policies, develop data masking templates, and create effective access control policies.
- **Data Masking:** Minimizing the amount of sensitive data your organization maintains helps 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:** Understanding and reporting on user activity, data access, and changes to database structures supports regulatory compliance requirements and can aid in post-incident investigations. Activity auditing collects audit records from databases and helps you manage audit policies. Audit insights make it simple to identify inefficient audit policies, while alerts based on audit data proactively notify you of risky activity.
- **SQL Firewall Management:** Protect against risks such as SQL injection attacks or compromised accounts. Oracle SQL Firewall is a new security capability built into the Oracle AI Database 26ai kernel and offers best-in-class protection against these risks. The SQL Firewall feature in Oracle Data Safe lets you centrally manage and monitor the SQL

Copyright © 2022, 2026, Oracle and/or its affiliates.

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.

Get Started with Oracle Data Safe

For more information about using Oracle Data Safe, see Administering Oracle Data Safe.

Security Zone Integration

This article describes the support of security zones by the Base Database Service.

A security zone is associated with one or more compartments in your tenancy and is created with a set of security policies called a security recipe. While you can create custom security recipes for particular use cases, this article concentrates on the Oracle-managed **Maximum Security Recipe**, which includes a number of curated security zone policies. The policies of a particular security recipe are enforced on any resource that is provisioned or moved into a security zone that uses the recipe. Thus, the only way to apply security zone policies is to control the compartment assignments of your Oracle Cloud Infrastructure resources.

For more information about security zones, including instructions for creating security zones and security recipes, see Overview of Security Zones.

For more information about compartments, see Compartment.

Restrictions on the Resources Located in Maximum Security Recipe Compartments

The Maximum Security Recipe includes all available security zone policies. For example, restrictions placed on databases in a security zone that uses the Maximum Security Recipe include:

- The database cannot allow public network access.
- The database must have automatic backups enabled.
- The database cannot have Data Guard associations that aren't in compartments within the same security zone.

For a complete list of the database restrictions implemented by the Maximum Security Recipe, see Security Zone Policies.

You can also create custom recipes that do not include all possible security restrictions for the resources and assign a custom recipe to a security zone. For more information, see Managing Recipes.

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 the OCI resources that you assign security attributes to. Security attributes are labels that ZPR uses to identify and organize OCI resources.

ZPR enforces policy at the network level each time access is requested, regardless of potential network architecture changes or misconfigurations.

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 policy. If any NSG, SCL, or ZPR rule or policy doesn't allow traffic, the request is dropped.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Manage ZPR

You can secure networks with Zero Trust Packet Routing (ZPR) in three steps: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.

> **Note**
>
> Administrators must set up security attribute namespaces and security attributes in a tenancy before users can apply security attributes to the DB systems.

For detailed information on ZPR, see Overview of Zero Trust Packet Routing.

Manage ZPR Policies

A ZPR policy is a rule that governs the communication between specific endpoints identified by their security attributes. ZPR policy can only be created in the root compartment of a tenancy.

The following policies are required for Base Database service to enable database service for all scenarios including backup and Data Guard.

Table 6-7 ZPR Policy Use Cases

Use Case	Policy	Notes
Enable database service for all scenarios (includes backup and Data Guard).	`in <security attribute of VCN> VCN allow <security attribute of source-compute> endpoints to connect to <security attribute of database service> endpoints with protocol='tcp/1521'`	This policy allows a compute VM to connect to a DB system.
	`in <security attribute of VCN> VCN allow <security attribute of database service> endpoints to connect to 'osn-services-ip-addresses'`	This policy allows the DB system to connect to OSN services.
	`in <security attribute of VCN> VCN allow <security attribute of database service> endpoints to connect to <security attribute of database service> endpoints`	This policy is required for RAC support.
	`in <security attribute of VCN> VCN allow <security attribute of source-compute> endpoints to connect to <Standby VCN CIDR> with protocol='tcp/1521'`	This policy allows Compute clients to connect to Data Guard Standby VCN.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 6-7 (Cont.) ZPR Policy Use Cases

Use Case	Policy	Notes
	`in <security attribute of Standby VCN> VCN allow <security attribute of database service> endpoints to connect to 'osn-services-ip-addresses'`	This policy allows the Data Guard Standby to connect to OSN services.
	`in <security attribute of VCN> VCN allow <security attribute of database service> endpoints to connect to <Standby VCN CIDR>` `in <security attribute of Standby VCN> VCN allow <VCN CIDR> to connect to <security attribute of database service> endpoints`	This policy allows Data Guard Primary to connect to the Data Guard Standby using CIDR, both egress and ingress in each VCN.
	`in <security attribute of VCN> VCN allow <Standby VCN CIDR> to connect to <security attribute of database service> endpoints` `in <security attribute of Standby VCN> VCN allow <security attribute of database service> endpoints to connect to <VCN CIDR>`	This policy allows Data Guard Standby to connect to the Data Guard Primary using CIDR.

For detailed instructions about deleting, updating, and viewing ZPR policies, see Managing Zero Trust Packet Routing Policies.

Manage Security Attributes

You can add, edit, or remove a security attribute for a DB system. For more information, see *Manage Security Attributes for the DB System*.

Manage Security Attributes for the DB System

This article provides the procedure to add, edit, and remove security attributes for a DB system.

General Information

- Administrators must set up security attribute namespaces and security attributes in a tenancy before users can apply security attributes to the DB systems.
- A security attribute is effective only with appropriate policies. If a security attribute without any policies is added, all access will be denied by default, even if allowed in the Security List or NSGs.
- If you use the security attribute, the security rules for the DB system will be a union of the rules in the security attributes along with any rules in the security list and the NSGs.
- You may need to use more than one security attribute. If you're not sure, contact your network administrator.
- A maximum of 3 security attributes can be specified for a DB system.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For more information about:

- creating a security attribute, see Creating a Security Attribute.
- defining appropriate ZPR policies, see Manage ZPR Policies.

Add a Security Attribute

Perform the following steps to add a security attribute for a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. Select the **Security** tab. A list of existing security attributes, if any, is displayed.
3. Select **Add security attributes**.
   - Alternatively, you can also select **Add security attributes** from the **Actions** menu in the DB system details page.
4. In the **Add security attributes** panel, enter the following details:
   - Select the namespace in which the required security attribute key is available.
   - Select the key and provide a value for the security attribute.
   - Optionally, to use additional security attributes, select **Add security attribute** and provide required details.
5. Select **Save**.

Edit a Security Attribute

Perform the following steps to edit a security attribute for a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. Select the **Security** tab. A list of existing security attributes, if any, is displayed.
3. Select the edit icon of the security attribute that you want to edit.
4. Enter a new value for the security attribute.
5. Select **Save**.

Remove a Security Attribute

Perform the following steps to remove a security attribute for a DB system using the OCI Console:

1. On the **DB Systems** list page, select the DB system that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. Select the **Security** tab. A list of existing security attributes, if any, is displayed.
3. Select the edit icon of the security attribute that you want to remove.
4. Select **Remove**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Policy Details for Base Database Service

This article provides the details for writing Oracle Cloud Infrastructure Identity and Access Management (IAM) policies to control access to Oracle Base Database Service resources.

> **Tip**
>
> For a sample policy, see Let database admins manage Oracle Cloud database systems.

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 `database-family` is equivalent to writing separate policies for the group that would grant access to the `db-systems`, `db-nodes`, `db-homes`, `databases`, `database-software-image`, and `db-backups` resource-types. For more information, see *Resource-Types* in How Policies Work.

Aggregate Resource-Type

- `database-family`

Individual Resource-Types

- `db-systems`
- `db-nodes`
- `db-homes`
- `databases`
- `pluggable databases`
- `db-backups`

Supported Variables

Only the general variables are supported. For more information, see *General Variables for All Requests* in Policy Reference.

Details for Verb + Resource-Type Combinations

The following tables show the permissions and API operations covered by each verb. 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` and `use` verbs for the `db-systems` resource-type cover no extra permissions or API operations compared to the `inspect` verb. However, the `manage` verb includes two more permissions and partially covers two more API operations.

Copyright © 2022, 2026, Oracle and/or its affiliates.

db-systems

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered

inspect	DB_SYSTEM_INSPECT	`ListDbSystems`
		`GetDbSystem`
		`ListDbSystemPatches`
		`ListDbSystemPatchHistoryEntries`
		`GetDbSystemPatch`
		`GetDbSystemPatchHistoryEntry`
			none

read	no extra	no extra	none
use	DB_SYSTEM_UPDATE	no extra	`ChangeDbSystemCompartment` (also needs `use db-homes`, `use databases`, and `inspect db-backups`)
manage	USE + DB_SYSTEM_CREATE DB_SYSTEM_DELETE	`UpdateDBSystem`	`LaunchDBSystem`, `TerminateDbSystem` (both also need `manage db-homes`, `manage databases`, `use vnics`, and `use subnets`)

db-nodes

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DB_NODE_INSPECT DB_NODE_QUERY	`GetDbNode`	none
read	no extra	no extra	none
use	no extra	no extra	none
manage	USE + DB_NODE_POWER_ACTIONS	`DbNodeAction`	none

db-homes

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered

inspect	DB_HOME_INSPECT	`ListDBHome`
		`GetDBHome`
		`ListDbHomePatches`
		`ListDbHomePatchHistoryEntries`
		`GetDbHomePatch`
		`GetDbHomePatchHistoryEntry`
			none

read	no extra	no extra	none

Copyright © 2022, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
use	DB_HOME_UPDATE	`UpdateDBHome`	`ChangeDbSystemCompartment` (also needs `use db-systems`, `use databases`, and `inspect db-backups`)
manage	USE + DB_HOME_CREATE DB_HOME_DELETE	no extra	`LaunchDBSystem`, `TerminateDbSystem` (both also need `manage db-systems`, `manage databases`, `use vnics`, and `use subnets`). If automatic backups are enabled on the default database, also needs `manage db-backups`. `CreateDbHome`, (also needs `use db-systems` and `manage databases`). If creating the Database Home by restoring from a backup, also needs `read db-backups`. `DeleteDbHome`, (also needs `use db-systems` and `manage databases`). If automatic backups are enabled on the default database, also needs `manage db-backups`. If the `performFinalBackup` option is selected, also needs `manage db-backups` and `read databases`.

databases

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered

inspect	DATABASE_INSPECT	`ListDatabases`
		`GetDatabase`
		`ListDataGuardAssociations`
		`GetDataGuardAssociation`
			none

read	no extra DATABASE_CONTENT_READ	no extra	none

use	READ + DATABASE_CONTENT_WRITE DATABASE_UPDATE	`UpdateDatabase`
		`SwitchoverDataGuardAssociation`
		`FailoverDataGuardAssociation`
		`ReinstateDataGuardAssociation`
		`RotateVaultKey`
		`MigrateVaultKey`
			`CreateDataGuardAssociation`
			`ChangeDbSystemCompartment` (also needs `use db-systems`, `use db-homes`, and `inspect db-backups`)

Copyright © 2022, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
manage	USE + DATABASE_CREATE DATABASE_DELETE	no extra	`LaunchDBSystem`, `TerminateDbSystem` (both also need `manage db-systems`, `manage db-homes`, `use vnics`, and `use subnets`)

pluggable databases

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	PLUGGABLE_DATABASE_INSPECT	`ListPluggableDatabases` `GetPluggableDatabase`	none
read	INSPECT + PLUGGABLE_DATABASE_CONTENT_READ	no extra	none

use	READ +
	PLUGGABLE_DATABASE_CONTENT_WRITE
	PLUGGABLE_DATABASE_UPDATE
		`UpdatePluggableDatabases`
		`StartPluggableDatabase`
		`StopPluggableDatabase`
			none

manage	USE + PLUGGABLE_DATABASE_CREATE PLUGGABLE_DATABASE_DELETE	no extra	`CreatePluggableDatabase`, `DeletePluggableDatabase`, `LocalClonePluggableDatabase`, `RemoteClonePluggableDatabase` (all also need `use databases`)

db-backups

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DB_BACKUP_INSPECT	`GetBackup` `ListBackups`	`ChangeDbSystemCompartment` (also needs `use db-systems`, `use db-homes`, and `use databases`)
read	INSPECT + DB_BACKUP_CONTENT_READ	none	`RestoreDatabase` (also needs `use databases`)
use	no extra	no extra	none
manage	USE + DB_BACKUP_CREATE DB_BACKUP_DELETE	`DeleteBackup`	`CreateBackup` (also needs `read databases`)

For more information on permissions and verbs, see Advanced Policy Features.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Permissions Required for Each API Operation

The following tables list the API operations for DB systems and pluggable databases in a logical order, grouped by resource type.

Database API Operations

API operation	Permissions required to use the operation
`ListDbSystems`	DB_SYSTEM_INSPECT
`GetDbSystem`	DB_SYSTEM_INSPECT
`LaunchDbSystem`	DB_SYSTEM_CREATE and DB_HOME_CREATE and DATABASE_CREATE and VNIC_CREATE and VNIC_ATTACH and SUBNET_ATTACH To enable automatic backups for the initial database, also need DB_BACKUP_CREATE and DATABASE_CONTENT_READ
`UpdateDbSystem`	DB_SYSTEM_INSPECT and DB_SYSTEM_UPDATE
`ChangeDbSystemCompartment`	DB_SYSTEM_UPDATE and DB_HOME_UPDATE and DATABASE_UPDATE and DB_BACKUP_INSPECT
`ListDbSystemPatches`	DB_SYSTEM_INSPECT
`ListDbSystemPatchHistoryEntries`	DB_SYSTEM_INSPECT
`GetDbSystemPatch`	DB_SYSTEM_INSPECT
`GetDbSystemPatchHistoryEntry`	DB_SYSTEM_INSPECT
`TerminateDbSystem`	DB_SYSTEM_DELETE and DB_HOME_DELETE and DATABASE_DELETE and VNIC_DETACH and VNIC_DELETE and SUBNET_DETACH If automatic backups are enabled for any database in the DB System, also need DB_BACKUP_DELETE
`GetDbNode`	DB_NODE_INSPECT
`DbNodeAction`	DB_NODE_POWER_ACTIONS
`ListDbHomes`	DB_HOME_INSPECT
`GetDbHome`	DB_HOME_INSPECT
`ListDbHomePatches`	DB_HOME_INSPECT
`ListDbHomePatchHistoryEntries`	DB_HOME_INSPECT
`GetDbHomePatch`	DB_HOME_INSPECT
`GetDbHomePatchHistoryEntry`	DB_HOME_INSPECT
`CreateDbHome`	DB_SYSTEM_INSPECT and DB_SYSTEM_UPDATE and DB_HOME_CREATE and DATABASE_CREATE To enable automatic backups for the database, also need DB_BACKUP_CREATE and DATABASE_CONTENT_READ
`UpdateDbHome`	DB_HOME_UPDATE
`DeleteDbHome`	DB_SYSTEM_UPDATE and DB_HOME_DELETE and DATABASE_DELETE If automatic backups are enabled, also need DB_BACKUP_DELETE If performing a final backup on termination, also need DB_BACKUP_CREATE and DATABASE_CONTENT_READ

Copyright © 2022, 2026, Oracle and/or its affiliates.

API operation	Permissions required to use the operation
`ListDatabases`	DATABASE_INSPECT
`GetDatabase`	DATABASE_INSPECT
`UpdateDatabase`	DATABASE_UPDATE To enable automatic backups, also need DB_BACKUP_CREATE and DATABASE_CONTENT_READ
`ListDbSystemShapes`	(no permissions required; available to anyone)
`ListDbVersions`	(no permissions required; available to anyone)
`GetDataGuardAssociation`	DATABASE_INSPECT
`ListDataGuardAssociations`	DATABASE_INSPECT
`CreateDataGuardAssociation`	DB_SYSTEM_UPDATE and DB_HOME_CREATE and DB_HOME_UPDATE and DATABASE_CREATE and DATABASE_UPDATE
`SwitchoverDataGuardAssociation`	DATABASE_UPDATE
`FailoverDataGuardAssociation`	DATABASE_UPDATE
`ReinstateDataGuardAssociation`	DATABASE_UPDATE
`MigrateVaultKey`	DATABASE_UPDATE
`RotateVaultKey`	DATABASE_UPDATE
`GetBackup`	DB_BACKUP_INSPECT
`ListBackups`	DB_BACKUP_INSPECT
`CreateBackup`	DB_BACKUP_CREATE and DATABASE_CONTENT_READ
`DeleteBackup`	DB_BACKUP_DELETE and DB_BACKUP_INSPECT
`RestoreDatabase`	DB_BACKUP_INSPECT and DB_BACKUP_CONTENT_READ and DATABASE_CONTENT_WRITE

Pluggable Database API Operations

API operation	Permissions required to use the operation
`ListPluggableDatabase`	PLUGGABLE_DATABASE_INSPECT
`GetPluggableDatabase`	PLUGGABLE_DATABASE_INSPECT
`CreatePluggableDatabase`	DATABASE_INSPECT* DATABASE_UPDATE* PLUGGABLE_DATABASE_CREATE Additional permissions required if auto-backups are enabled on the CDB and includes this PDB: PLUGGABLE_DATABASE_CONTENT_READ
`UpdatePluggableDatabase`	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE Additional permissions required if auto-backups are enabled on the CDB and includes this PDB: PLUGGABLE_DATABASE_CONTENT_READ
`StartPluggableDatabase`	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE

Copyright © 2022, 2026, Oracle and/or its affiliates.

API operation	Permissions required to use the operation
`StopPluggableDatabase`	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE
`DeletePluggableDatabase`	DATABASE_INSPECT (exists) DATABASE_UPDATE (exists) PLUGGABLE_DATABASE_DELETE
`LocalClonePluggableDatabase`	DATABASE_INSPECT* DATABASE_UPDATE* PLUGGABLE_DATABASE_INSPECT PLUGGABLE_DATABASE_UPDATE PLUGGABLE_DATABASE_CONTENT_READ PLUGGABLE_DATABASE_CREATE PLUGGABLE_DATABASE_CONTENT_WRITE
`RemoteClonePluggableDatabase`	DATABASE_INSPECT* DATABASE_UPDATE* PLUGGABLE_DATABASE_INSPECT PLUGGABLE_DATABASE_UPDATE PLUGGABLE_DATABASE_CONTENT_READ PLUGGABLE_DATABASE_CREATE PLUGGABLE_DATABASE_CONTENT_WRITE

For more information on permissions and verbs, see Advanced Policy Features.

Transparent Data Encryption

This article describes Transparent Data Encryption for Base Database Service.

All user-created tablespaces in a DB system database are encrypted by default, using Transparent Data Encryption (TDE).

- On a 1- or 2-node RAC DB system, you can use the TDE commands to update the master encryption key for a database.
- You must create and activate a master encryption key for any PDBs that you create. After creating or plugging in a new PDB on a 1- or 2-node RAC DB System, use the `dbcli update-tdekey` command to create and activate a master encryption key for the PDB. Otherwise, you might encounter the error `ORA-28374: typed master key not found in wallet` when attempting to create tablespaces in the PDB. In a multitenant environment, each PDB has its own master encryption key which is stored in a single keystore used by all containers.
- For version 12c databases, if you don't want your tablespaces encrypted, you can set the `ENCRYPT_NEW_TABLESPACES` database initialization parameter to `DDL`.

For more information about:

- changing an existing TDE wallet password using the OCI Console, see *Manage Administrator and TDE Wallet Passwords*.
- TDE commands, see TDE Commands.
- TDE, see Quick-start Setup Guide for Wallet Based Transparent Data Encryption.

Copyright © 2022, 2026, Oracle and/or its affiliates.

7

Reference

Oracle Database CLI Reference

The Database CLI (`dbcli`) is a command line interface available for Base Database Service. After you connect to the DB system in the Base Database Service, you can use the `dbcli` to perform tasks such as creating Oracle database homes and databases.

Operational Notes

- The database CLI commands must be run as the root user.
- `dbcli` is in the `/opt/oracle/dcs/bin/` directory.

  This directory is included in the path for the root user's environment.
- Oracle Database maintains logs of the `dbcli` command output in the `dcscli.log` and `dcs-agent.log` files in the `/opt/oracle/dcs/log/` directory.
- The database CLI commands and most parameters are case sensitive and should be typed as shown. A few parameters are not case sensitive, as indicated in the parameter descriptions, and can be typed in uppercase or lowercase.

Syntax

The database CLI commands use the following syntax:

```
dbcli command [parameters]
```

where:

- `command` is a verb-object combination such as `list-databases`.
- `parameters` include additional options for the command. Most parameter names are preceded with two dashes, for example, `--help`. Abbreviated parameter names are preceded with one dash, for example, `-h`.
- User-specified parameter values are shown in red text within angle brackets, for example, <*db_home_id*>. Omit the angle brackets when specifying these values.
- The help parameter is available with every command.

The remainder of this topic contains syntax and other details about the commands.

CLI Update Command

Occasionally, new commands are added to the database CLI and other commands are updated to support new features. You can use the following command to update the database CLI:

Copyright © 2022, 2026, Oracle and/or its affiliates.

cliadm update-dbcli

Use the `cliadm update-dbcli` command to update the database CLI with the latest new and updated commands.

> **Note**
>
> On RAC DB systems, execute the `cliadm update-dbcli` command on each node in the cluster.

Syntax

```
cliadm update-dbcli [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

Example

The following command updates the `dbcli`:

```
cliadm update-dbcli
```

Output:

```
{
  "jobId" : "dc9ce73d-ed71-4473-99cd-9663b9d79bfd",
  "status" : "Created",
  "message" : "Dcs cli will be updated",
  "reports" : [ ],
  "createTimestamp" : "January 18, 2017 10:19:34 AM PST",
  "resourceList" : [ ],
  "description" : "dbcli patching",
  "updatedTime" : "January 18, 2017 10:19:34 AM PST"
}
```

Agent Commands

The following command is available to manage agents:

- `dbcli ping-agent`

dbcli ping-agent

Use the `dbcli ping-agent` command to test the reachability of an agent.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Syntax

```
dbcli ping-agent [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

Autologcleanpolicy Commands

The following commands are available to manage policies for automatic cleaning (purging) of logs.

- `dbcli create-autoLogCleanPolicy`
- `dbcli list-autoLogCleanPolicy`

dbcli create-autoLogCleanPolicy

Use the `dbcli create-autoLogCleanPolicy` command to create policies for automatic cleaning (purging) of logs.

Syntax

```
dbcli create-autoLogCleanPolicy
  [-c {gi|database|dcs}]
  [-f <number>]
  [-o <number>]
  [-u {Day|Hour|Minute}]
  [-uMB <number>]
  [-uPer <number>]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-c	--components	(Optional) Components to purge. Possible values are gi, database, and dcs. Separate multiple values with commas. Example: `gi,dcs`
-f	--freeSpaceBelowPercentage	(Optional) Purges logs when the free disk space is below the specified percentage of the total partition size. Valid range: 20-50. Default: 20.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-o	--olderthan	(Optional) Quantity portion of time interval. Default: 30. Cleans logs older than the specified time interval (`-o` and `-u`).
-u	--olderThanUnit	(Optional) Unit portion of time interval. Possible values: Day, Hour, or Minute. Default: Day. Cleans logs older than the specified time interval (`-o` and `-u`).

Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameter	Full Name	Description
-uMB	--usageOverMB	(Optional) Purges logs when log usage exceeds the specified number of MegaBytes (MB). Valid range: 10 to 50% of total partition size.
-uPer	--usageOverPercentage	(Optional) Purges logs when log usage exceeds the specified percentage of the total partition size. Valid range: 10-50.

dbcli list-autoLogCleanPolicy

Use the `dbcli list-autoLogCleanPolicy` command to list policies for automatic cleaning of logs.

Syntax

```
dbcli list-autoLogCleanPolicy
  [-c {gi|database|dcs}]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-c	--components	(Optional) Components. Possible values are gi, database, and dcs. Separate multiple values with commas. Example: `gi,dcs`
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

Backup Commands

The following commands are available to back up databases:

- `dbcli create-backup`
- `dbcli getstatus-backup`
- `dbcli schedule-backup`

> **Note**
>
> Instead of using `dbcli`, you can use the Console or the API to manage backing up the databases your DB system to the Object Storage. However, if you switch from using `dbcli` to using managed backups, a new backup configuration is created and associated with your database, and backups you created by using `dbcli` will not be accessible from the managed backup interfaces. For information about managed backups, see *Back Up and Recovery in Base Database Service*.

Before you can back up a database by using the `dbcli create-backup` command, you'll need to:

1. Create a backup configuration by using the `dbcli create-backupconfig` command.

Copyright © 2022, 2026, Oracle and/or its affiliates.

2. Associate the backup configuration with the database by using the `dbcli update-database` command.

After a database is associated with a backup configuration, you can use the `dbcli create-backup` command in a `cron` job to run backups automatically. You can use a cron utility such as CronMaker to help build expressions. For more information, see CronMaker.

dbcli create-backup

Use the `dbcli create-backup` command to create a backup of a database.

Syntax

```
dbcli create-backup
  -in <db_name>
  -i <db_id>
  [-bt {Regular-L0|Regular-L1|Longterm|ArchiveLog}]
  [-c {Database|TdeWallet}]
  [-k <n>]
  [-t <tag>]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-bt	--backupType	(Optional) Backup type. Possible values are Regular-L0, Regular-L1, Longterm, and ArchiveLog. Regular-L0 and Regular L1 correspond to incremental L0 and L1 backups. Longterm corresponds to Full backup. ArchiveLog corresponds to archived redo logs backup. The default value is Regular-L1. Values are not case-sensitive. If omitted, the default value is used.
-c	--component	(Optional) Component. Possible values are Database and TdeWallet. The default value is Database. The value TdeWallet backs up TDE wallets. Values are not case-sensitive. If omitted, the default value is used. Note that the TDE wallets are automatically backed up in the following situations: • A database is created with an Object Storage backup configuration. • A database that has an Object Storage backup configuration is updated. • An Object Storage backup configuration is updated. • A backup of the type Longterm is created. • The TDE key for a database is rotated. • A database is backed up and no TDE wallet backups exist yet.
-h	--help	(Optional) Displays help for using the command.
-i	--dbid	The ID of the database to back up. Use the `dbcli list-databases` command to get the database's ID.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameter	Full Name	Description
-in	--dbName	The name of the database to back up. Use the `dbcli list-databases` command to get the database's name.
-j	--json	(Optional) Displays JSON output.
-k	--keepDays	(Optional) Specifies the time until which the backup or copy must be kept. After this time the backup is obsolete, regardless of the backup retention policy settings. For Longterm backup type only.
-t	--tag	(Required for Longterm backup type) Specifies a user-specified tag name for a backup set and applies this tag to the output files generated by the command. This value is not case sensitive. Valid number of characters: 1 to 30. The characters are limited to the characters that are valid in file names on the target file system. For example, ASM does not support the use of the hyphen (-) character in the file names it uses internally, so weekly-incremental is not a valid tag name for backups in ASM disk groups. Environment variables are not valid in the TAG parameter.

Examples

The following command creates a backup of the specified database using the database ID.

```
dbcli create-backup -i 573cadb2-0cc2-4c1c-9c31-595ab8963d5b
```

The following command creates a backup of the specified database using the database name ("mydb").

```
dbcli create-backup -in mydb
```

dbcli getstatus-backup

Use the `dbcli getstatus-backup` command to display the status of a backup.

Syntax

```
dbcli getstatus-backup
  -t <backup_type>
  [i <id>]
  [-in <name>]
  [-l] [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--dbId	(Optional) Database Resource ID.
-in	--dbName	(Optional) Database Resource Name.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameter	Full Name	Description
-j	--json	(Optional) Displays JSON output.
-l	--isLatestBackupReport	(Optional) Latest backup report. Default: true.
-t	--backupType	Backup type.

dbcli schedule-backup

Use the `dbcli schedule-backup` command to schedule a backup of a database.

Syntax

```
dbcli schedule-backup
  -t <backup_type>
  -f <number>
  [i <id>]
  [-in <name>]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-f	--frequency	Frequency in minutes.
-h	--help	(Optional) Displays help for using the command.
-i	--dbId	(Optional) Database Resource ID.
-in	--dbName	(Optional) Database Resource Name.
-j	--json	(Optional) Displays JSON output.
-t	--backupType	Backup type.

Backupconfig Commands

A backup configuration determines the backup destination and recovery window for database backups. You create the backup configuration and then associate it with a database by using the `dbcli update-database` command.

> **Note**
>
> Backups that were configured using the Console may become unusable if you make changes using these commands. For backups configured using the Console, use these commands with support guidance only.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> Instead of using `dbcli`, you can use the Console or the API to manage backing up the databases in your DB system to the Object Storage. For information about managed backups, see *Back Up and Recovery in Base Database Service*.

After a database is associated with a backup configuration, you can use the `dbcli create-backup` command in a `cron` job to run backups automatically. You can use a cron utility such as CronMaker to help build expressions. For more information, see CronMaker.

The following commands are available to manage backup configurations:

- `dbcli create-backupconfig`
- `dbcli list-backupconfigs`
- `dbcli describe-backupconfig`
- `dbcli update-backupconfig`
- `dbcli delete-backupconfig`

dbcli create-backupconfig

Use the `dbcli create-backupconfig` command to create a backup configuration that defines the backup destination and recovery windows.

Syntax

```
dbcli create-backupconfig
  -d {DISK|OBJECTSTORE|NONE}
  -c <bucket>
  -o <object_store_swift_id>
  -on <object_store_swift_name>
  -w <n>
  -n <name>
  [-cr|-no-cr]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-c	--container	The name of an existing bucket in the Oracle Cloud Infrastructure Object Storage service. You can use the Console or the Object Storage API to create the bucket. For more information, see Managing Buckets. You must also specify `--backupdestination objectstore` and the `--objectstoreswiftId` parameter.
-cr -no-cr	--crosscheck --no-crosscheck	(Optional) Indicates whether to enable the crosscheck operation. This operation determines if the files on the disk or in the media management catalog correspond to data in the RMAN repository. If omitted, the default setting is used (crosscheck is enabled by default).

Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameter	Full Name	Description
-d	--backupdestination	The backup destination as one of the following (these values are **not** case sensitive): DISK - The local Fast Recovery Area. OBJECTSTORE - The Oracle Cloud Infrastructure Object Storage service. You must also specify the `--container` and `--objectstoreswiftId` parameters. NONE - Disables the backup.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-n	--name	The name of the backup configuration.
-o	--objectstoreswiftId	The ID of the object store that contains the endpoint and credentials for the Oracle Cloud Infrastructure Object Storage service. Use the `dbcli list-objectstoreswifts` command to get the object store ID. Use the `dbcli create-objectstoreswift` command to create an object store. You must also specify `--backupdestination objectstore` and the `--container` parameter.
-on	--objectstoreswiftName	The name of the object store that contains the endpoint and credentials for the Oracle Cloud Infrastructure Object Storage service. Use the `dbcli list-objectstoreswifts` command to get the object store ID. Use the `dbcli create-objectstoreswift` command to create an object store. You must also specify `--backupdestination objectstore` and the `--container` parameter.
-w	--recoverywindow	The number of days for which backups and archived redo logs are maintained. The interval always ends with the current time and extends back in time for the number of days specified. For a DISK backup destination, specify 1 to 14 days. For an OBJECTSTORE backup destination, specify 1 to 30 days.

Example

The following command creates a backup configuration named 'dbbkcfg1':

```
dbcli create-backupconfig -d Disk -w 7 -n dbbkcfg1
```

Output:

```
{
  "jobId" : "4e0e6011-db53-4142-82ef-eb561658a0a9",
  "status" : "Success",
  "message" : null,
  "reports" : [ {
    "taskId" : "TaskParallel_919",
    "taskName" : "persisting backup config metadata",
    "taskResult" : "Success",
    "startTime" : "November 18, 2016 20:21:25 PM UTC",
    "endTime" : "November 18, 2016 20:21:25 PM UTC",
    "status" : "Success",
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
    "taskDescription" : null,
    "parentTaskId" : "TaskSequential_915",
    "jobId" : "4e0e6011-db53-4142-82ef-eb561658a0a9",
    "tags" : [ ],
    "reportLevel" : "Info",
    "updatedTime" : "November 18, 2016 20:21:25 PM UTC"
  } ],
  "createTimestamp" : "November 18, 2016 20:21:25 PM UTC",
  "description" : "create backup config:dbbkcfg1",
  "updatedTime" : "November 18, 2016 20:21:25 PM UTC"
}
```

dbcli list-backupconfigs

Use the `dbcli list-backupconfigs` command to list all the backup configurations in the DB system.

Syntax

```
dbcli list-backupconfigs [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

Example

The following command lists a backup configuration:

```
dbcli list-backupconfigs
```

Output:

```
ID                                       Name
RecoveryWindow     BackupDestination CreateTime
---------------------------------------- --------------------
------------------ ----------------- --------------------------
ccdd56fe-a40b-4e82-b38d-5f76c265282d     dbbkcfg1
7                  Disk              July 10, 2016 12:24:08 PM UTC
```

dbcli describe-backupconfig

Use the `dbcli describe-backupconfig` command to show details about a specific backup configuration.

Syntax

```
dbcli describe-backupconfig -i <id> -in <name> [-h] [-j]
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--backupconfigid	The backup configuration ID. Use the `dbcli list-backupconfigs` command to get the ID.
-in	--backupconfigname	The backup configuration name. Use the `dbcli list-backupconfigs` command to get the name.
-j	--json	(Optional) Displays JSON output.

Example

The following command displays details about a backup configuration:

```
dbcli describe-backupconfig -i ccdd56fe-a40b-4e82-b38d-5f76c265282d
```

Output:

```
Backup Config details
----------------------------------------------------------------
                     ID: ccdd56fe-a40b-4e82-b38d-5f76c265282d
                   Name: dbbkcfg1
         RecoveryWindow: 7
      BackupDestination: Disk
            CreatedTime: July 10, 2016 12:24:08 PM UTC
            UpdatedTime: July 10, 2016 12:24:08 PM UTC
```

dbcli update-backupconfig

Use the `dbcli update-backupconfig` command to update an existing backup configuration.

Syntax

```
dbcli update-backupconfig
  -i <id>
  -in <name>
  -w <n>
  -d {DISK|OBJECTSTORE|NONE}
  -c <bucket>
  -o <object_store_swift_id>
  -on <object_store_swift_name>
  [-cr|-no-cr]
  [-h] [-j]
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameters

Parameter	Full Name	Description
-c	--container	The name of an existing bucket in the Oracle Cloud Infrastructure Object Storage service. You can use the Console or the Object Storage API to create the bucket. For more information, see Managing Buckets. You must also specify `--backupdestination objectstore` and the `--objectstoreswiftId` parameter.
-cr -no-cr	--crosscheck --no-crosscheck	(Optional) Indicates whether to enable the crosscheck operation. This operation determines if the files on the disk on in the media management catalog correspond to data in the RMAN repository. If omitted, the default setting is used (crosscheck is enabled by default).
-h	--help	(Optional) Displays help for using the command.
-i	--backupconfigid	The ID of the backup configuration to update. Use the `dbcli list-backupconfigs` command to get the ID.
-in	--backupconfigname	The name of the backup configuration to update. Use the `dbcli list-backupconfigs` command to get the name.
-j	--json	(Optional) Displays JSON output.
-o	--objectstoreswiftId	The ID of the object store that contains the endpoint and credentials for the Oracle Cloud Infrastructure Object Storage service. Use the `dbcli list-objectstoreswifts` command to get the object store ID. Use the `dbcli create-objectstoreswift` command to create an object store. You must also specify `--backupdestination objectstore` and the `--container` parameter.
-on	--objectstoreswiftname	The name of the object store that contains the endpoint and credentials for the Oracle Cloud Infrastructure Object Storage service. Use the `dbcli list-objectstoreswifts` command to get the object store ID. Use the `dbcli create-objectstoreswift` command to create an object store. You must also specify `--backupdestination objectstore` and the `--container` parameter.
-w	--recoverywindow	The new disk recovery window. For a DISK backup destination, specify 1 to 14 days. For an OBJECTSTORE backup destination, specify 1 to 30 days.

Example

The following command updates the recovery window for a backup configuration:

```
dbcli update-backupconfig -i ccdd56fe-a40b-4e82-b38d-5f76c265282d -w 5
```

Output:

```
{
  "jobId" : "0e849291-e1e1-4c7a-8dd2-62b522b9b807",
  "status" : "Created",
  "message" : null,
  "reports" : [ ],
  "createTimestamp" : 1468153731699,
  "description" : "update backup config: dbbkcfg1",
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
  "updatedTime" : 1468153731700
}
```

dbcli delete-backupconfig

Use the `dbcli delete-backupconfig` command to delete a backup configuration.

Syntax

```
dbcli delete-backupconfig -i <id> -in <name> [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--id	The backup configuration ID to delete. Use the `dbcli list-backupconfigs` command to get the ID.
-in	--backupconfigname	The name of the backup configuration to delete. Use the `dbcli list-backupconfigs` command to get the name.
-j	--json	(Optional) Displays JSON output.

Example

The following command deletes the specified backup configuration:

```
dbcli delete-backupconfig -i ccdd56fe-a40b-4e82-b38d-5f76c265282d
```

Component Command

dbcli describe-component

Your DB system might not include this newer command. If you have trouble running the command, use the `CLI Update Command` command to update the database CLI and then retry the command.

> **Note**
>
> The `dbcli describe-component` command is not available on 2-node RAC DB systems. Patching 2-node systems from Object Storage is not supported.

Use the `dbcli describe-component` command to show the installed and available patch versions for the server, storage, and/or database home components in the DB system.

This command requires a valid Object Storage credentials configuration. If the configuration is missing or invalid, the command fails with the error: `Failed to connect to the object store. Please provide valid details.`

For more information about updating the CLI, creating the credentials configuration, and applying patches, see *Update a DB System*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Syntax

```
dbcli describe-component
  [-s <server_group>]
  [-d <db_group>]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-d	--dbhomes	(Optional) Lists the installed and available patch versions for only the database home components.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-s	--server	(Optional) Lists the installed and available patch versions for only the server components.

Example

The following command to show the current component versions and the available patch versions in the object store:

```
dbcli describe-component
```

Output:

```
System Version
--------------
12.1.2.10.0

Component                                 Installed Version    Available
Version
----------------------------------------- -------------------
--------------------
OAK                                       12.1.2.10.0          up-to-date
GI                                        12.1.0.2.161018      up-to-date
ORADB12102_HOME1                          12.1.0.2.161018      up-to-date
ORADB12102_HOME2, ORADB12102_HOME3        12.1.0.2.160719      12.1.0.2.161018
```

Database Commands

The following commands are available to manage databases:

- `dbcli clone-database`
- `dbcli describe-database`
- `dbcli list-databases`
- `dbcli modify-database`
- `dbcli recover-database`

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `dbcli update-database`

dbcli clone-database

Use the `dbcli clone-database` command to clone a database.

Syntax

```
dbcli clone-database
  -f <name>
  -u <name>
  -n <name>
  [-s <shape>]
  [-t <type>]
  [m <sys_password>]
  [-p <tde_password>]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-f	--sourcedbname	Source database name.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-m	--syspassword	(Optional) Password for SYS.
-n	--dbname	Database name.
-p	--tdepassword	(Optional) Password for source TDE wallet.
-s	--dbshape	(Optional) Database shape. Examples: odb1, odb2.
-t	--dbtype	(Optional) Database Type: SI
-u	--databaseUniqueName	Database unique name.

dbcli describe-database

Use the `dbcli describe-database` command to display database details.

Syntax

```
dbcli describe-database
  -i <db_id>
  -in <db_name>
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--dbid	The ID of the database to display. Use the `dbcli list-databases` command to get the database ID.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameter	Full Name	Description
-in	--dbName	The name of the database to display. Use the `dbcli list-databases` command to get the database name.
-j	--json	(Optional) Displays JSON output.

Example

The following command displays information for a database named b727bf80-c99e-4846-ac1f-28a81a725df6:

```
dbcli describe-dbhome -i b727bf80-c99e-4846-ac1f-28a81a725df6
```

Output:

```
DB Home details
------------------------------------------------------------
                     ID: b727bf80-c99e-4846-ac1f-28a81a725df6
                   Name: OraDB12102_home1
                Version: 12.1.0.2
          Home Location: /u01/app/orauser/product/12.1.0.2/dbhome_1
                Created: Jun 2, 2016 10:19:23 AM
```

dbcli list-databases

Use the `dbcli list-databases` command to list all databases on the DB system.

Syntax

```
dbcli  list-databases [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

Example

The following command displays a list of databases:

```
dbcli  list-databases
```

Output:

```
ID                                       DB Name    DB Version
CDB        Class    Shape    Storage    Status
---------------------------------------- ---------- --------------------
---------- -------- -------- ---------- ----------
80ad855a-5145-4f8f-a08f-406c5e4684ff     dbst       12.1.0.2
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
true       OLTP     odb2     ACFS       Configured
6f4e36ae-120b-4436-b0bf-d0c4aef9f7c9     db11tsta   11.2.0.4
false      OLTP     odb1     ACFS       Configured
d8e31790-84e6-479c-beb0-ef97207091a2     db11tstb   11.2.0.4
false      OLTP     odb1     ACFS       Configured
cce096c7-737b-447a-baa1-f4c2a330c030     pdbtst     12.1.0.2
true       OLTP     odb1     ACFS       Configured
```

The following command displays the JSON output for a database:

```
dbcli list-databases -j
```

Output:

```
[ {
  "id" : "80ad855a-5145-4f8f-a08f-406c5e4684ff",
  "name" : "dbtst",
  "dbName" : "dbtst",
  "databaseUniqueName" : "dbtst_phx1cs",
  "dbVersion" : "12.1.0.2",
  "dbHomeId" : "2efe7af7-0b70-4e9b-ba8b-71f11c6fe287",
  "instanceOnly" : false,
  "registerOnly" : false,
  "dbId" : "167525515",
  "isCdb" : true,
  "pdBName" : "pdb1",
  "pdbAdminUserName" : "pdbuser",
  "enableTDE" : true,
  "dbType" : "SI",
  "dbTargetNodeNumber" : "0",
  "dbClass" : "OLTP",
  "dbShape" : "odb2",
  "dbStorage" : "ACFS",
  "dbCharacterSet" : {
    "characterSet" : "US7ASCII",
    "nlsCharacterset" : "AL16UTF16",
    "dbTerritory" : "AMERICA",
    "dbLanguage" : "AMERICAN"
  },
  "dbConsoleEnable" : false,
  "backupConfigId" : null,
  "backupDestination" : "NONE",
  "cloudStorageContainer" : null,
  "state" : {
    "status" : "CONFIGURED"
  },
  "createTime" : "November 09, 2016 17:23:05 PM UTC",
  "updatedTime" : "November 09, 2016 18:00:47 PM UTC"
}
```

dbcli modify-database

Use the `dbcli modify-database` command to modify a database.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Syntax

```
dbcli modify-database
  -i <db_id>
  -dh <destination_db_home_id>
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-dh	--destdbhomeid	Destination database home ID.
-h	--help	(Optional) Displays help for using the command.
-i	--databaseid	Database ID.
-j	--json	(Optional) Displays JSON output.

dbcli recover-database

Use the `dbcli recover-database` command to recover a database.

Syntax

```
dbcli recover-database
  [-br <json>]
  [-in <db_name>]
  [-i <db_id>]
  [-r <time>]
  [-t {Latest|PITR|SCN}]
  [-s]
  [-l <location>]
  [-tp <tde_password>]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-br	--backupReport	(Optional) JSON input for backup report.
-h	--help	(Optional) Displays help for using the command.
-i	--dbid	(Optional) Database resource ID.
-in	--dbName	(Optional) Database name.
-j	--json	(Optional) Displays JSON output.
-l	--tdeWalletLocation	(Optional) TDE wallet backup location. TDE wallet should be backed up in tar.gz format.
-r	--recoveryTimeStamp	(Required when recovery type is PITR) Recovery timestamp in the format mm/dd/yyyy hh:mi:ss. Default: [ ]
-s	--scn	(Required when recovery type is SCN) SCN.
-t	--recoverytype	(Required when backup report is provided) Recovery type. Possible values are Latest, PITR, and SCN.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameter	Full Name	Description
-tp	--tdeWalletPassword	(Optional) TDE wallet password.

dbcli update-database

Use the `dbcli update-database` command to associate a backup configuration with a database.

Syntax

```
dbcli update-database
  -i <db_id>
  -bi <bkup_config_id>
  -bin <bkup_config_name>;
  [-id <id>]
  -in <name>
  [-no-ab]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-bi	--backupconfigid	Defines the backup configuration ID. Use the `dbcli list-backupconfigs` command to get the ID.
-bin	--backupconfigname	Defines the backup configuration name for future use. Use the `dbcli list-backupconfigs` command to get the name.
-id	--databaseid	(Optional.) Specifies the DBID, which is a unique 32-bit identification number computed when the database is created. RMAN displays the DBID upon connection to the target database. You can obtain the DBID by querying the V$DATABASE view or the RC_DATABASE and RC_DATABASE_INCARNATION recovery catalog views.
-in	--dbName	Defines the database name to be updated. Use the `dbcli list-databases` command to get the database name.
-h	--help	(Optional) Displays help for using the command.
-i	--dbid	Defines the database ID to be updated. Use the `dbcli list-databases` command to get the database ID.
-j	--json	(Optional) Displays JSON output.
-no-ab	--noautobackup	(Optional) Disables automatic backups for the specified database. Note that, once disabled, automatic backup cannot be re-enabled using the CLI. To re-enable automatic backup, use the Console.

Example

The following command associates a backup configuration file with a database:

```
dbcli update-database -bi 78a2a5f0-72b1-448f-bd86-cf41b30b64ee -i
71ec8335-113a-46e3-b81f-235f4d1b6fde
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Output:

```
{
  "jobId" : "2b104028-a0a4-4855-b32a-b97a37f5f9c5",
  "status" : "Created",
  "message" : null,
  "reports" : [ ],
  "createTimestamp" : 1467775842977,
  "description" : "update database id:71ec8335-113a-46e3-b81f-235f4d1b6fde",
  "updatedTime" : 1467775842978
}
```

Dbhome Commands

The following commands are available to manage database homes:

- `dbcli create-dbhome`
- `dbcli describe-dbhome`
- `dbcli delete-dbhome`
- `dbcli list-dbhomes`
- `dbcli update-dbhome`

dbcli create-dbhome

Use the `dbcli create-dbhome` command to create an Oracle Database Home.

Syntax

```
dbcli create-dbhome -v <version> [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-v	--version	Defines the Database Home version.

Example

The following command creates an Oracle Database Home version 12.1.0.2:

```
dbcli create-dbhome -v 12.1.0.2
```

dbcli describe-dbhome

Use the `dbcli describe-dbhome` command to display Oracle Database Home details.

Syntax

```
dbcli describe-dbhome -i <db_home_id> [-h] [-j]
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--dbhomeid	Identifies the database home ID. Use the `dbcli list-dbhomes` command to get the ID.
-j	--json	(Optional) Displays JSON output.

Example

The following output is an example of using the display Oracle Database Home details command.

```
dbcli describe-dbhome -i 52850389-228d-4397-bbe6-102fda65922b
```

Output:

```
DB Home details
------------------------------------------------------------
                     ID: 52850389-228d-4397-bbe6-102fda65922b
                   Name: OraDB12102_home1
                Version: 12.1.0.2
          Home Location: /u01/app/oracle/product/12.1.0.2/dbhome_1
                Created: June 29, 2016 4:36:31 AM UTC
```

dbcli delete-dbhome

Use the `dbcli delete-dbhome` command to delete a database home from the DB system.

Syntax

```
dbcli delete-dbhome -i <db_home_id> [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--dbhomeid	Identifies the database home ID to be deleted. Use the `dbcli list-dbhomes` command to get the ID.
-j	--json	(Optional) Displays JSON output.

dbcli list-dbhomes

Use the `dbcli list-dbhomes` command to display a list of Oracle Home directories.

Syntax

```
dbcli list-dbhomes [-h] [-j]
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameter

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

Example

The following command displays a list of Oracle Home directories.

```
dbcli  list-dbhomes
```

Output:

```
ID                                       Name                DB Version Home
Location
---------------------------------------- ----------------  ----------
----------------------------------------
b727bf80-c99e-4846-ac1f-28a81a725df6 OraDB12102_home1     12.1.0.2  /u01/app/
orauser/product/12.1.0.2/dbhome_1
```

dbcli update-dbhome

> **Note**
>
> Your DB system might not include this newer command. If you have trouble running the command, use the `CLI Update` command to update the database CLI and then retry the command.

Use the `dbcli update-dbhome` command to apply the DBBP bundle patch to a database home. For more information about applying patches, see *Update a DB System*.

Syntax

```
dbcli update-dbhome
  -i <db_home_id>
  -n <node>
  [--local]
  [--precheck]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--dbhomeid	The ID of the database home. Use the `dbcli list-dbhomes` command to get the ID.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameter	Full Name	Description
-j	--json	(Optional) Displays JSON output.
-n	--node	(Optional) Node number to be updated. Use the `dbcli list-nodes` command to get the node number.
-l	--local	(Optional) Performs the operation on the local node of a multi-node high availability (HA) system. This parameter is not needed to perform the operation on a single-node system.
-p	--precheck	(Optional) Runs precheck operations to check prerequisites.

Example

The following commands update the database home and show the output from the update job:

```
dbcli update-dbhome -i e1877dac-a69a-40a1-b65a-d5e190e671e6
```

Output:

```
{
  "jobId" : "493e703b-46ef-4a3f-909d-bbd123469bea",
  "status" : "Created",
  "message" : null,
  "reports" : [ ],
  "createTimestamp" : "January 19, 2017 10:03:21 AM PST",
  "resourceList" : [ ],
  "description" : "DB Home Patching: Home Id is e1877dac-a69a-40a1-b65a-
d5e190e671e6",
  "updatedTime" : "January 19, 2017 10:03:21 AM PST"
}
```

```
dbcli describe-job -i 493e703b-46ef-4a3f-909d-bbd123469bea
```

Output:

```
Job details
----------------------------------------------------------------
                     ID:  493e703b-46ef-4a3f-909d-bbd123469bea
            Description:  DB Home Patching: Home Id is e1877dac-a69a-40a1-
b65a-d5e190e671e6
                 Status:  Running
                Created:  January 19, 2017 10:03:21 AM PST
                Message:

Task Name                                Start Time
End Time                           Status
---------------------------------------- ---------------------------------
---------------------------------- ----------
Create Patching Repository Directories   January 19, 2017 10:03:21 AM PST
January 19, 2017 10:03:21 AM PST    Success
Download latest patch metadata           January 19, 2017 10:03:21 AM PST
January 19, 2017 10:03:21 AM PST    Success
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
Update System version                    January 19, 2017 10:03:21 AM PST
January 19, 2017 10:03:21 AM PST    Success
Update Patching Repository               January 19, 2017 10:03:21 AM PST
January 19, 2017 10:03:31 AM PST    Success
Opatch updation                          January 19, 2017 10:03:31 AM PST
January 19, 2017 10:03:31 AM PST    Success
Patch conflict check                     January 19, 2017 10:03:31 AM PST
January 19, 2017 10:03:31 AM PST    Running
```

Dbstorage Commands

The following commands are available to manage database storage:

- `dbcli list-dbstorages`
- `dbcli describe-dbstorage`

dbcli list-dbstorages

Use the `dbcli list-dbstorages` command to list the database storage in the DB system.

Syntax

```
dbcli list-dbstorages [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

Example

The following command displays details about database storage:

```
dbcli list-dbstorages
```

Output:

```
ID                                       Type   DBUnique Name        Status
---------------------------------------- ------ --------------------
----------
afb4a1ce-d54d-4993-a149-0f28c9fb33a4     Acfs   db1_2e56b3a9b815
Configured
d81e8013-4551-4d10-880b-d1a796bca1bc     Acfs   db11xp
Configured
```

dbcli describe-dbstorage

Use the `dbcli describe-dbstorage` command to show detailed information about a specific database storage resource.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Syntax

```
dbcli describe-dbstorage -i <db_storage_id> [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--id	Defines the database storage ID. Use the `dbcli list-dbstorages` command to get the database storage ID.
-j	--json	(Optional) Displays JSON output.

Example

The following command displays the database storage details for 105a2db2-625a-45ba-8bdd-ee46da0fd83a:

```
dbcli describe-dbstorage -i 105a2db2-625a-45ba-8bdd-ee46da0fd83a
```

Output:

```
DBStorage details
----------------------------------------------------------------

                    ID: 105a2db2-625a-45ba-8bdd-ee46da0fd83a
               DB Name: db1
            DBUnique Name: db1
          DB Resource ID: 439e7bd7-f717-447a-8046-08b5f6493df0
            Storage Type:
            DATA Location: /u02/app/oracle/oradata/db1
            RECO Location: /u03/app/oracle/fast_recovery_area/
            REDO Location: /u03/app/oracle/redo/
                   State: ResourceState(status=Configured)
                 Created: July 3, 2016 4:19:21 AM UTC
             UpdatedTime: July 3, 2016 4:41:29 AM UTC
```

Dgconfig Commands

dbcli list-dgconfigs

Use the `dbcli list-dgconfigs` command to list DG configurations.

Syntax

```
dbcli list-dgconfigs [-h] [-j]
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

Job Commands

The following commands are available to manage jobs:

- `dbcli describe-job`
- `dbcli list-jobs`

dbcli describe-job

Use the `dbcli describe-job` command to display details about a specific job.

Syntax

```
dbcli describe-job -i <job_id> [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--jobid	Identifies the job. Use the `dbcli list-jobs` command to get the jobid.
-j	--json	(Optional) Displays JSON output.

Example

The following command displays details about the specified job ID:

```
dbcli describe-job -i 74731897-fb6b-4379-9a37-246912025c17
```

Output:

```
Job details
----------------------------------------------------------------
                     ID:  74731897-fb6b-4379-9a37-246912025c17
            Description:  Backup service creation with db name: dbtst
                 Status:  Success
                Created:  November 18, 2016 8:33:04 PM UTC
                Message:

Task Name                                Start Time
End Time                                 Status
---------------------------------------- ----------------------------------
---------------------------------- ----------
Backup Validations                       November 18, 2016 8:33:04 PM UTC
November 18, 2016 8:33:13 PM UTC    Success
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
validate recovery window                November 18, 2016 8:33:13 PM UTC
November 18, 2016 8:33:17 PM UTC    Success
Db cross check                          November 18, 2016 8:33:17 PM UTC
November 18, 2016 8:33:23 PM UTC    Success
Database Backup                         November 18, 2016 8:33:23 PM UTC
November 18, 2016 8:34:22 PM UTC    Success
Backup metadata                         November 18, 2016 8:34:22 PM UTC
November 18, 2016 8:34:22 PM UTC    Success
```

dbcli list-jobs

Use the `dbcli list-jobs` command to display a list of jobs, including the job IDs, status, and the job

created date and time stamp.

Syntax

```
dbcli list-jobs [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

Example

The following command displays a list of jobs:

```
dbcli list-jobs
```

Output:

```
ID
Description
Created                             Status
----------------------------------------
--------------------------------------------------------------------------------
---------------------------------- ----------
0a362dac-0339-41b5-9c9c-4d229e363eaa     Database service creation with db
name: db11                               November 10, 2016 11:37:54 AM UTC
Success
9157cc78-b487-4ee9-9f46-0159f10236e4     Database service creation with db
name: jhfpdb                             November 17, 2016 7:19:59 PM UTC
Success
013c408d-37ca-4f58-a053-02d4efdc42d0     create backup
config:myBackupConfig                                        November 18,
2016 8:28:14 PM UTC    Success
921a54e3-c359-4aea-9efc-6ae7346cb0c2     update database
id:80ad855a-5145-4f8f-a08f-406c5e4684ff                    November 18, 2016
8:32:16 PM UTC    Success
74731897-fb6b-4379-9a37-246912025c17     Backup service creation with db
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
name: dbtst                                    November 18, 2016 8:33:04 PM
UTC    Success
40a227b1-8c47-46b9-a116-48cc1476fc12     Creating a report for database
80ad855a-5145-4f8f-a08f-406c5e4684ff         November 18, 2016 8:41:39 PM
UTC    Success
```

Latestpatch Command

dbcli describe-latestpatch

> **Note**
>
> - Your DB system might not include this newer command. If you have trouble running the command, use the `CLI Update` command to update the database CLI and then retry the command.
> - The `dbcli describe-latestpatch` command is not available on 2-node RAC DB systems. Patching 2-node systems from Object Storage is not supported.

Use the `dbcli describe-latestpatch` command show the latest patches applicable to the DB system and available in Oracle Cloud Infrastructure Object Storage.

This command requires a valid Object Storage credentials configuration. If the configuration is missing or invalid, the command fails with the error: `Failed to connect to the object store. Please provide valid details`.

For more information about updating the CLI, creating the credentials configuration, and applying patches, see *Update a DB System*.

Syntax

```
dbcli describe-latestpatch [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

Example

The following command displays patches available in the object store:

```
dbcli describe-latestpatch
```

Output:

```
componentType   availableVersion
--------------- --------------------
gi              12.1.0.2.161018
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
db              11.2.0.4.161018
db              12.1.0.2.161018
oak             12.1.2.10.0
```

Logcleanjob Commands

The following commands are available to manage log cleaning jobs:

- `dbcli create-logCleanJob`
- `dbcli describe-logCleanJob`
- `dbcli list-logCleanJobs`

dbcli create-logCleanJob

Use the `dbcli create-logCleanJob` command to create a log cleaning job.

Syntax

```
dbcli create-logCleanJob
  [-c {gi|database|dcs}]
  [-o <number>]
  [u {Day|Hour|Minute}]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-c	--components	(Optional) Components. Possible values are gi, database, and dcs. Separate multiple values by commas.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-o	--olderThan	(Optional) Quantity portion of time interval. Default: 30. Cleans logs older than the specified time interval (`-o` and `-u`).
-u	--unit	(Optional) Unit portion of time interval. Possible values: Day, Hour, or Minute. Default: Day. Cleans logs older than the specified time interval (`-o` and `-u`).

dbcli describe-logCleanJob

Use the `dbcli describe-logCleanJob` command to display the summary for a log cleaning job.

Syntax

```
dbcli describe-logCleanJob
  -i <job_id>
  [-h] [-j]
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--jobid	ID of log cleaning job for which to display the summary.
-j	--json	(Optional) Displays JSON output.

dbcli list-logCleanJobs

Use the `dbcli list-logCleanJobs` command to list log cleaning jobs.

Syntax

```
dbcli list-logCleanJobs [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

Netsecurity Commands

The following commands are available to manage network encryption on the DB system:

- `dbcli describe-netsecurity`
- `dbcli update-netsecurity`

dbcli describe-netsecurity

Use the `dbcli describe-netsecurity` command to display the current network encryption setting for a database home.

Syntax

```
dbcli describe-netsecurity -H <db_home_id> [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-H	--dbHomeId	Defines the database home ID. Use the `dbcli list-dbhomes` command to get the dbhomeid.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Example

The following command displays the encryption setting for specified database home:

```
dbcli describe-netsecurity -H 16c96a9c-f579-4a4c-a645-8d4d22d6889d
```

Output:

```
NetSecurity Rules
-----------------------------------------------------------------
         DatabaseHomeID: 16c96a9c-f579-4a4c-a645-8d4d22d6889d

                   Role: Server
   EncryptionAlgorithms: AES256 AES192 AES128
    IntegrityAlgorithms: SHA1
         ConnectionType: Required

                   Role: Client
   EncryptionAlgorithms: AES256 AES192 AES128
    IntegrityAlgorithms: SHA1
         ConnectionType: Required
```

dbcli update-netsecurity

Use the `dbcli update-netsecurity` command to update the Oracle Net security configuration on the DB system.

Syntax

```
dbcli update-netsecurity
  {-c|-s}
  -t {REJECTED|ACCEPTED|REQUESTED|REQUIRED}
  -H <db_home_id>
  -e {AES256|AES192|AES128}
  -i {SHA1|SHA512|SHA384|SHA256}
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-c	--client	Indicates that the specified data encryption or data integrity configuration is for the client. (`--client` and `--server` are mutually exclusive.)
-e	--encryptionAlgorithms	Defines the algorithm to be used for encryption. Specify either AES256, AES192, or AES128.
-H	--dbHomeId	Defines the database home ID. Use the `dbcli list-dbhomes` command to get the dbHomeId.
-h	--help	(Optional) Displays help for using the command.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameter	Full Name	Description
-i	--integrityAlgorithms	Defines the algorithm to be used for integrity. Specify either SHA1, SHA512, SHA384, or SHA256. For Oracle Database 11g, the only accepted value is SHA1.
-j	--json	(Optional) Displays JSON output.
-s	--server	Indicates that the specified data encryption or data integrity configuration is for the server. (`--client` and `--server` are mutually exclusive.)
-t	--connectionType	Specifies how Oracle Net Services data encryption or data integrity is negotiated with clients. The following values are listed in the order of increasing security: REJECTED - Do not enable data encryption or data integrity, even if required by the client. ACCEPTED - Enable data encryption or data integrity if required or requested by the client. REQUESTED - Enable data encryption or data integrity if the client permits it. REQUIRED - Enable data encryption or data integrity or preclude the connection.

Example

The following command updates the connection type to ACCEPTED:

```
dbcli update-netsecurity -H a2ffbb07-c9c0-4467-a458-bce4d3b76cd5 -t ACCEPTED
```

Objectstoreswift Commands

You can back up a database to an existing bucket in the Oracle Cloud Infrastructure Object Storage service by using the `dbcli create-backup` command, but first you'll need to:

1. Create an object store on the DB system, which contains the endpoint and credentials to access Object Storage, by using the `dbcli create-objectstoreswift` command.
2. Create a backup configuration that refers to the object store ID and the bucket name by using the `dbcli create-backupconfig` command.
3. Associate the backup configuration with the database by using the `dbcli update-database` command.

The following commands are available to manage object stores.

- `dbcli create-objectstoreswift`
- `dbcli describe-objectstoreswift`
- `dbcli list-objectstoreswifts`

dbcli create-objectstoreswift

Use the `dbcli create-objectstoreswift` command to create an object store.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Syntax

```
dbcli create-objectstoreswift
  -n <object_store_name>
  -t <object_storage_namespace>
  -u <user_name>
  -e https://swiftobjectstorage.<region_name>.oraclecloud.com/v1
  -p <password>
  [-h] [-j]
```

where <*object_storage_namespace*> is your tenancy's Object Storage namespace.

Parameters

Parameter	Full Name	Description
-e	--endpointurl	The following endpoint URL. `https://swiftobjectstorage.<region_name>.oraclecloud.com/v1`
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-n	--name	The name for the object store to be created.
-p	--swiftpassword	The auth token that you generated by using the Console or IAM API. For information about generating an auth token for use with Swift, see Managing User Credentials. This is **not** the password for the Oracle Cloud Infrastructure user. Specify `-p` (with no password) to be prompted. Specify `-hp "<password>"` in quotes to provide the password (auth token) in the command.
-t	--tenantname	The Object Storage namespace of your tenancy.

-u	--username	The user name for the Oracle Cloud Infrastructure user account, for example: `-u djones@example.com`
		This is the user name you use to sign in to the Console.
		The user name must have tenancy-level access to the Object Storage. An easy way to do this is to add the user name to the Administrators group. However, that allows access to *all* of the cloud services. Instead, an administrator can create a policy that allows tenancy-level access to just Object Storage. The following is an example of such a policy.
		`Allow group DBAdmins to manage buckets in tenancy`
		`Allow group DBAdmins to manage objects in tenancy`
		For more information about adding a user to a group, see Managing Groups. For more information about policies, see Getting Started with Policies.

Example

The following command creates an object store and prompts for the Swift password:

```
dbcli create-objectstoreswift -n r2swift -t MyObjectStorageNamespace -u
djones@example.com -e https://
swiftobjectstorage.<region_name>.oraclecloud.com/v1 -p
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Output:

```
Password for Swift:
{
  "jobId" : "c565bb71-f67b-4fab-9d6f-a34eae36feb7",
  "status" : "Created",
  "message" : "Create object store swift",
  "reports" : [ ],
  "createTimestamp" : "January 19, 2017 11:11:33 AM PST",
  "resourceList" : [ {
    "resourceId" : "8a0fe039-f5d4-426a-8707-256c612b3a30",
    "resourceType" : "ObjectStoreSwift",
    "jobId" : "c565bb71-f67b-4fab-9d6f-a34eae36feb7",
    "updatedTime" : "January 19, 2017 11:11:33 AM PST"
  } ],
  "description" : "create object store:biyanr2swift",
  "updatedTime" : "January 19, 2017 11:11:33 AM PST"
}
```

dbcli describe-objectstoreswift

Use the `dbcli describe-objectstoreswift` command to display details about an object store.

Syntax

```
dbcli describe-objectstoreswift
  -i <object_store_swift_id>
  -in <object_store_swift_name>
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--objectstoreswiftid	The object store ID. Use the `dbcli list-objectstoreswifts` command to get the ID.
-in	--objectstoreswiftName	The object store name. Use the `dbcli list-objectstoreswifts` command to get the name.
-j	--json	(Optional) Displays JSON output.

Example

The following command displays details about an object store:

```
dbcli describe-objectstoreswift -i 910e9e2d-25b4-49b4-b88e-ff0332f7df87
```

Output:

```
Object Store details
----------------------------------------------------------------
                     ID: 910e9e2d-25b4-49b4-b88e-ff0332f7df87
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
                  Name: objstrswift15
              UserName: djones@example.com
            TenantName: CompanyABC
          endpoint URL: https://
swiftobjectstorage.<region_name>.oraclecloud.com/v1
           CreatedTime: November 16, 2016 11:25:34 PM UTC
           UpdatedTime: November 16, 2016 11:25:34 PM UTC
```

dbcli list-objectstoreswifts

Use the `dbcli list-objectstoreswifts` command to list the object stores on a DB system.

Syntax

```
dbcli list-objectstoreswifts [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

Example

The following command lists the object stores on the DB system:

```
dbcli list-objectstoreswifts
```

Output:

```
ID                                       Name
UserName             TenantName
Url                                                          createTime
---------------------------------------- --------------------
-------------------  -------------- ------
--------------------------------------------------
----------------------------------
2915bc6a-6866-436a-a38c-32302c7c4d8b     swiftobjstr1
djones@example.com    LargeComputers https://
swiftobjectstorage.<region_name>.oraclecloud.com/v1 November 10, 2016 8:42:18
PM UTC
910e9e2d-25b4-49b4-b88e-ff0332f7df87     objstrswift15
djones@example.com    LargeComputers https://
swiftobjectstorage.<region_name>.oraclecloud.com/v1 November 16, 2016
11:25:34 PM UTC
```

Pendingjob Command

dbcli list-pendingjobs

Use the `dbcli list-pendingjobs` command to display a list of pending jobs.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Syntax

```
dbcli list-pendingjobs [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

Rmanbackupreport Commands

The following commands are available to manage RMAN backup reports:

- `dbcli create-rmanbackupreport`
- `dbcli delete-rmanbackupreport`
- `dbcli describe-rmanbackupreport`
- `dbcli list-rmanbackupreports`

dbcli create-rmanbackupreport

Use the `dbcli create-rmanbackupreport` command to create an RMAN backup report.

Syntax

```
dbcli create-rmanbackupreport
  -w {summary|detailed}
  -rn <name>
  [-i <db_id>]
  [-in <db_name>]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--dbid	(Optional) Database resource ID.
-in	--dbname	(Optional) Database resource name.
-j	--json	(Optional) Displays JSON output.
-rn	--rptname	RMAN backup report name. Maximum number of characters: 30. Wrap name in single quotes when special characters are used.
-w	--reporttype	RMAN backup report type. Possible values: summary or detailed.

dbcli delete-rmanbackupreport

Use the `dbcli delete-rmanbackupreport` command to delete an RMAN backup report.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Syntax

```
dbcli delete-rmanbackupreport
  [-d <db_id>]
  [-dn <db_name>]
  [-n <number>]
  [-i <rpt_id>]
  [-in <rpt_name>]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-d	--dbid	(Optional) Database resource ID.
-dn	--dbname	(Optional) Database resource name.
-h	--help	(Optional) Displays help for using the command.
-i	--reportid	(Optional) RMAN backup report ID
-in	--rptname	(Optional) RMAN backup report name
-j	--json	(Optional) Displays JSON output.
-n	--numofday	(Optional) Number of days since created (provided with Database ID/ Database Name)

dbcli describe-rmanbackupreport

Use the `dbcli describe-rmanbackupreport` command to

Syntax

```
dbcli describe-rmanbackupreport
  [-i <rpt_id>]
  [-in <rpt_name>]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--id	(Optional) RMAN backup report ID
-in	--name	(Optional) RMAN backup report name
-j	--json	(Optional) Displays JSON output.

dbcli list-rmanbackupreports

Use the `dbcli list-rmanbackupreports` command to

Syntax

```
dbcli list-rmanbackupreports
  [-i <db_id>]
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
[-in <db_name>]
[-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--dbid	(Optional) Database resource ID.
-in	--dbName	(Optional) Database resource name.
-j	--json	(Optional) Displays JSON output.

Schedule Commands

The following commands are available to manage schedules:

- `dbcli describe-schedule`
- `dbcli list-schedules`
- `dbcli update-schedule`

dbcli describe-schedule

Use the `dbcli describe-schedule` command to describe a schedule.

Syntax

```
dbcli describe-schedule -i <id> [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--scheduleid	Schedule ID.
-j	--json	(Optional) Displays JSON output.

dbcli list-schedules

Use the `dbcli list-schedules` command to list schedules.

Syntax

```
dbcli list-schedules [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

Copyright © 2022, 2026, Oracle and/or its affiliates.

dbcli update-schedule

Use the `dbcli update-schedule` command to update a schedule.

Syntax

```
dbcli update-schedule
  -i <id>
  [-x <expression>]
  [-t <description>]
  [-d]
  [-e]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-d	--disable	(Optional) Disables the schedule.
-e	--enable	(Optional) Enables the schedule.
-h	--help	(Optional) Displays help for using the command.
-i	--scheduleid	Schedule ID.
-j	--json	(Optional) Displays JSON output.
-t	--description	(Optional) Description
-x	--cronExpression	(Optional) Cron expression. Use cronmaker.com to generate a valid cron expression.

Scheduledexecution Command

dbcli list-scheduledExecutions

Use the `dbcli list-scheduledExecutions` command to list scheduled executions.

Syntax

```
dbcli list-scheduledExecutions
  [-e <execution_id>]
  [-i <schedule_id>]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-e	--executionid	(Optional) Execution ID.
-h	--help	(Optional) Displays help for using the command.
-i	--scheduleid	(Optional) Schedule ID.
-j	--json	(Optional) Displays JSON output.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Server Command

dbcli update-server

> **Note**
>
> Your DB system might not include this newer command. If you have trouble running the command, use the `CLI Update` command to update the database CLI and then retry the command.

Use the `dbcli update-server` command to apply patches to the server components in the DB system. For more information about applying patches, see *Update a DB System*.

Syntax

```
dbcli update-server
  [-n <number>]
  [--local]
  [--precheck]
  [-v]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-l	--local	(Optional) Performs the operation on the local node of a multi-node high availability (HA) system. This parameter is not needed to perform the operation on a single-node system.
-n	--node	(Optional) Node number to be updated. Use the `dbcli list-nodes` command to get the node number.
-p	--precheck	(Optional) Runs precheck operations to check prerequisites.
-v	--version	(Optional) Version to be updated.

Examples

The following commands update the server and show the output from the update job:

```
dbcli update-server
```

Output:

```
{
  "jobId" : "9a02d111-e902-4e94-bc6b-9b820ddf6ed8",
  "status" : "Created",
  "reports" : [ ],
  "createTimestamp" : "January 19, 2017 09:37:11 AM PST",
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
  "resourceList" : [ ],
  "description" : "Server Patching",
  "updatedTime" : "January 19, 2017 09:37:11 AM PST"
}
```

```
dbcli describe-job -i 9a02d111-e902-4e94-bc6b-9b820ddf6ed8
```

Output:

```
Job details
----------------------------------------------------------------
                     ID:  9a02d111-e902-4e94-bc6b-9b820ddf6ed8
            Description:  Server Patching
                 Status:  Running
                Created:  January 19, 2017 9:37:11 AM PST
                Message:

Task Name                                Start Time
End Time                              Status
---------------------------------------- --------------------------------
---------------------------------- ----------
Create Patching Repository Directories   January 19, 2017 9:37:11 AM PST
January 19, 2017 9:37:11 AM PST     Success
Download latest patch metadata           January 19, 2017 9:37:11 AM PST
January 19, 2017 9:37:11 AM PST     Success
Update System version                    January 19, 2017 9:37:11 AM PST
January 19, 2017 9:37:11 AM PST     Success
Update Patching Repository               January 19, 2017 9:37:11 AM PST
January 19, 2017 9:38:35 AM PST     Success
oda-hw-mgmt upgrade                      January 19, 2017 9:38:35 AM PST
January 19, 2017 9:38:58 AM PST     Success
Opatch updation                          January 19, 2017 9:38:58 AM PST
January 19, 2017 9:38:58 AM PST     Success
Patch conflict check                     January 19, 2017 9:38:58 AM PST
January 19, 2017 9:42:06 AM PST     Success
apply clusterware patch                  January 19, 2017 9:42:06 AM PST
January 19, 2017 10:02:32 AM PST    Success
Updating GiHome version                  January 19, 2017 10:02:32 AM PST
January 19, 2017 10:02:38 AM PST    Success
```

The following command updates node 0 of the server only, with precheck:

```
dbcli update-server -n 0 -p
```

Output:

```
{
  "jobId" : "3e2a1e3c-83d3-4101-86b8-4d525f3f8c18",
  "status" : "Created",
  "message" : null,
  "reports" : [ ],
  "createTimestamp" : "April 26, 2019 06:07:27 AM UTC",
  "resourceList" : [ ],
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
  "description" : "Server Patching Prechecks",
  "updatedTime" : "April 26, 2019 06:07:27 AM UTC"
}
```

System Command

dbcli describe-system

Use the `dbcli describe-system` command to display details about the system. On a 2-node RAC DB system, the command provides information about the local node.

Syntax

```
dbcli describe-system [-b] [-d] [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-b	--bom	(Optional) Displays BOM information.
-d	--details	(Optional) Displays additional information about the DB system, including dcs CLI and agent version information.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

TDE Commands

The following commands are available to manage TDE-related items (backup reports, keys, and wallets):

- `dbcli list-tdebackupreports`
- `dbcli update-tdekey`
- `dbcli recover-tdewallet`

dbcli list-tdebackupreports

Use the `dbcli list-tdebackupreports` command to list backup reports for TDE wallets.

Syntax

```
dbcli  list-tdebackupreports
  [-i <db_id>]
  [-in <db_name>]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameter	Full Name	Description
-i	--dbResid	(Optional) Displays the TDE Wallet backup reports for the specified database resource ID. Use the `dbcli list-databases` command to get the database resource ID.
-in	--dbResname	(Optional) Displays the TDE Wallet backup reports for the specified database resource name. Use the `dbcli list-databases` command to get the database resource name.
-j	--json	(Optional) Displays JSON output.

Example

The following command lists the backup reports for TDE wallets:

```
dbcli list-tdebackupreports
```

Output:

```
DbResID                                  OraDbId            BackupLocation
---------------------------------------- -------------------
----------------------------------------
538ca5b1-654d-4418-8ce1-f49b6c987a60    1257156075         https://
swiftobjectstorage.us-phoenix-1.oraclecloud.com/v1/dbaasimage/backuptest/
host724007/tdewallet/Testdb5/1257156075/2017-08-17/
TDEWALLET_BMC60_2017-08-17_10-58-17.0990.tar.gz
538ca5b1-9fb2-4245-b157-6e25d7c988c5     704287483         https://
swiftobjectstorage.us-phoenix-1.oraclecloud.com/v1/dbaasimage/backuptest/
host724007/tdewallet/Testdb1/704287483/2017-08-17/
TDEWALLET_AUTO_2017-08-17_11-03-25.0953.tar.gz
538ca5b1-9fb2-4245-b157-6e25d7c988c5     704287483         https://
swiftobjectstorage.us-phoenix-1.oraclecloud.com/v1/dbaasimage/backuptest/
host724007/tdewallet/Testdb1/704287483/2017-08-17/
TDEWALLET_BMC62_2017-08-17_11-04-41.0264.tar.gz
19714ffa-de1b-4433-9188-c0592887e609     1157116855         https://
swiftobjectstorage.us-phoenix-1.oraclecloud.com/v1/dbaasimage/backuptest/
host724007/tdewallet/Testdb7/1157116855/2017-08-17/
TDEWALLET_AUTO_2017-08-17_11-57-47.0605.tar.gz
```

dbcli update-tdekey

Use the `dbcli update-tdekey` command to update the TDE encryption key inside the TDE wallet. You can update the encryption key for Pluggable Databases (if -pdbNames are specified), and/or the Container Database (if -rootDatabase is specified).

Syntax

```
dbcli update-tdekey
  -i <db_id>
  -p [-all]
  -n <pdbname1,pdbname2>
  [-r|-no-r]
  -t <tag_name>
  [-h] [-j]
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Parameters

Parameter	Full Name	Description
-all	--allPdbNames	(Optional) Flag to rotate (update) all PDB names. To update all instead of specified PDB names, use this parameter instead of `-n`. Default: false.
-i	--databaseId	Defines the database ID for which to update the key.
-p	--password	Defines the TDE Admin wallet password. Specify `-p` with no password. You will be prompted for the password. If you must provide the password in the command, for example in a script, use `-hp <password>` instead of -p.
-n	--pdbNames	Defines the PDB names to be rotated (updated).
-r -no-r	--rootDatabase --no-rootDatabase	Indicates whether to rotate the key for the root database if it is a container database.
-t	-tagName	Defines the TagName used to backup the wallet. The default is OdaRotateKey.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

Example

The following command updates the key for pdb1 and pdb2 only:

```
dbcli update-tdekey -dbid ee3eaab6-a45b-4e61-a218-c4ba665503d9 -p -n pdb1,pdb2
```

Output:

```
TDE Admin wallet password:
{
  "jobId" : "08e5edb1-42e1-4d16-a47f-783c0afa4778",
  "status" : "Created",
  "message" : null,
  "reports" : [ ],
  "createTimestamp" : 1467876407035,
  "description" : "TDE update",
  "updatedTime" : 1467876407035
}
```

The following command updates pdb1, pdb2, and the container database:

```
dbcli update-tdekey -dbid ee3eaab6-a45b-4e61-a218-c4ba665503d9 -p -n
pdb1,pdb2 -r
```

Output:

```
TDE Admin wallet password:
{
  "jobId" : "c72385f0-cd81-42df-a8e8-3a1e7cab1278",
  "status" : "Created",
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
  "message" : null,
  "reports" : [ ],
  "createTimestamp" : 1467876433783,
  "description" : "TDE update",
  "updatedTime" : 1467876433783
}
```

dbcli recover-tdewallet

Use the `dbcli recover-tdewallet` command to recover a TDE wallet.

Syntax

```
dbcli recover-tdewallet
  -in <db_name>
  -tp <password>
  [-l <location>]
  [-h] [-j]
```

Parameters

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-in	--dbName	Database name.
-j	--json	(Optional) Displays JSON output.
-l	--tdeWalletBackuplocation	(Optional) TDE wallet backup location. TDE wallet should b ebacked up in tar.gz format.
-tp	--tdeWalletPassword	Defines the TDE Admin wallet password.

Compartment Quotas

Use Compartment Quotas in Oracle Base Database service to control resource consumption within compartments.

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 OCI Console.

General Information

Using compartment quotas, you can restrict, on a compartment by compartment basis, the availability of the following Base Database resources in your tenancy:

- Available data storage
- Number of available OCPUs

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> Quotas are enforced only for provisioning databases and manually scaling their allocated CPUs 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 `vm-standard-a1-ocpu-count` quota to 40 OCPUs in the compartment `MyCompartment`:

```
set database quota vm-standard-a1-ocpu-count to 40 in compartment
MyCompartment
```

unset

Resets quotas back to the default service limits.

For example, the following statement resets the quota for `vm-standard-a1-ocpu-count` in the entire tenancy:

```
unset database quota vm-standard-a1-ocpu-count in tenancy
```

zero

Removes access to a cloud resource for a compartment.

For example, the following statement disallows use of outbound email or notifications in the compartment `MyCompartment`:

```
Zero email-delivery quotas in compartment MyCompartment
Zero notifications quotas in compartment MyCompartment
```

For more information about quotas, see Overview of Compartment Quotas.

Available Quotas

The following table lists the quotas for Base Database service.

Family name: `database`

Table 7-1 Available Quotas

Name	Scope	Description
`vm-block-storage-gb`	Availability domain	Total size (in GB) of block storage attachments across all DB systems.
`vm-standard1-ocpu-count`	Availability domain	Number of VM.Standard1.x OCPUs for the Intel X7 processor (fixed shape) based DB systems.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 7-1 (Cont.) Available Quotas

Name	Scope	Description
`vm-standard2-ocpu-count`	Availability domain	Number of VM.Standard2.x OCPUs for the Intel X7 processor (fixed shape) based DB systems.
`vm-standard3-ocpu-count`	Availability domain	Number of VM.Standard3.Flex OCPUs for the Intel X9 processor (flexible shape) based DB systems.
`vm-standard-a1-ocpu-count`	Availability domain	Number of VM.Standard.A1.Flex OCPUs for the Arm-based Ampere A1 processor (flexible shape) based DB systems.
`vm-standard-e4-ocpu-count`	Availability domain	Number of VM.Standard.E4.flex OCPUs for the AMD processor (flexible shape) based DB systems.
`vm-standard-x86-ecpu-count`	Availability domain	Number of ECPUs for VM.Standard.x86 based DB systems.
`vm-basedb-x86-ecpu-count`	Availability domain	Number of ECPUs for VM.BaseDB.x86 based DB systems.

Manage Time Zone

The **Time Zone** field in the Console and in the API allows you to launch DB System resources with a time zone other than UTC (the default).

The time zone that you specify when you create the database system applies to the host and to the Oracle Grid Infrastructure (if the system has Grid Infrastructure), and controls the time zone of the database log files. The time zone of the database itself is not affected. However, the database's time zone affects only the `timestamp` datatype. You can change the database time zone manually, but Oracle recommends that you keep it as UTC (the default) to avoid data conversion and improve performance when data is transferred among databases. This configuration is especially important for distributed databases, replication, and export and import operations.

Although UTC is the recommended time zone to use, having a common time zone for your database clients and application hosts can simplify management and troubleshooting for the database administrator.

> **Note**
>
> Time zones are largely used for display purposes or to handle user input. Changing time zone does not change the time on the system clock.

Time Zone Options

Whether you use the Console or the API, the time zone options you can select from are represented in the named region format, for example, *America/Los_Angeles*. The Console allows you to select UTC, the time zone detected in your browser (if your browser supports time zone detection), or an alternate time zone.

To specify an alternate time zone (the **Select another time zone** option), you first select a value in the **Region or country** field to narrow the list of time zones to select from in the **Time zone** field. In the *America/Los_Angeles* example, *America* is the time region and *Los_Angeles* is the time zone. The options you see in these two fields roughly correlate with the time zones

Copyright © 2022, 2026, Oracle and/or its affiliates.

supported in both the `Java.util.TimeZone` class and on the Linux operating system. If you do not see the time zone you are looking for, try selecting "Miscellaneous" in the **Region and country** field.

> **Tip**
>
> If you are using the API and would like to see a list of supported time zones, you can examine the time zone options in the Console. These options appear on the **Create DB System** page when you show advanced options after you select a DB system shape.

View the Current Time Zone

Perform the following steps to view the time zone information.

1. Log on to the host as the grid user.

   Example:

   ```
   [opc@rc ~]$ sudo su - grid
   ```

2. Use the following command to view the time zone.

   ```
   timedatectl
   ```

   Example:

   ```
   [grid@rc ~]$ timedatectl
           Local time: Thu 2023-03-16 08:00:25 UTC
       Universal time: Thu 2023-03-16 08:00:25 UTC
             RTC time: Thu 2023-03-16 08:00:21
            Time zone: UTC (UTC, +0000)
          NTP enabled: yes
     NTP synchronized: no
      RTC in local TZ: no
           DST active: n/a
   ```

   Notice that the system is currently in the UTC time zone.

3. Log on to the database.

   Example:

   ```
   [grid@rc ~]$ sqlplus / as sysdba

   SQL*Plus: Release 21.0.0.0.0 - Production on Thu Mar 16 08:00:33 2023
   Version 21.9.0.0.0

   Copyright (c) 1982, 2022, Oracle.  All rights reserved.

   Connected to:
   Oracle Database 21c Enterprise Edition Release 21.0.0.0.0 -Production
   Version 21.9.0.0.0
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

4. You can view the current database time using the following command.

   ```
   sysdate
   ```

   Example:

   ```
   SQL> select to_char(sysdate,'DD-MON-YYYY HH24:MI:SS') from dual;

   TO_CHAR(SYSDATE,'DD-
   --------------------
   16-MAR-2023 08:00:40
   ```

5. Exit.

Change the Time Zone of the DB System

Perform the following steps to set the time zone of the DB System.

1. Log on to the host as the root user.

   Example:

   ```
   [opc@rc ~]$ sudo su -
   ```

2. Identify the correct time zone (TZ) value to be set using the following command. For example, if you want to identify the correct TZ value for the US eastern time zone (EST/EDT), you can use the `tzselect` command and choose Americas/United States/Eastern.

   ```
   tzselect
   ```

   Example:

   ```
   [root@rc ~]# tzselect
   Please identify a location so that time zone rules can be set correctly.
   Please select a continent or ocean.
    1) Africa
    2) Americas
    3) Antarctica
    4) Arctic Ocean
    5) Asia
    6) Atlantic Ocean
    7) Australia
    8) Europe
    9) Indian Ocean
   10) Pacific Ocean
   11) none - I want to specify the time zone using the Posix TZ format.
   #? 2
   Please select a country.
    1) Anguilla              19) Dominican Republic    37) Peru
    2) Antigua & Barbuda     20) Ecuador               38) Puerto Rico
    3) Argentina             21) El Salvador           39) St Barthelemy
    4) Aruba                 22) French Guiana         40) St Kitts & Nevis
    5) Bahamas               23) Greenland             41) St Lucia
    6) Barbados              24) Grenada               42) St Maarten (Dutch)
    7) Belize                25) Guadeloupe            43) St Martin (French)
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
 8) Bolivia                26) Guatemala             44) St Pierre &
Miquelon
 9) Brazil                 27) Guyana                45) St Vincent
10) Canada                 28) Haiti                 46) Suriname
11) Caribbean NL           29) Honduras              47) Trinidad & Tobago
12) Cayman Islands         30) Jamaica               48) Turks & Caicos Is
13) Chile                  31) Martinique            49) United States
14) Colombia               32) Mexico                50) Uruguay
15) Costa Rica             33) Montserrat            51) Venezuela
16) Cuba                   34) Nicaragua             52) Virgin Islands
(UK)
17) Curaçao                35) Panama                53) Virgin Islands
(US)
18) Dominica               36) Paraguay
#? 49
Please select one of the following time zone regions.
 1) Eastern (most areas)              16) Central - ND (Morton rural)
 2) Eastern - MI (most areas)         17) Central - ND (Mercer)
 3) Eastern - KY (Louisville area)    18) Mountain (most areas)
 4) Eastern - KY (Wayne)              19) Mountain - ID (south); OR (east)
 5) Eastern - IN (most areas)         20) MST - Arizona (except Navajo)
 6) Eastern - IN (Da, Du, K, Mn)      21) Pacific
 7) Eastern - IN (Pulaski)            22) Alaska (most areas)
 8) Eastern - IN (Crawford)           23) Alaska - Juneau area
 9) Eastern - IN (Pike)               24) Alaska - Sitka area
10) Eastern - IN (Switzerland)        25) Alaska - Annette Island
11) Central (most areas)              26) Alaska - Yakutat
12) Central - IN (Perry)              27) Alaska (west)
13) Central - IN (Starke)             28) Aleutian Islands
14) Central - MI (Wisconsin border)   29) Hawaii
15) Central - ND (Oliver)
#? 1
The following information has been given:
        United States
        Eastern (most areas)
Therefore TZ='America/New_York' will be used.
    Local time is now:    Thu Mar 16 04:01:39 EDT 2023.
Universal Time is now:    Thu Mar 16 08:01:39 UTC 2023.
Is the above information OK?
1) Yes
2) No
#? 1
You can make this change permanent for yourself by appending the line
        TZ='America/New_York'; export TZ
to the file '.profile' in your home directory; then log out and log in
again.
Here is that TZ value again, this time on standard output so that you
can use the /bin/tzselect command in shell scripts:
        America/New_York
[root@rc ~]#
```

From the `tzselect` command, we have identified the TZ value as `America/New_York` for the EDT time zone.

Copyright © 2022, 2026, Oracle and/or its affiliates.

3. Update the following command to set the new time zone.

   ```
   timedatectl set-timezone <new_time_zone>
   ```

   Example:

   ```
   [root@rc ~]# timedatectl set-timezone America/New_York
   ```

4. Verify if the time zone has been updated using the following command.

   ```
   timedatectl
   ```

   Example:

   ```
   [root@rc ~]# timedatectl
            Local time: Thu 2023-03-16 04:02:15 EDT
        Universal time: Thu 2023-03-16 08:02:15 UTC
              RTC time: Thu 2023-03-16 08:02:10
             Time zone: America/New_York (EDT, -0400)
           NTP enabled: yes
      NTP synchronized: no
       RTC in local TZ: no
            DST active: yes
       Last DST change: DST began at
                        Sun 2023-03-12 01:59:59 EST
                        Sun 2023-03-12 03:00:00 EDT
       Next DST change: DST ends (the clock jumps one hour backwards) at
                        Sun 2023-11-05 01:59:59 EDT
                        Sun 2023-11-05 01:00:00 EST
   ```

   Notice that the system is now in the America/New_York (EDT) time zone.

5. Exit.

Change the Time Zone of the Host on DB Systems that Use Grid Infrastructure

The time zone of the Oracle Grid Infrastructure determines the time zone of the database log files.

Perform the following steps to update the time zone information of the host of the DB systems that use grid infrastructure for storage management.

> **Note**
>
> - Steps 1 through 10 should be performed on all the nodes in a RAC cluster.
> - Steps 11 through 15 can be performed on any node in a RAC cluster.
> - This procedure does not apply to fast provisioned DB systems, which use Logical Volume Manager instead of Grid Infrastructure for storage management.

Copyright © 2022, 2026, Oracle and/or its affiliates.

1. Log on to the host as the `grid` user.

   Example:

   ```
   [opc@rc ~]$ sudo su - grid
   ```

2. Identify the correct time zone (TZ) value to be set using the following command. For example, if you want to identify the correct TZ value for the US Pacific time zone (PST/PDT), you can use the `tzselect` command and choose Americas/United States/Pacific.

   ```
   tzselect
   ```

   Example:

   ```
   [grid@rc ~]$ tzselect
   Please identify a location so that time zone rules can be set correctly.
   Please select a continent or ocean.
    1) Africa
    2) Americas
    3) Antarctica
    4) Arctic Ocean
    5) Asia
    6) Atlantic Ocean
    7) Australia
    8) Europe
    9) Indian Ocean
   10) Pacific Ocean
   11) none - I want to specify the time zone using the Posix TZ format.
   #? 2
   Please select a country.
    1) Anguilla              19) Dominican Republic    37) Peru
    2) Antigua & Barbuda     20) Ecuador               38) Puerto Rico
    3) Argentina             21) El Salvador           39) St Barthelemy
    4) Aruba                 22) French Guiana         40) St Kitts & Nevis
    5) Bahamas               23) Greenland             41) St Lucia
    6) Barbados              24) Grenada               42) St Maarten (Dutch)
    7) Belize                25) Guadeloupe            43) St Martin (French)
    8) Bolivia               26) Guatemala             44) St Pierre &
   Miquelon
    9) Brazil                27) Guyana                45) St Vincent
   10) Canada                28) Haiti                 46) Suriname
   11) Caribbean NL          29) Honduras              47) Trinidad & Tobago
   12) Cayman Islands        30) Jamaica               48) Turks & Caicos Is
   13) Chile                 31) Martinique            49) United States
   14) Colombia              32) Mexico                50) Uruguay
   15) Costa Rica            33) Montserrat            51) Venezuela
   16) Cuba                  34) Nicaragua             52) Virgin Islands
   (UK)
   17) Curaçao               35) Panama                53) Virgin Islands
   (US)
   18) Dominica              36) Paraguay
   #? 49
   Please select one of the following time zone regions.
    1) Eastern (most areas)             16) Central - ND (Morton rural)
    2) Eastern - MI (most areas)        17) Central - ND (Mercer)
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
 3) Eastern - KY (Louisville area)    18) Mountain (most areas)
 4) Eastern - KY (Wayne)              19) Mountain - ID (south); OR (east)
 5) Eastern - IN (most areas)         20) MST - Arizona (except Navajo)
 6) Eastern - IN (Da, Du, K, Mn)      21) Pacific
 7) Eastern - IN (Pulaski)            22) Alaska (most areas)
 8) Eastern - IN (Crawford)           23) Alaska - Juneau area
 9) Eastern - IN (Pike)               24) Alaska - Sitka area
10) Eastern - IN (Switzerland)        25) Alaska - Annette Island
11) Central (most areas)              26) Alaska - Yakutat
12) Central - IN (Perry)              27) Alaska (west)
13) Central - IN (Starke)             28) Aleutian Islands
14) Central - MI (Wisconsin border)   29) Hawaii
15) Central - ND (Oliver)
#? 21
The following information has been given:
        United States
        Pacific
Therefore TZ='America/Los_Angeles' will be used.
    Local time is now:    Thu Mar 16 01:08:57 PDT 2023.
Universal Time is now:    Thu Mar 16 08:08:57 UTC 2023.
Is the above information OK?
1) Yes
2) No
#? 1
You can make this change permanent for yourself by appending the line
        TZ='America/Los_Angeles'; export TZ
to the file '.profile' in your home directory; then log out and log in
again.
Here is that TZ value again, this time on standard output so that you
can use the /bin/tzselect command in shell scripts:
        America/Los_Angeles
[grid@rc ~]$
```

From the `tzselect` command, we have identified the TZ value as `America/Los_Angeles` for the PDT time zone.

3. Run the following command to get the host name.

```
hostname
```

Example:

```
[grid@rc ~]$ hostname
rc
```

4. Run the following command to get the database name.

```
srvctl config database -v
```

Example:

```
[grid@rc ~]$ srvctl config database -v
SMDB0316_iad1cx    /u01/app/oracle/product/21.0.0.0/dbhome_1    21.0.0.0.0
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

5. Log out as a grid user and sign in as a root user.

   Example:

   ```
   [grid@rc ~]$ exit
   logout
   [opc@rc ~]$ sudo su -
   ```

6. Navigate to the following directory.

   ```
   cd $GRID_HOME/crs/install
   ```

   Example:

   ```
   [root@rc ~]# cd /u01/app/21.0.0.0/grid/crs/install/
   ```

7. Edit the TZ content of the following file. Save and close the file after editing it.

   ```
   vim s_crsconfig_<hostname>_env.txt
   ```

   Example:

   ```
   [root@rc install]# vim s_crsconfig_rc_env.txt
   ```

8. Verify if the TZ entry in the file is updated.

   ```
   cat s_crsconfig_<hostname>_env.txt
   ```

   Example:

   ```
   [root@rc install]# cat s_crsconfig_rc_env.txt
   ########################################################################
   #This file can be used to set values for the NLS_LANG and TZ environment
   #variables and to set resource limits for Oracle Clusterware and
   #Database processes.
   #1. The NLS_LANG environment variable determines the language and
   #   characterset used for messages. For example, a new value can be
   #   configured by setting NLS_LANG=JAPANESE_JAPAN.UTF8
   #2. The Time zone setting can be changed by setting the TZ entry to
   #   the appropriate time zone name. For example, TZ=America/New_York
   #3. Resource limits for stack size, open files and number of processes
   #   can be specified by modifying the appropriate entries.
   #
   #Do not modify this file except as documented above or under the
   #direction of Oracle Support Services.
   ########################################################################
   TZ=America/Los_Angeles
   NLS_LANG=AMERICAN_AMERICA.AL32UTF8
   CRS_LIMIT_STACK=2048
   CRS_LIMIT_OPENFILE=65536
   CRS_LIMIT_NPROC=65536
   TNS_ADMIN=
   [root@rc install]#
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

9. Set the new time zone using the following command.

```
timedatectl set-timezone <new_time_zone>
```

Example:

```
[root@rc install]# timedatectl set-timezone America/Los_Angeles
```

> **Note**
>
> - Steps 1 through 9 should be performed on all the nodes in a RAC cluster.
> - Steps 10 through 15 can be performed on any node in a RAC cluster.

10. Navigate to the following directory.

```
cd $GRID_HOME/bin
```

Example:

```
[root@rc install]# cd /u01/app/21.0.0.0/grid/bin/
```

11. Change the TZ value at the database level to the desired time zone using the following command.

```
srvctl setenv database -d <database_name> -t "TZ=<new_time_zone>"
```

Example:

```
[root@rc bin]# ./srvctl setenv database -d SMDB0316_iad1cx -t "TZ=America/
Los_Angeles"
```

12. Verify the TZ value set using the following command.

```
srvctl getenv database -d <database_name>
```

Example:

```
[root@rc bin]# ./srvctl getenv database -d SMDB0316_iad1cx
SMDB0316_iad1cx:
TZ=America/Los_Angeles
```

13. Stop the CRS stack on all of the compute nodes.

```
crsctl stop cluster
```

Example:

```
[root@rc bin]# ./crsctl stop cluster
CRS-2673: Attempting to stop 'ora.crsd' on 'rc'
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
.
.
.
CRS-2673: Attempting to stop 'ora.cssd' on 'rc'
CRS-2677: Stop of 'ora.cssd' on 'rc' succeeded
[root@rc bin]#
```

14. Start the CRS stack on all of the compute nodes.

    ```
    crsctl start cluster
    ```

    Example:

    ```
    [root@rc bin]# ./crsctl start cluster
    CRS-2672: Attempting to start 'ora.cssdmonitor' on 'rc'
    .
    .
    .
    CRS-2672: Attempting to start 'ora.crsd' on 'rc'
    CRS-2676: Start of 'ora.crsd' on 'rc' succeeded
    [root@rcbin]#
    ```

15. Verify if the time zone has been updated using the following command.

    ```
    timedatectl
    ```

    Example:

    ```
    [root@rc bin]# timedatectl
             Local time: Thu 2023-03-16 01:19:15 PDT
         Universal time: Thu 2023-03-16 08:19:15 UTC
               RTC time: Thu 2023-03-16 08:19:10
              Time zone: America/Los_Angeles (PDT, -0700)
            NTP enabled: yes
       NTP synchronized: no
        RTC in local TZ: no
             DST active: yes
        Last DST change: DST began at
                         Sun 2023-03-12 01:59:59 PST
                         Sun 2023-03-12 03:00:00 PDT
        Next DST change: DST ends (the clock jumps one hour backwards) at
                         Sun 2023-11-05 01:59:59 PDT
                         Sun 2023-11-05 01:00:00 PST
    ```

    Notice that the system is now in the America/Los_Angeles (PDT) time zone.

Network Time Protocol

This article describes Network Time Protocol for Base Database Service.

Oracle recommends that you run a Network Time Protocol (NTP) daemon on your 1-node DB systems to keep system clocks stable during rebooting.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Oracle recommends that you configure NTP on both nodes in a 2-node RAC DB system to synchronize time across the nodes. If you do not configure NTP, then Oracle Clusterware configures and uses the Cluster Time Synchronization Service (CTSS), and the cluster time might be out of sync with applications that use NTP for time synchronization.

For information about configuring NTP, see Setting Network Time Protocol for Cluster Time Synchronization.

Software Images

This article provides the details and procedure to manage software images.

Software images are resources within your tenancy that you create before provisioning or updating the Oracle Grid Infrastructure or Oracle Database software.

There are two types of software image resources:

- **Grid Infrastructure software image**: Grid Infrastructure software images are resources containing Oracle Grid Infrastructure software used to update the Grid Infrastructure in the DB system.
- **Database software image**: Database software images are resources containing Oracle Database software used to provision and update Oracle Databases and Oracle Database homes in the DB system.

Software images are either Oracle-published software releases or custom software images created by you that include the desired release updates (RU) and additional one-off (interim) updates or patches referenced from specific Oracle home inventory file.

There is no limit on the number of software images you can create in your tenancy, and you can create your software images with any supported Oracle Grid Infrastructure or Oracle Database release update available in the OCI Console.

Software images are automatically stored in Oracle-managed Object Storage and can be viewed and managed in the OCI Console.

Software images are regional-level resources and can be accessed from any region within your tenancy.

> **Note**
>
> The software images incur Object Storage usage costs.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be granted security access in a **policy** by an administrator. This access is required whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you get a message that you don’t have permission or are unauthorized, verify with your administrator what type of access you have and which **compartment** to work in.

For administrators: The policy in Let database admins manage Oracle Cloud database systems lets the specified group do everything with databases and related Database resources.

If you're new to policies, see Getting Started with Policies and Common Policies. If you want to dig deeper into writing policies for databases, see Details for the Database Service.

Copyright © 2022, 2026, Oracle and/or its affiliates.

List the Software Images

Perform the following steps to view the list of software images in your tenancy using the Console:

1. Open the **navigation menu**, select **Oracle Database**, and then select **Oracle Base Database Service**.

   The **DB Systems** list page opens. All DB systems in the selected compartment are displayed in a table.

2. On the **DB Systems** list page, under **Resources**, select **Software images**.

   A list containing custom Grid Infrastructure and database software images is displayed.

3. To view the software images in a different compartment, use the **Compartment** filter to switch compartments.

Create a Grid Infrastructure Software Image

Perform the following steps to create a Grid Infrastructure software image using the Console:

1. On the **Software images** list page, select **Create software image**. If you need help finding the list page, see List the Software Images.
2. In the **Create software image** page, enter the following details:
   - For the **Image type**, select **Grid Infrastructure software image**.
   - Enter a display name for the software image. Avoid entering confidential information.
   - Select a compartment from the list.
   - Select a Grid Infrastructure release for the software image from the list.
   - Select a Grid Infrastructure version for the software image from the list. You can create a software image using any supported release update (RU).
   - Optionally, you can enter a comma-separated list of one-off patch numbers.
   - Optionally, you can upload an Oracle home inventory file from an existing DB system. For instructions on creating an inventory file using OPatch, see Verify the Updates Applied to an Oracle Home.

     > **Note**
     >
     > Oracle Home inventory contains a list of updates installed in a Grid Infrastructure or database home. To ensure that a new software image contains, at a minimum, all updates and patches installed in an existing home, you can upload a file containing the output of the "opatch lsinventory" command run in the existing home.

   - Expand **Advanced options** to provide the advanced options for this resource.
     - You can add free-form tags or defined tags to this resource. You must have permission to use the tag namespace for defined tags. For information about using tags to manage your OCI resources, see Resource Tags.
3. Select **Create**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Create a Database Software Image

Perform the following steps to create a database software image using the Console:

1. On the **Software images** list page, select **Create software image**. If you need help finding the list page, see List the Software Images.
2. In the **Create software image** page, enter the following details:
   - For the **Image type**, select **Database software image**.
   - Enter a display name for the software image. Avoid entering confidential information.
   - Select a compartment from the list.
   - Select a database release for the software image from the list.
   - Select a database version for the software image from the list. You can create a software image using any supported release update (RU).
   - Optionally, you can enter a comma-separated list of one-off patch numbers.
   - Optionally, you can upload an Oracle home inventory file from an existing DB system. For instructions on creating an inventory file using OPatch, see Verify the Updates Applied to an Oracle Home.

     > **Note**
     >
     > Oracle Home inventory contains a list of updates installed in a Grid Infrastructure or database home. To ensure that a new software image contains, at a minimum, all updates and patches installed in an existing home, you can upload a file containing the output of the "opatch lsinventory" command run in the existing home.

   - Expand **Advanced options** to provide the advanced options for this resource.
     - You can add free-form tags or defined tags to this resource. You must have permission to use the tag namespace for defined tags. For information about using tags to manage your OCI resources, see Resource Tags.
3. Select **Create**.

Create a Database Software Image from an Existing Database

Perform the following steps to create a database software image from a database in an existing DB system using the Console:

1. On the **DB Systems** list page, select the DB system that contains the database that you want to work with. If you need help finding the list page or the DB system, see *List the DB Systems*.
2. On the **Databases** tab, select the database that you want to work with.
3. On the database details page, from the **Actions** menu, select **Create image from database**.
4. Enter a display name for the software image. Avoid entering confidential information.
5. Select a compartment from the list. You can also select a different compartment from the one you are currently working in to store the database software image.

Copyright © 2022, 2026, Oracle and/or its affiliates.

6. Select **Create**.

View Details of a Software Image

You can view the details, such as version, release update, one-off (interim) updates, and so on, of the database and Grid Infrastructure software images.

Perform the following steps to view the software image details using the Console:

1. On the software images list page, select the software image that you want to view. If you need help finding the list page, see List the Software Images
2. On the software image details page, on the **Software image information** tab, you can view the following details:
   - **Lifecycle state**: The lifecycle state of the software image, such as available, provisioning, or failed.
   - **Compartment**: The compartment where the software image is located.
   - **OCID**: The unique Oracle Cloud Identifier (OCID) for this software image. You can select **Show** to view the complete OCID or select **Copy** to copy the OCID to the clipboard.
   - **Created**: Date of creation of this software image.
   - **Image type**: Type of the software image, either Grid Infrastructure or database.
   - **Service**: Name of the Oracle Cloud Service using which this software image was created.
   - **Release**: The release number of the database or Grid Infrastructure, such as 19.0.0.0 for 19c.
   - **Version**: The version number of the database or Grid Infrastructure, such as 19.22.0.0 for the database and 19.16.0.0.0 for Grid Infrastructure.
   - **One-Off updates (patches)**: The list of one-off updates included in the software image, if any. The list includes all updates specified when creating the software image (including updates listed in `lsinventory`).
     – To view all the included updates, select **View all items**. A list of all one-off updates is displayed.

Delete a Software Image

Perform the following steps to delete a software image using the Console:

1. On the software images list page, select the software image that you want to delete. If you need help finding the list page, see List the Software Images
2. On the software image details page, from the **Actions**, menu, select **Delete** and then confirm the deletion.

Move a Software Image to a Different Compartment

Perform the following steps to move a software image to a different compartment of your choice using the OCI Console:

1. On the software images list page, select the software image that you want to move. If you need help finding the list page, see List the Software Images
2. On the software image details page, from the **Actions** menu, select **Move resource**.

Copyright © 2022, 2026, Oracle and/or its affiliates.

3. Select a destination compartment from the list.
4. Select **Move resource**.

Provision a Database Using a Database Software Image

You can use a database software image to provision a DB system. For more information, see *Create a DB System*.

Update a Database Using a Database Software Image

You can use a database software image to update the database of an existing DB system. This is sometimes referred to as an "in-place update". For information about using a custom database software image to update a database in a DB system, and to determine if a database has been updated with a particular database software image, see *Update a Database*.

Update Grid Infrastructure using a Grid Infrastructure Software Image

You can use a Grid Infrastructure software image to update the Grid Infrastructure of an existing DB system. For more information, see Update the Grid Infrastructure of a DB System.

Oracle Data Guard Associations

You can use a database software image for in-place updates on both the primary and standby database instances to ensure that both databases have the same updates. For more information about Oracle Data Guard associations, see *Use Oracle Data Guard on a DB System*.

Verify the Updates Applied to an Oracle Home

The OPatch utility enables you to apply and manage interim patches for your Oracle Database software. Using the `lsinventory` command provided by OPatch, you can create a file that lists the interim patches applied to an Oracle Database Home. This file can then be uploaded to the OCI Console during the creation of a custom database software image to add the exact set of patches used by the source database home to the list of patches included in the software image. You can find the OPatch utility in the `$ORACLE_HOME/Opatch` directory. The following example shows how to use the `lsinventory` command to create the lsinventory file.

```
ORACLE_HOME/OPatch/opatch lsinventory


Oracle Interim Patch Installer version 12.2.0.1.21
Copyright (c) 2021, Oracle Corporation. All rights reserved.

Oracle Home : /u02/app/oracle/product/19.0.0.0/dbhome_2
Central Inventory : /u01/app/oraInventory
from : /u02/app/oracle/product/19.0.0.0/dbhome_2/oraInst.loc
OPatch version : 12.2.0.1.21
OUI version : 12.2.0.7.0
Log file location : /u02/app/oracle/product/19.0.0.0/dbhome_2/cfgtoollogs/
opatch/opatch2021-01-21_09-22-45AM_1.log

Lsinventory Output file location : /u02/app/oracle/product/19.0.0.0/dbhome_2/
cfgtoollogs/opatch/lsinv/lsinventory2021-01-21_09-22-45AM.txt
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Using the API

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 the following APIs to manage database updates.

- CreateDatabaseSoftwareImage
- ListDatabaseSoftwareImages
- GetDatabaseSoftwareImage
- DeleteDatabaseSoftwareImage
- ChangeDatabaseSoftwareImageCompartment

Policy Details for Database Software Images

This topic provides the details for writing Oracle Cloud Infrastructure Identity and Access Management (IAM) policies to control access to database software images used by Base Database resources.

> **Tip**
>
> For a sample policy, see Let database admins manage Oracle Cloud database systems.

Resource-Types

The `database-software-image` resource-type covers the Oracle and custom database software images available through the database software image feature. The `database-family` aggregate resource-type covers the database software image resource-type as well as other resources related to Base Database instances. For more information, see *Resource-Types* in How Policies Work.

Aggregate Resource-Type

There are no aggregate resource-types for database software images.

Individual Resource-Types

- `database-software-image` (covered under the `database-family` aggregate resource-type)

Supported Variables

Only the general variables are supported. For more information, see *General Variables for All Requests* in Policy Reference.

Details for Verb + Resource-Type Combinations

The following tables show the permissions and API operations covered by each verb. 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 preceding it, whereas "no extra" indicates no incremental access.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 7-2 Details for verb + database-software-image combinations

Verbs	Permissions	APIs fully covered	APIs partially covered
inspect	DB_SOFTWARE_IMG_INSPECT	`ListDatabaseSoftwareImages` `GetDatabaseSoftwareImage`	none
read	no extra	none	none
use	READ + DB_SOFTWARE_IMG_UPDATE	`UpdateDatabaseSoftwareImage` `ChangeDatabaseSoftwareImageCompartment`	none
manage	USE + DB_SOFTWARE_IMG_CREATE DB_SOFTWARE_IMG_DELETE	`CreateDatabaseSoftwareImage` `DeleteDatabaseSoftwareImage`	none

Permissions Required for Each API Operation

The following tables list the API operations for database software images in a logical order, grouped by resource type.

Database Software Image API Operations

Table 7-3 Database Software Image API Operations

API operation	Permissions required to use the operation
`ListDatabaseSoftwareImages`	DB_SOFTWARE_IMG_INSPECT
`GetDatabaseSoftwareImage`	DB_SOFTWARE_IMG_INSPECT
`UpdateDatabaseSoftwareImage`	DB_SOFTWARE_IMG_INSPECT and DB_SOFTWARE_IMG_UPDATE
`ChangeDatabaseSoftwareImageCompartment`	DB_SOFTWARE_IMG_INSPECT and DB_SOFTWARE_IMG_UPDATE
`CreateDatabaseSoftwareImage`	DB_SOFTWARE_IMG_INSPECT and DB_SOFTWARE_IMG_CREATE
`DeleteDatabaseSoftwareImage`	DB_SOFTWARE_IMG_INSPECT and DB_SOFTWARE_IMG_DELETE

For more information about permissions and verbs, see Advanced Policy Features.

Tags for Base Database Service Resources

Tagging is a powerful foundational service for Oracle Cloud Infrastructure (OCI) that enables users to search, control access, and do bulk actions on a set of resources based on the tag.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Importance of Tagging

Using the Oracle Cloud Infrastructure (OCI) tagging system, you can tag resources as per your organizational scheme allowing you to group resources, manage costs, and give insights into usage. Tags also help you build a governance model around security and Maximum Availability Architecture (MAA). As your organization expands its cloud footprint, it can become challenging to keep track of the deployment architectures, security best practices, MAA, application tier, etc. Using metadata tags to identify workload attributes can help keep up with the security and availability of your tenancy without cost overruns.

To enable customers to manage OCI resources securely and cost-effectively, Oracle provides a set of predefined tags in line with best practices for tagging resources. These tags are grouped into two namespaces - The `OracleStandard` namespace and the `OracleApplicationName` namespace. You can think of a tag namespace as a container for your tag keys.

Consider a scenario where your organization has multiple cloud resources such as, DB system, database, compute, network, and load balancers across multiple compartments in your tenancy. Suppose you wish to track these cloud resources for specific purposes, report on them, or take bulk actions. In that case, you will need a system that lets you group these resources based on different criteria such as environment, criticality, target users, application, etc. You can achieve this by applying appropriate tags to these resources.

For example, you may tag all resources in your development stack with `Oracle-Standard.Environment=Dev` or for a business critical application stack, set `Oracle-Standard.Criticality=High` or `Extreme`. In the event of service disruptions due to various reasons, you would then be able to quickly identify all OCI resources associated with an application or business function or be able to separate critical and non-critical workloads.

Tagging can also help you deploy optimized configurations based on workload attributes identified via tags. For example, database deployments for the Peoplesoft application require a specific configuration. By setting the `ApplicationName` and `AppMajorVersion` tags, while deploying a database, can ensure that the database is configured ready for the particular application (Example: Peoplesoft) out of the box.

Moreover, integration with the Cloud Advisor OCI service can provide you with direct, deep insight into how well your cloud services adhere to the corporate guidelines and help your management govern with a vision. For more information, see Cloud Advisor Overview.

Adding Tags

You can tag resources using the Console, the CLI, or the SDK.

There are many cloud resources that can be tagged in a DB system. DB systems, VM clusters, databases, are some of them. Tags can either be applied while creating the resources or modified later. For example, you can apply tags to an VM cluster while provisioning or add them later from its **Details** page.

Tagging integrates with OCI authorization system. You can use IAM policy controls to enable delegation or restriction of tag manipulation. For more information about the permissions required to work with defined and free-form tags, see *Authentication and Authorization* in Tagging Overview.

Your tenancies come with a library of standard tags that would apply to most resources. These tags are currently available as a set of Tag Namespaces that your governance administrators can deploy. OCI best practices recommend applying these tags to all resources a standard tag

Copyright © 2022, 2026, Oracle and/or its affiliates.

can be applied to. Besides reporting and governance, OCI service automation can deliver workload-specific optimizations based on standard tag values.

For example, database deployments for the Peoplesoft application require a specific configuration. By setting the appropriate application tag key in the `Oracle-ApplicationName` tag namespace while deploying a database, can ensure that the database is configured ready for the particular application (Example: Peoplesoft) out of the box.

Oracle Standard Tags

Your tenancy governance administrators can deploy the standard tags at the tenancy level and may also mark certain tags as required, thereby enforcing tags on resources in those compartments. The following are the standard tags defined in the namespace called `OracleStandard`. For more information, see Understanding Standard Tags.

Table 7-4 Oracle Standard Tags

Tag Key	Tag Value Options	Description
`OracleStandard.Criticality`	Extreme High Medium Low	Enables tiering of resource in line with corporate application classification standards. Customer governance can use this tag for reporting and ensuring resources are configured as per the guideline for the tier they belong to. For example, a database resource with `OracleStandard.Criticality` set to `Extreme` or `High` may require the highest availability SLA, and may need to be configured with Data Guard.
`OracleStandard.Environment`	Dev Test Prod Pre-prod Staging Trial Sandbox User Testing	Denotes a resource lifecycle. In the case of databases, it helps determine consolidation density, database distribution across containers, set maintenance plans, and manage clones.
`OracleStandard.Sensitivity`	Public Internal Sensitive Highly Sensitive Extremely Sensitive	An application or database classification tag. `OracleStandard.Sensitivity` set to `Highly Sensitive` may indicate that an access control list or certain Network Security Group (NSG) enforcement is mandatory to restrict access.
`OracleStandard.Regulation`	For values, see List of Compliance Regulations.	Denotes one or more compliance regulations that a resource must adhere to. Tag administrators may add values to the list from the OCI Governance and Administration console. For more information see Using Predefined Values.
`OracleStandard.TargetUsers`	Public Customers Partners Company Division Department Workgroup	Denotes the end users of a resource. Another form of resource classification that helps determine target users and allow governance teams to set corporate standards based on user or application type.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 7-4 (Cont.) Oracle Standard Tags

Tag Key	Tag Value Options	Description
`OracleStandard.EndUserCount`	1 10 100 1000 10000 100000 1000000 1000000 10000000	An approximate count of end-users. This tag helps determine the number of users impacted or blast radius during an availability or security event. This also help prioritize recovery efforts in the event of major outages affecting a large number of cloud resources.
`OracleStandard.OwnerEmail`	Free form tag.	Denotes the email address of the resource owner.
`OracleStandard.Org`	HR, Finance, Marketing, Sales, Legal, R&D, Customer Support, Internal Support, Manufacturing	Identifies the customer's line of business or department that owns or uses the resource. This may help with cost aggregation reports and determining usage across business units. Tag administrators may add relevant values to the list from the OCI Governance and Administration console. For more information, see Using Predefined Values.
`OracleStandard.CostCenter`	12345, WebMarketing	Freeform field for cost center.
`OracleStandard.RecoveryTimeObjectiveMinutes`	0-10080	Time in minutes. Denotes the maximum time within which the resource is required to recover from a failure.
`OracleStandard.RecoveryPointObjectiveMinutes`	0-1440	Time in minutes. Maximum data loss tolerance for a data store resource such as a database or a storage device.

List of Compliance Regulations

The table below lists the valid values that you can apply for the `OracleStandard.Regulation` tag.

Table 7-5 List of Compliance Regulations

Regulation	Description
PCI DSS	Payment Card Industry Data Security Standard
HIPAA	Health Insurance Portability and Accountability Act
ISO	International Standards Organization
SOC1	System and Organization Controls 1
SOC 2	System and Organization Controls 2
FedRamp	Federal Risk and Authorization Management Program
GLBA	Gramm–Leach–Bliley Act
CCPA	California Consumer Privacy Act
SOX	Sarbanes Oxley
NIST	National Institute of Standards and Technology - Cyber Security
FISMA	Federal Information Security Management

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table 7-5 (Cont.) List of Compliance Regulations

Regulation	Description
HITECH	Health Information Technology for Economic and Clinical Health Act
FERPA	Family Educational Rights and Privacy Act (Student privacy)
FACTA	Fair and Accurate Credit Transaction Act
Texas HB300	Texas Medical Records Privacy Act
CIS	Center for Internet Security
CJIS	Criminal Justice Information Services Security Policy
C-TPAT	Customs-Trade Partnership Against Terrorism
COPPA	Children's Online Privacy Protection Act
PIPED Act, or PIPEDA	Personal Information Protection and Electronic Documents Act
GDPR	General Data Protection Regulation
PIPL	Personal Information Protection Law

Oracle Application Name Tags

The following are the application name tags defined in the `Oracle-ApplicationName` namespace.

Table 7-6 Oracle Application Name Tags

Tag Key	Tag Value Options	Description
Hyperion	11.2 11.1	Denotes the version of the Hyperion application.
JD Edwards	9.2 9.1 9.0	Denotes the version of the JD Edwards application.
Oracle_E-Business_Suite	12.2 12.1 12.0 11i	Denotes the version of the Oracle E-Business Suite application.
PeopleSoft	9.2 9.1	Denotes the version of the Peoplesoft application.
Siebel	8.2 8.1	Denotes the version of the Siebel application.
Other_Oracle_Application	Free form tag in string format.	Can be used to denote any application other than those listed above. You can enter the application name as a string value.

This article explains how to export data from Oracle Cloud Infrastructure (OCI) Database as a Service (DBaaS) to on-premises systems or other cloud environments.

Customers using **OCI Database Services**—including **Exadata Database Service on Cloud@Customer, Exadata Database Service on Dedicated Infrastructure**, **Exadata**

Copyright © 2022, 2026, Oracle and/or its affiliates.

Database Service on Exascale Infrastructure, and **Base Database Service**—can export their data to on-premises environments or to other public cloud providers.

OCI DBaaS customers have multiple, fully supported options to move their data out of Oracle Cloud, ensuring flexibility, portability, and no lock-in. The supported export methods include:

- **Oracle Data Pump (expdp/impdp):** Export database schemas or full databases to dump files, stage them in OCI Object Storage or local file systems, and transfer them to on-premises or other cloud object stores.
- **RMAN Backup/Restore:** Create backup sets and copy them to OCI Object Storage, NFS, or directly to on-premises storage. Backups can then be restored in any Oracle Database environment outside OCI.
- **SQL*Loader/External Tables:** Unload data in formats such as CSV or Parquet to file or Object Storage for easy integration into on-premises data lakes or third-party cloud analytics platforms.
- **Oracle GoldenGate:** Replicate data in real time from OCI DBaaS to on-premises Oracle Databases or heterogeneous targets in other cloud providers.
- **OCI Database Migration Service (DMS):** A managed service that leverages Oracle GoldenGate for online or offline migrations to on-premises or other cloud providers.
- **Transportable Databases or Tablespaces:** Oracle Database files can be directly copied and mounted by an Oracle instance on-premises or in a different cloud.

With these tools, customers can move data securely and efficiently—whether they need one-time exports, continuous replication, or large-scale backups—giving them the freedom to integrate Oracle Databases into a multi-cloud or hybrid-cloud architecture.

Common Database Export Mechanisms

Regardless of which OCI DBaaS flavor is used, Oracle Databases provide multiple ways to move data:

Oracle Data Pump (expdp/impdp)

Use Data Pump Export (`expdp`) to export a schema or an entire database into dump files.

Dump files can be written to:

- OCI Object Storage
  - Using the `DBMS_CLOUD` package to stage dump files in Object Storage
  - Or via Oracle Data Pump integration with pre-authenticated URLs
- Local Filesystem (DB Node Storage)
  - Write dump files locally on the DB server node
  - Copy them out using standard tools such as `scp` or `rsync`

RMAN Backup/Restore

RMAN provides a reliable and flexible solution for backing up and restoring Oracle Databases across cloud and on-premises environments.

- Backup Options:
  - RMAN can back up databases to OCI Object Storage using the Database Backup Cloud Service module.

Copyright © 2022, 2026, Oracle and/or its affiliates.

  - Backup pieces can also be copied to on-premises environments or other object storage repositories as needed.
- Restore Options: RMAN restore can be performed in any supported target database environment.

SQL*Loader / External Tables

Use SQL*Loader or External Tables to handle partial or large-scale data unloads.

- Options for Exporting Data:
  - `UTL_FILE`: A PL/SQL package that allows the database to read and write files on the server's filesystem, commonly used to generate CSV or text files.
  - `ORACLE_DATAPUMP` access driver: Enables direct access to Data Pump export files from external tables for fast and efficient data reads.
  - `DBMS_CLOUD.EXPORT_DATA`: OCI-native PL/SQL procedure to export data directly to Object Storage in formats like CSV or Parquet.
- Data Movement: Download or migrate the generated files to on-premises environments or other cloud storage as needed.

Oracle GoldenGate

Oracle GoldenGate is a high-performance, real-time replication solution for Oracle Databases across different environments.

- Supports all three OCI DBaaS flavors.
- Can replicate transactions to:
  - On-premises Oracle Databases.
  - Oracle Databases hosted in other cloud environments.

OCI Database Migration Service (DMS)

Database Migration Service (DMS) is an OCI-native service that leverages Oracle GoldenGate to simplify data movement.

- Provides a managed pathway to migrate data from OCI DBaaS to on-premises targets or other cloud environments.
- Handles data replication, synchronization, and migration with minimal manual effort.
- Ensures data consistency and supports migrations with minimal downtime.

Transportable Databases or Tablespaces

Customers can physically copy datafiles associated with tablespaces and entire databases.

- Oracle Databases on OCI can be directly copied from OCI Block Volume via `scp`, `rsync`, or `sftp`.
- Backup pieces and snapshots can be copied/exported out of OCI Block Storage for off-cloud restore.
- Suitable for very large databases where logical export is not practical.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Transport the Data

Customers can choose from multiple options to transport exported or backed-up data out of OCI, depending on data volume, time constraints, and network availability:

- Over the Internet
  - Suitable for small- to medium-sized datasets.
  - Data can be securely copied using `scp`, `rsync`, `sftp`, or `HTTPS` uploads/downloads from Object Storage.
  - Pre-authenticated requests (PARs) allow external tools and scripts to fetch data directly without requiring IAM credentials.
  - Ensure TLS encryption in transit and use compression to optimize transfer performance.
- Direct Networking (FastConnect / Cloud Interconnect)
  - For high-throughput and low-latency transfers, customers can use private interconnects between OCI and their on-premises environment or between OCI and other cloud providers.
  - Recommended for customers with high-volume, recurring data transfers.
  - Provides a private, dedicated connection between OCI and on-premises or third-party data centers.
  - Enables faster throughput, lower latency, and enhanced security compared to internet-based transfers. Ideal for scenarios where exports or backups need to be routinely synchronized with on-premises or hybrid environments.
- Data Transfer Appliance (DTA)
  - Designed for very large-scale data movement (multi-TB to PB).
  - Encrypted appliances are physically shipped to the customer, filled with data, and returned to Oracle for direct import/export.
  - Removes dependency on limited network bandwidth.
  - Suitable for initial large migrations, off-cloud archival, or DR seeding operations.

Security and Compliance

- All exports and backups support encryption:
  - RMAN encrypted backups for secure transport.
  - Data Pump encryption (`ENCRYPTION=ALL`).
  - Oracle GoldenGate with SSL/TLS for replication streams.
- Ensure keys and wallet files are transferred securely if TDE is used.

Performance Best Practices

- For Data Pump:
  - Use the `PARALLEL` option and multiple dump files for faster throughput.
  - Write directly to OCI Object Storage with pre-authenticated URLs to avoid local storage bottlenecks.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- For RMAN:
  - Enable backup compression and parallelism.
  - Write to multiple channels for higher throughput.
- For Oracle GoldenGate:
  - Use parallel replication streams where supported.
  - Deploy Oracle GoldenGate Microservices for scale-out replication.

Service-Specific Notes

Exadata Database Service on Cloud@Customer

- Data resides physically in the customer's data center but is managed by Oracle.
- Exports/Backups can go to customer's on-premises Object Storage, NFS, or tape, in addition to OCI Object Storage.
- Oracle GoldenGate can replicate to other on-premises or cloud targets directly.

Exadata Database Service on Dedicated Infrastructure and Exascale Infrastructure

- Full Exadata Cloud environment in OCI regions.
- Data Pump, RMAN to OCI Object Storage, and Oracle GoldenGate are the primary exit paths.
- Customers typically export backup and export files to Object Storage, then download or sync to other providers.

Base Database Service

- Functions similarly to traditional on-premises databases but is hosted in OCI.
- Customers have direct SSH access to DB hosts, so they can:
  - Run `expdp` RMAN backups.
  - Transfer data using `scp`, `rsync`, or `sftp`.
  - Write backups/data files directly to mounted Object Storage.

References and Resources

Export Tools

- Oracle Data Pump
- DBMS_CLOUD Package

Backup and Recovery

- Introduction to Backup and Recovery
- Configuring the RMAN Environment
- Encrypting RMAN Backups
- Block Volume Backups

Copyright © 2022, 2026, Oracle and/or its affiliates.

Replication and Migration

- Oracle GoldenGate
- Database Migration

Known Issues

This article provides information about known issues in the Base Database service.

These known issues are identified in the Base Database Service:

- Existing pluggable databases in new database
- PDB in existing Data Guard configuration
- Billing issue when changing license type
- Backing up to Object Storage using dbcli fails due to certificate change
- Backing up to Object Storage using RMAN fails due to certificate change
- Breaking changes in Database service SDKs
- Unable to use Managed Backups in your DB system
- Managed Automatic Backups fail on the VM.Standard1.1 shape due to a process crash
- Oracle Data Pump operations return "ORA-00439: feature not enabled"
- Unable to connect to the EM Express console from your 1-node DB system
- OCI Console information not synced for Data Guard enabled databases when using dbaascli
- Grid Infrastructure does not start after offlining and onlining a disk
- DB system storage scaling failure
- DB system shape scaling failure

Existing pluggable databases in new database

Details

Existing pluggable databases (PDB) do not appear in a newly created database, and it may take up to a few hours before they appear in the OCI Console. This includes the default PDB for a new database and existing PDBs for cloned or restored databases. In case of an in-place restore to an older version, the PDB list is updated similarly and may have some delay.

Workaround

None.

PDB in existing Data Guard configuration

Details

Creating and cloning a PDB in the primary database is not allowed via the OCI Console or the API.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Workaround

You can use SQL*Plus to create or clone PDBs in the primary database, and they will be synced later in the OCI Console.

Billing issue when changing license type

Details

When you change the license type of your database or DB system from BYOL to license included, or the other way around, you are billed for both types of licenses for the first hour. After that, you are billed according to your updated license type.

Workaround

We're working on a resolution.

Backing up to Object Storage using dbcli fails due to certificate change

Details

Unmanaged backups to Object Storage using the database CLI (dbcli) fail with the following errors:

```
-> Oracle Error Codes found:
-> ORA-19554: error allocating device, device type: SBT_TAPE, device name:
-> ORA-19511: non RMAN, but media manager or vendor specific failure, error
text:
-> KBHS-00712: ORA-29024 received from local HTTP service
-> ORA-27023: skgfqsbi: media manager protocol error
```

In response to policies implemented by two common web browsers regarding Symantec certificates, Oracle recently changed the certificate authority used for Oracle Cloud Infrastructure. The resulting change in SSL certificates can cause backups to Object Storage to fail if the Oracle Database Cloud Backup Module still points to the old certificate.

Workaround

Check the log files for the errors listed and, if found, update the backup module.

1. Review the RMAN backup log files for the errors listed above:
   a. Run the following command to determine the ID of the failed backup job.

      ```
      dbcli list-jobs
      ```

      In this example output, the failed backup job ID is "f59d8470-6c37-49e4-a372-4788c984ea59".

      ```
      root@<node name> ~]# dbcli list-jobs

      ID
      Description
              Created                             Status
      ----------------------------------------
      ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
--------------------------------------------------------------------------
------- ---------------------------------- ----------
cbe852de-c0f3-4807-85e8-7523647ec78c     Authentication key update for
DCS_ADMIN                                          March 30, 2018
4:10:21 AM UTC       Success
db83fdc4-5245-4307-88a7-178f8a0efa48     Provisioning service
creation                                                   March 30,
2018 4:12:01 AM UTC       Success
c1511a7a-3c2e-4e42-9520-f156b1b4cf0e     SSH keys
update
March 30, 2018 4:48:24 AM UTC       Success
22adf146-9779-4a2c-8682-7fd04d7520b2     SSH key
delete
March 30, 2018 4:50:02 AM UTC       Success
6f2be750-9823-4ed5-b5ff-8e49f136dd22     create object
store:bV0wqIaoLA4xLT4dGjOu                                      March
30, 2018 5:33:38 AM UTC       Success
0716f464-1a10-40df-a303-cadee0302b1b     create backup
config:bV0wqIaoLA4xLT4dGjOu_BC                                  March
30, 2018 5:33:49 AM UTC       Success
e08b21c3-cd09-4e3a-944c-d1da96cb21d8     update database :
hfdb1                                                     March 30,
2018 5:34:04 AM UTC       Success
1c3d7c58-79c3-4039-8f48-787057ce7c6e     Create Longterm Backup with
TAG-DBTLongterm<identity number> for Db:<dbname>    March 30, 2018
5:37:11 AM UTC       Success
f59d8470-6c37-49e4-a372-4788c984ea59     Create Longterm Backup with
TAG-DBTLongterm<identity number> for Db:<dbname>    March 30, 2018
5:43:45 AM UTC       Failure
```

b. Use the ID of the failed job to obtain the location of the log file to review.

```
dbcli describe-job -i <failed_job_ID>
```

Relevant output from the `describe-job` command should look like this:

```
Message: DCS-10001:Internal error encountered: Failed to run Rman
statement.
Refer log in Node <node_name>: /opt/oracle/dcs/log/<node_name>/rman/
bkup/<db_unique_name>/rman_backup/<date>/rman_backup_<date>.log.
```

2. Update the Oracle Database Cloud Backup Module:

   a. Determine the Swift object store ID and user the database is using for backups.

      i. Run the `dbcli list-databases` command to determine the ID of the database.

      ii. Use the database ID to determine the backup configuration ID (`backupConfigId`).

```
dbcli list-databases
dbcli describe-database -i <database_ID> -j
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

iii. Using the backup configuration ID you noted from the previous step, determine the object store ID (`objectStoreId`).

```
dbcli list-backupconfigs
dbcli describe-backupconfig –i <backupconfig_ID> -j
```

iv. Using the object store ID you noted from the previous step, determine the object store user (`userName`).

```
dbcli list-objectstoreswifts
dbcli describe-objectstoreswift –i <objectstore_ID> -j
```

b. Using the object store credentials you obtained from step 1, update the backup module.

```
dbcli update-objectstoreswift –i <objectstore_ID> -p –u <user_name>
```

Backing up to Object Storage using RMAN fails due to certificate change

Details

Unmanaged backups to Object Storage using the RMAN fail with the following errors:

```
-> Oracle Error Codes found:
-> ORA-19554: error allocating device, device type: SBT_TAPE, device name:
-> ORA-19511: non RMAN, but media manager or vendor specific failure, error
text:
-> KBHS-00712: ORA-29024 received from local HTTP service
-> ORA-27023: skgfqsbi: media manager protocol error
```

In response to policies implemented by two common web browsers regarding Symantec certificates, Oracle recently changed the certificate authority used for Oracle Cloud Infrastructure. The resulting change in SSL certificates can cause backups to Object Storage to fail if the Oracle Database Cloud Backup Module still points to the old certificate.

Workaround

Check the RMAN log files for the error messages listed. If found, log on to the host as the oracle user, and use your Swift credentials to reinstall the backup module.

> **Note**
>
> Swift passwords are now called "Auth tokens." For more information, see Managing User Credentials.

```
java -jar <opc_install.jar_path> -opcId '<swift_user_ID>' -opcPass
'<auth_token>' -container <objectstore_container> -walletDir
<wallet_directory> -configfile <config_file> -host https://
swiftobjectstorage.<region_name>.oraclecloud.com/v1/
<object_storage_namespace> -import-all-trustcerts
```

For a multi-node DB system, perform the workaround on all nodes in the cluster.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For more information about Oracle Database Cloud Backup Module, see Using Oracle Database Backup Cloud Service.

Breaking changes in Database service SDKs

Details

The SDKs released on October 18, 2018 introduce code-breaking changes to the database size and the database edition attributes in the database backup APIs.

Workaround

Refer to the language-specific documentation for more details about the breaking changes, and update your existing code as applicable.

Unable to use Managed Backups in your DB system

Details

Backup and restore operations might not work in your DB system when you use the OCI Console or the API.

Workaround

Install the Oracle Database Cloud Backup Module, and then contact Oracle Support Services for further instructions.

Perform the following steps to install the Oracle Database Cloud Backup Module:

1. SSH to the DB system, and log in as `opc`.

   ```
   ssh -i <SSH_key> opc@<DB_system_IP address>
   login as: opc
   ```

   Alternatively, you can use `opc@<DB_system_hostname>` to log in.
2. Download the Oracle Database Cloud Backup Module from Oracle Database Cloud Backup Module.
3. Extract the contents of `opc_installer.zip` to a target directory, for example, `/home/opc`.
4. In your tenancy, create a temporary user, and grant them privileges to access the tenancy's Object Storage.
5. For this temporary user, create an auth token and note down the password. For more information about creating auth token, see Managing User Credentials.

   > **Note**
   >
   > Swift passwords are now called "Auth tokens.
6. Verify that credentials work by running the following curl command:

   ```
   curl -v -X HEAD -u  <user_id>:'<auth_token>' https://
   swiftobjectstorage.<region_name>.oraclecloud.com/v1/
   <object_storage_namespace>
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

To determine the correct region, see Object Storage FAQ.

The command should return either the `HTTP 200` or the `HTTP 204 No Content` success status response code. Any other status code indicates a problem connecting to Object Storage.

7. Run the following command:

```
java -jar opc_install.jar -opcid <user_id> -opcPass '<auth_token>' -libDir
<target_dir> -walletDir <target_dir> -host https://
swiftobjectstorage.<region_name>.oraclecloud.com/v1/
<object_storage_namespace> -configFile config.txt
```

Note that `<target_dir>` is the directory to which you extracted `opc_installer.zip` in step 3.

This command might take a few minutes to complete because it downloads `libopc.so` and other files. Once the command completes, you should see several files (including `libopc.so`) in your target directory.

8. Change directory to your target directory, and copy the `lipopc.so` and `opc_install.jar` files into the `/opt/oracle/oak/pkgrepos/oss/odbcs` directory.

```
cp libopc.so /opt/oracle/oak/pkgrepos/oss/odbcs
```

```
cp opc_install.jar /opt/oracle/oak/pkgrepos/oss/odbcs
```

(You might have to use `sudo` with the copy commands to run them as `root`.)

9. Run the following command to check whether the directory indicated exists:

```
ls /opt/oracle/dcs/commonstore/pkgrepos/oss/odbcs
```

If this directory exists, perform the following steps:

   a. Back up the files in the `/opt/oracle/dcs/commonstore/pkgrepos/oss/odbcs` directory.
   b. Run these two commands to replace the existing `libopc.so` and `opc_install.jar` files in that directory:

```
cp libopc.so /opt/oracle/dcs/commonstore/pkgrepos/oss/odbcs
cp opc_install.jar /opt/oracle/dcs/commonstore/pkgrepos/oss/odbcs
```

10. Verify the version of `opc_install.jar`.

```
java -jar /opt/oracle/oak/pkgrepos/oss/odbcs/opc_install.jar |grep -i build
```

If `/opt/oracle/dcs/commonstore/pkgrepos/oss/odbcs` exists, also run the following command:

```
java -jar /opt/oracle/dcs/commonstore/pkgrepos/oss/odbcs/opc_install.jar |
grep -i build
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Both commands should return the following output:

```
Oracle Database Cloud Backup Module Install Tool, build MAIN_2017-08-16.
```

11. (Optional) Delete the temporary user and the target directory you used to install the backup module.

After you complete the procedure, contact Oracle Support or your tenant administrator for further instructions. You must provide the OCID of the DB system for which you would like to enable backups.

Managed Automatic Backups fail on the VM.Standard1.1 shape due to a process crash

Details

Memory limitations of host machines running the VM.Standard1.1 shape can cause failures for automatic database backup jobs managed by Oracle Cloud Infrastructure (jobs managed by using either the OCI Console or the API). You can change the system's memory parameters to resolve this issue.

Workaround

Change the systems' memory parameters as follows:

1. Switch to the oracle user in the operating system.

```
[opc@hostname ~]$ sudo su - oracle
```

2. Set the environment variable to login to the database instance. For example:

```
oracle@hostname ~]$ . oraenv
ORACLE_SID = [oracle] ? orcl
```

3. Start SQL*Plus.

```
[oracle@hostname ~]$ sqlplus / as sysdba
```

4. Change the initial memory parameters as follows:

```
SQL> ALTER SYSTEM SET SGA_TARGET = 1228M scope=spfile;
SQL> ALTER SYSTEM SET PGA_AGGREGATE_TARGET = 1228M;
SQL> ALTER SYSTEM SET PGA_AGGREGATE_LIMIT = 2457M;
SQL> exit
```

5. Restart the database instance.

```
[oracle@hostname ~]$ srvctl stop database -d db_unique_name -o immediate
[oracle@hostname ~]$ srvctl start database -d db_unique_name -o open
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Oracle Data Pump operations return "ORA-00439: feature not enabled"

Details

On High Performance and Extreme Performance DB systems, Data Pump utility operations that use compression and/or parallelism might fail and return the error ORA-00439: feature not enabled. This issue affects database versions 12.1.0.2.161018 and 12.1.0.2.170117.

Workaround

Apply patch 25579568 or 25891266 to Oracle Database homes for database versions 12.1.0.2.161018 or 12.1.0.2.170117, respectively. Alternatively, use the OCI Console to apply the April 2017 patch to the DB system and database home.

> **Note**
>
> To determine the version of a database in a database home, run either `$ORACLE_HOME/OPatch/opatch lspatches` as the *oracle* user or `dbcli list-dbhomes` as the *root* user.

Unable to connect to the EM Express console from your 1-node DB system

Details

You might get a "Secure Connection Failed" error message when you try to connect to the EM Express console from your 1-node DB system because the correct permissions were not applied automatically.

Workaround

Add read permissions for the `asmadmin` group on the wallet directory of the DB system, and then retry the connection. Perform the following steps to add read permissions.

1. SSH to the DB system host, log in as `opc`, and `sudo` to the `grid` user.

   ```
   [opc@dbsysHost ~]$ sudo su - grid
   [grid@dbsysHost ~]$ . oraenv
   ORACLE_SID = [+ASM1] ?
   The Oracle base has been set to /u01/app/grid
   ```

2. Get the location of the `wallet` directory. It is the value of the `my_wallet_directory` in the following command output.

   ```
   [grid@dbsysHost ~]$ lsnrctl status | grep xdb_wallet

   (DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)
   (HOST=dbsysHost.sub04061528182.dbsysapril6.oraclevcn.com)(PORT=5500))
   (Security=(my_wallet_directory=/u01/app/oracle/admin/dbsys12_phx3wm/
   xdb_wallet))(Presentation=HTTP)(Session=RAW))
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

3. Return to the `opc` user, switch to the `oracle` user, and change to the `wallet` directory.

```
[opc@dbsysHost ~]$ sudo su - oracle
[oracle@dbsysHost ~]$ cd /u01/app/oracle/admin/dbsys12_phx3wm/xdb_wallet
```

4. List the directory contents and note the permissions.

```
[oracle@dbsysHost xdb_wallet]$ ls -ltr
total 8
-rw------- 1 oracle asmadmin 3881 Apr  6 16:32 ewallet.p12
-rw------- 1 oracle asmadmin 3926 Apr  6 16:32 cwallet.sso
```

5. Change the permissions.

```
[oracle@dbsysHost xdb_wallet]$ chmod 640 /u01/app/oracle/admin/
dbsys12_phx3wm/xdb_wallet/*
```

6. Verify that read permissions were added.

```
[oracle@dbsysHost xdb_wallet]$ ls -ltr
total 8
-rw-r----- 1 oracle asmadmin 3881 Apr  6 16:32 ewallet.p12
-rw-r----- 1 oracle asmadmin 3926 Apr  6 16:32 cwallet.sso
```

OCI Console information not synced for Data Guard enabled databases when using dbaascli

Details

The OCI Console synchronizes and displays information about databases that are created and managed by using the `dbaasapi` and `dbaascli` utilities. However, databases with Data Guard configured do not display correct information in the OCI Console under the following conditions:

- If Data Guard was enabled by using the OCI Console, and then a change is made to the primary or standby database by using `dbaascli` (such as moving the database to a different home), the result is not reflected in the OCI Console.
- If Data Guard was configured manually, the OCI Console does not show a Data Guard association between the two databases.

Workaround

We are aware of the issue and working on a resolution. In the meantime, Oracle recommends that you manage your Data Guard enabled databases by using either only the OCI Console or only command line utilities.

Grid Infrastructure does not start after offlining and onlining a disk

Details

This is a clusterware issue that occurs only when the Oracle Grid Infrastructure (GI) version is 12.2.0.1 without any bundle patch. The problem is caused by corruption of a voting disk after you offline then online the disk.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Workaround

Determine the version of the GI, and whether the voting disk is corrupted. Repair the disk, if applicable, and then apply the latest GI bundle.

Perform the following steps:

1. Verify the GI version is 12.2.0.1 without any bundle patch applied:

```
[root@rmstest-udaau1 ~]# su - grid
[grid@rmstest-udaau1 ~]$ . oraenv
ORACLE_SID = [+ASM1] ? +ASM1
The Oracle base has been set to /u01/app/grid
[grid@rmstest-udaau1 ~]$ $ORACLE_HOME/OPatch/opatch lsinventory
Oracle Interim Patch Installer version 12.2.0.1.6
Copyright (c) 2018, Oracle Corporation.  All rights reserved.

Oracle Home       : /u01/app/12.2.0.1/grid
Central Inventory : /u01/app/oraInventory
   from           : /u01/app/12.2.0.1/grid/oraInst.loc
OPatch version    : 12.2.0.1.6
OUI version       : 12.2.0.1.4
Log file location : /u01/app/12.2.0.1/grid/cfgtoollogs/opatch/
opatch2018-01-15_22-11-10PM_1.log

Lsinventory Output file location : /u01/app/12.2.0.1/grid/cfgtoollogs/
opatch/lsinv/lsinventory2018-01-15_22-11-10PM.txt

--------------------------------------------------------------------------
-----
Local Machine Information::
Hostname: rmstest-udaau1.exaagclient.sretest.oraclevcn.com
ARU platform id: 226
ARU platform description:: Linux x86-64

Installed Top-level Products (1):

Oracle Grid Infrastructure 12c
12.2.0.1.0
There are 1 products installed in this Oracle Home.

There are no Interim patches installed in this Oracle Home.

--------------------------------------------------------------------------
-----

OPatch succeeded.
```

2. Check the `/u01/app/grid/diag/crs/<hostname>/crs/trace/ocssd.trc` file for evidence that the GI failed to start due to voting disk corruption:

```
ocssd.trc

2017-01-17 23:45:11.955 :    CSSD:3807860480: clssnmvDiskCheck::
configured
Sites = 1, Incative sites = 1, Mininum Sites required = 1
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
2017-01-17 23:45:11.955 :    CSSD:3807860480:
(:CSSNM00018:)clssnmvDiskCheck:
Aborting, 2 of 5 configured voting disks available, need 3
......
.
2017-01-17 23:45:11.956 :    CSSD:3807860480:
clssnmCheckForNetworkFailure:
skipping 31 defined 0
2017-01-17 23:45:11.956 :    CSSD:3807860480: clssnmRemoveNodeInTerm: node
4,
slcc05db08 terminated. Removing from its own member and connected bitmaps
2017-01-17 23:45:11.956 :    CSSD:3807860480:
###################################
2017-01-17 23:45:11.956 :    CSSD:3807860480: clssscExit: CSSD aborting
from
thread clssnmvDiskPingMonitorThread
2017-01-17 23:45:11.956 :    CSSD:3807860480:
###################################
2017-01-17 23:45:11.956 :    CSSD:3807860480: (:CSSSC00012:)clssscExit: A
fatal error occurred and the CSS daemon is terminating abnormally

------------

2017-01-19 19:00:32.689 :    CSSD:3469420288: clssnmFindVF: Duplicate
voting disk found in the queue of previously configured disks
queued(o/192.168.10.18/PCW_CD_02_slcc05cel10|[66223efc-29254fbb-
bf901601-21009
cbd]),
found(o/192.168.10.18/PCW_CD_02_slcc05cel10|[66223efc-29254fbb-
bf901601-21009c
bd]), is not corrupted
2017-01-19 19:01:06.467 :    CSSD:3452057344: clssnmvVoteDiskValidation:
Voting disk(o/192.168.10.19/PCW_CD_02_slcc05cel11) is corrupted
```

3. You can also use SQL*Plus to confirm that the voting disks are corrupted:
   a. Log in as the grid user, and set the environment to ASM.

      ```
      [root@rmstest-udaau1 ~]# su - grid
      [grid@rmstest-udaau1 ~]$ . oraenv
      ORACLE_SID = [+ASM1] ? +ASM1
      The Oracle base has been set to /u01/app/grid
      ```

   b. Log in to SQL*Plus as `SYSASM`.

      ```
      $ORACLE_HOME/bin/sqlplus / as sysasm
      ```

   c. Run the following two queries:

      ```
      SQL> select name, voting_file from v$asm_disk where VOTING_FILE='Y' and
      group_number !=0;
      SQL> select  CC.name, count(*) from x$kfdat AA JOIN (select
      disk_number, name from v$asm_disk where VOTING_FILE='Y' and
      group_number !=0) CC ON CC.disk_number = AA.NUMBER_KFDAT where
      AA.FNUM_KFDAT= 1048572 group by CC.name;
      ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

If the system is healthy, the results should look like the following example.

```
Query 1 Results
NAME                           VOTING_FILE
------------------------------ --------------
DBFSC3_CD_02_SLCLCX0788        Y
DBFSC3_CD_09_SLCLCX0787        Y
DBFSC3_CD_04_SLCLCX0786        Y

Query 2 Results
NAME                           COUNT(*)
------------------------------ --------------
DBFSC3_CD_02_SLCLCX0788        8
DBFSC3_CD_09_SLCLCX0787        8
DBFSC3_CD_04_SLCLCX0786        8
```

In a healthy system, every voting disk returned in the first query should also be returned in the second query and the counts for all the disks should be non-zero. Otherwise, one or more of your voting disks are corrupted.

4. If a voting disks is corrupted, offline the grid disk that contains the voting disk. The cells will automatically move the bad voting disk to the other grid disk and online that voting disk.

   a. The following command offlines a grid disk named DATAC01_CD_05_SCAQAE08CELADM13.

   ```
   SQL> alter diskgroup DATAC01 offline disk
   DATAC01_CD_05_SCAQAE08CELADM13;
        Diskgroup altered.
   ```

   b. Wait 30 seconds and then rerun the two queries in step 3c to verify that the voting disk migrated to the new grid disk and that it is healthy.

   c. Verify the grid disk you offlined is now online:

   ```
   SQL> select name, mode_status, voting_file from v$asm_disk where
   name='DATAC01_CD_05_SCAQAE08CELADM13';
   ```

   The mode_status should be ONLINE, and the voting_file should NOT be Y.

   Repeat steps 4a through 4c for each remaining grid disk that contains a corrupt voting disk.

> **Note**
>
> If the CRS does not start because of the voting disk corruption, start it using Exclusive mode before you execute the command in step 4.
>
> ```
> crsctl start crs -excl
> ```

5. If you are using Oracle GI version 12.2.0.1 without any bundle patch, you must upgrade the GI version to the latest GI bundle, whether or not a voting disk was corrupted.

Copyright © 2022, 2026, Oracle and/or its affiliates.

DB system storage scaling failure

Details

Attempting to scale from a value less than 10,240 GB to a value higher than 10,240 GB in a single operation can lead to a failure of the scaling operation.

Workaround

If you are scaling either regular data storage or recovery area (RECO) storage from a value less than 10,240 GB (10 TB) to a value exceeding 10,240 GB, perform the scaling in two operations. First, scale the system to 10,240 GB. After this first scaling operation is complete and the system is in the "available" state, perform a second scaling operation, specifying your target storage value above 10,240 GB.

DB system shape scaling failure

Details

When scaling a DB system to use a larger system shape, the scaling operation fails if a `DB_Cache_nX` parameter is not set to 0 (zero).

Workaround

When scaling a DB system, ensure that all `DB_Cache_nX` parameters (for example, `DB_nK_CACHE_SIZE`) are set to 0 (zero).

Troubleshoot

This article provides troubleshooting steps to assess, manage, and monitor the health of the DCS stack.

These steps help identify service-level issues within the DCS stack and assist in restoring normal operation.

Check, stop, and start the DCS services

1. List installed RPM packages for DCS services.

   ```
   rpm -qa --last | egrep "dcs|mysql|zoo"
   ```

   Example:

   ```
   [root@test ~]$ rpm -qa --last | egrep "dcs|mysql|zoo"
   dcs-agent-25.4.1.0.0_251008.0630-4.x86_64     Mon Nov 24 00:00:32 2025
   dcs-admin-25.4.1.0.0_251008.0630-1.x86_64     Sat Nov 22 04:03:02 2025
   mysql-8.0.44_el7-8.0.44.x86_64                Sat Nov 22 00:00:29 2025
   zookeeper-3.9.3-1.x86_64                      Thu Feb 27 00:05:21 2025
   dcs-cli-24.2.1.0.0_240524.1355-1.x86_64       Tue Jun  4 05:40:19 2024
   ```

2. Check, stop, start, and restart the DCS Agent service.

Copyright © 2022, 2026, Oracle and/or its affiliates.

a. List the status of DCS stack services.

```
systemctl --state=running|egrep -i "dcs|mysql|zoo|tfa"
```

Example:

```
[root@test ~]# systemctl --state=running|egrep -i "dcs|mysql|zoo|tfa"
initdcsadmin.service     loaded active running Oracle dcs-admin startup
initdcsagent.service     loaded active running Oracle dcs-agent startup
mysqld.service           loaded active running MySQL Server
oracle-tfa.service       loaded active running Oracle Trace File
Analyzer
```

b. List all processes that are listening on DCS designated ports.

```
sudo lsof -nP -iTCP -sTCP:LISTEN | grep -E ':(7060|7070|3306)\b'
```

Example:

```
[root@test ~]# sudo lsof -nP -iTCP -sTCP:LISTEN | grep -E ':(7060|7070|
3306)\b'
java    63835   root  403u  IPv4 3563681668    0t0  TCP *:7070 (LISTEN)
mysqld  76121  mysql  120u  IPv4 1268462598    0t0  TCP 127.0.0.1:3306
(LISTEN)
java    76285   root  383u  IPv4 3563956774    0t0  TCP *:7060 (LISTEN)
```

c. Verify DCS Agent status using `dbcli`.

```
dbcli ping-agent
```

Example:

```
[root@test ~]# dbcli ping-agent
Agent is ready to serve the requests.
```

d. Check, stop, start, and restart the MySQL service.

```
systemctl status mysqld
systemctl stop mysqld
systemctl start mysqld
systemctl restart mysqld
```

Example:

```
[root@test ~]# systemctl status mysqld
mysqld.service - MySQL Server
   Loaded: loaded (/etc/systemd/system/mysqld.service; enabled; vendor
preset: disabled)
   Active: active (running) since Sat 2025-11-22 00:00:37 UTC; 3 months
4 days ago
     Docs: man:mysqld(8)
           http://dev.mysql.com/doc/refman/en/using-systemd.html
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
 Main PID: 76121 (mysqld)
   Status: "Server is operational"
    Tasks: 36 (limit: 79998)
   Memory: 218.0M
   CGroup: /system.slice/mysqld.service
           76121 /opt/oracle/mysql_compact/mysql/bin/mysqld
                --defaults-file=/opt/oracle/mysql_compact/my.sim.cnf

Warning: Journal has been rotated since unit was started.
Log output is incomplete or unavailable.

[root@test ~]# systemctl stop mysqld
[root@test ~]# systemctl status mysqld
mysqld.service - MySQL Server
   Loaded: loaded (/etc/systemd/system/mysqld.service; enabled; vendor
preset: disabled)
   Active: inactive (dead) since Thu 2026-02-26 01:22:51 UTC; 3s ago
     Docs: man:mysqld(8)
           http://dev.mysql.com/doc/refman/en/using-systemd.html
  Process: 76121 ExecStart=/opt/oracle/mysql_compact/mysql/bin/mysqld
                --defaults-file=/opt/oracle/mysql_compact/my.sim.cnf
                $MYSQLD_OPTS (code=exited, status=0/SUCCESS)
 Main PID: 76121 (code=exited, status=0/SUCCESS)
   Status: "Server shutdown complete"

Feb 26 01:22:47 test systemd[1]: Stopping MySQL Server...
Feb 26 01:22:51 test systemd[1]: mysqld.service: Succeeded.
Feb 26 01:22:51 test systemd[1]: Stopped MySQL Server.
Warning: Journal has been rotated since unit was started.
Log output is incomplete or unavailable.

[root@test ~]# systemctl start mysqld
[root@test ~]# systemctl status mysqld
mysqld.service - MySQL Server
   Loaded: loaded (/etc/systemd/system/mysqld.service; enabled; vendor
preset: disabled)
   Active: active (running) since Thu 2026-02-26 01:23:02 UTC; 3s ago
     Docs: man:mysqld(8)
           http://dev.mysql.com/doc/refman/en/using-systemd.html
 Main PID: 86644 (mysqld)
   Status: "Server is operational"
    Tasks: 33 (limit: 79998)
   Memory: 111.7M
   CGroup: /system.slice/mysqld.service
           86644 /opt/oracle/mysql_compact/mysql/bin/mysqld
                --defaults-file=/opt/oracle/mysql_compact/my.sim.cnf

Feb 26 01:23:01 test systemd[1]: Starting MySQL Server...
Feb 26 01:23:02 test systemd[1]: Started MySQL Server.

[root@test ~]# systemctl restart mysqld
[root@test ~]# systemctl status mysqld
mysqld.service - MySQL Server
   Loaded: loaded (/etc/systemd/system/mysqld.service; enabled; vendor
preset: disabled)
   Active: active (running) since Thu 2026-02-26 01:23:14 UTC; 2s ago
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
     Docs: man:mysqld(8)
           http://dev.mysql.com/doc/refman/en/using-systemd.html
 Main PID: 86821 (mysqld)
   Status: "Server is operational"
    Tasks: 33 (limit: 79998)
   Memory: 107.5M
   CGroup: /system.slice/mysqld.service
           86821 /opt/oracle/mysql_compact/mysql/bin/mysqld
                --defaults-file=/opt/oracle/mysql_compact/my.sim.cnf

Feb 26 01:23:13 test systemd[1]: Starting MySQL Server...
Feb 26 01:23:14 test systemd[1]: Started MySQL Server.
```

e. Check, stop, start, and restart the DCS Agent service.

```
systemctl status initdcsagent
systemctl stop initdcsagent
systemctl start initdcsagent
systemctl restart initdcsagent
```

Example:

```
[root@test ~]# systemctl status initdcsagent
initdcsagent.service - Oracle dcs-agent startup
   Loaded: loaded (/etc/systemd/system/initdcsagent.service; enabled;
vendor preset: disabled)
   Active: active (running) since Mon 2025-11-24 00:00:50 UTC; 3 months
2 days ago
 Main PID: 7956 (sh)
    Tasks: 1353 (limit: 79998)
   Memory: 21.5G
   CGroup: /system.slice/initdcsagent.service
           7956 /bin/sh -c . /opt/oracle/dcs/bin/setupJre.sh;
        $JAVA -Xms128m -Xmx512m -XX:MetaspaceSize=128m -
XX:MaxMetaspaceSize=512m
        -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDum>
           7958 /opt/oracle/dcs/java/dcs-agent/jdk-11.0.29/bin/java
        -Xms128m -Xmx512m -XX:MetaspaceSize=128m -
XX:MaxMetaspaceSize=512m
        -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDu>

Feb 25 18:30:02 test su[27582]: pam_unix(su:session): session closed
for user oracle
Feb 25 18:30:03 test su[27690]: (to oracle) root on none
Feb 25 18:30:03 test su[27690]: pam_unix(su:session): session opened
for user oracle by (uid=0)
Feb 25 18:30:03 test su[27690]: pam_unix(su:session): session closed
for user oracle
Feb 25 18:30:03 test su[27694]: (to oracle) root on none
Feb 25 18:30:03 test su[27694]: pam_unix(su:session): session opened
for user oracle by (uid=0)
Feb 25 18:30:03 test su[27694]: pam_unix(su:session): session closed
for user oracle
Feb 25 18:30:03 test su[27752]: (to oracle) root on none
Feb 25 18:30:03 test su[27752]: pam_unix(su:session): session opened
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
for user oracle by (uid=0)
Feb 25 18:30:04 test su[27752]: pam_unix(su:session): session closed
for user oracle

[root@test ~]# systemctl stop initdcsagent
[root@test ~]# systemctl status initdcsagent
initdcsagent.service - Oracle dcs-agent startup
   Loaded: loaded (/etc/systemd/system/initdcsagent.service; enabled;
vendor preset: disabled)
   Active: inactive (dead) since Wed 2026-02-25 19:09:22 UTC; 11s ago
  Process: 7956 ExecStart=/bin/sh -c . /opt/oracle/dcs/bin/setupJre.sh;
        $JAVA -Xms128m -Xmx512m -XX:MetaspaceSize=128m -
XX:MaxMetaspaceSize=512m
        -XX:+HeapDumpOnOutOfMemoryError -XX>
 Main PID: 7956 (code=killed, signal=TERM)

Feb 25 19:05:10 test su[58800]: pam_unix(su:session): session closed
for user grid
Feb 25 19:05:10 test su[58839]: (to grid) root on none
Feb 25 19:05:10 test su[58839]: pam_unix(su:session): session opened
for user grid by (uid=0)
Feb 25 19:05:10 test su[58839]: pam_unix(su:session): session closed
for user grid
Feb 25 19:05:15 test su[59381]: (to grid) root on none
Feb 25 19:05:15 test su[59381]: pam_unix(su:session): session opened
for user grid by (uid=0)
Feb 25 19:05:15 test su[59381]: pam_unix(su:session): session closed
for user grid
Feb 25 19:09:21 test systemd[1]: Stopping Oracle dcs-agent startup...
Feb 25 19:09:22 test systemd[1]: initdcsagent.service: Succeeded.
Feb 25 19:09:22 test systemd[1]: Stopped Oracle dcs-agent startup.

[root@test ~]# systemctl start initdcsagent
[root@test ~]# systemctl status initdcsagent
initdcsagent.service - Oracle dcs-agent startup
   Loaded: loaded (/etc/systemd/system/initdcsagent.service; enabled;
vendor preset: disabled)
   Active: active (running) since Wed 2026-02-25 19:09:51 UTC; 4s ago
 Main PID: 62798 (sh)
    Tasks: 1268 (limit: 79998)
   Memory: 20.2G
   CGroup: /system.slice/initdcsagent.service
           62798 /bin/sh -c . /opt/oracle/dcs/bin/setupJre.sh;
        $JAVA -Xms128m -Xmx512m -XX:MetaspaceSize=128m -
XX:MaxMetaspaceSize=512m
        -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDu>
           62799 /opt/oracle/dcs/java/dcs-agent/jdk-11.0.29/bin/java
        -Xms128m -Xmx512m -XX:MetaspaceSize=128m -
XX:MaxMetaspaceSize=512m
        -XX:+HeapDumpOnOutOfMemoryError -XX:HeapD>

Feb 25 19:09:51 test systemd[1]: Started Oracle dcs-agent startup.

[root@test ~]# systemctl restart initdcsagent
[root@test ~]# systemctl status initdcsagent
initdcsagent.service - Oracle dcs-agent startup
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
   Loaded: loaded (/etc/systemd/system/initdcsagent.service; enabled;
vendor preset: disabled)
   Active: active (running) since Wed 2026-02-25 19:10:09 UTC; 11s ago
 Main PID: 63834 (sh)
    Tasks: 1192 (limit: 79998)
   Memory: 20.2G
   CGroup: /system.slice/initdcsagent.service
           63834 /bin/sh -c . /opt/oracle/dcs/bin/setupJre.sh;
        $JAVA -Xms128m -Xmx512m -XX:MetaspaceSize=128m -
XX:MaxMetaspaceSize=512m
        -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDu>
           63835 /opt/oracle/dcs/java/dcs-agent/jdk-11.0.29/bin/java -
Xms128m -Xmx512m
        -XX:MetaspaceSize=128m -XX:MaxMetaspaceSize=512m
        -XX:+HeapDumpOnOutOfMemoryError -XX:HeapD>

Feb 25 19:10:16 test su[64523]: pam_unix(su:session): session closed
for user oracle
Feb 25 19:10:17 test su[64648]: (to oracle) root on none
Feb 25 19:10:17 test su[64648]: pam_unix(su:session): session opened
for user oracle by (uid=0)
Feb 25 19:10:17 test su[64648]: pam_unix(su:session): session closed
for user oracle
Feb 25 19:10:17 test su[64652]: (to oracle) root on none
Feb 25 19:10:17 test su[64652]: pam_unix(su:session): session opened
for user oracle by (uid=0)
Feb 25 19:10:17 test su[64652]: pam_unix(su:session): session closed
for user oracle
Feb 25 19:10:17 test su[64865]: (to oracle) root on none
Feb 25 19:10:17 test su[64865]: pam_unix(su:session): session opened
for user oracle by (uid=0)
Feb 25 19:10:18 test su[64865]: pam_unix(su:session): session closed
for user oracle
```

f. Check, stop, start, and restart the DCS Admin service.

```
systemctl status initdcsadmin
systemctl stop initdcsadmin
systemctl start initdcsadmin
systemctl restart initdcsadmin
```

Example:

```
[root@test ~]# systemctl status initdcsadmin
initdcsadmin.service - Oracle dcs-admin startup
   Loaded: loaded (/etc/systemd/system/initdcsadmin.service; enabled;
vendor preset: disabled)
   Active: active (running) since Sat 2025-11-22 03:17:28 UTC; 3 months
4 days ago
 Main PID: 48268 (sh)
   Tasks: 313 (limit: 79998)
   Memory: 1.0G
   CGroup: /system.slice/initdcsadmin.service
           48268 /bin/sh -c . /opt/oracle/dcs/bin/setupAdminJre.sh;
        $JAVA -Xms128m -Xmx512m -XX:MetaspaceSize=128m -
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
XX:MaxMetaspaceSize=512m
        -XX:+DisableExplicitGC -XX:ParallelGC>
           48269 /opt/oracle/dcs/java/dcs-admin/jdk-11.0.29/bin/java -
Xms128m -Xmx512m
        -XX:MetaspaceSize=128m -XX:MaxMetaspaceSize=512m
        -XX:+DisableExplicitGC -XX:ParallelGCThre>

Feb 25 11:13:58 test su[50418]: (to mysql) root on none
Feb 25 11:13:58 test su[50418]: pam_unix(su:session): session opened
for user mysql by (uid=0)
Feb 25 11:13:58 test su[50418]: pam_unix(su:session): session closed
for user mysql
Feb 25 17:13:59 test su[62504]: (to mysql) root on none
Feb 25 17:13:59 test su[62504]: pam_unix(su:session): session opened
for user mysql by (uid=0)
Feb 25 17:13:59 test su[62504]: pam_unix(su:session): session closed
for user mysql
Warning: Journal has been rotated since unit was started. Log output is
incomplete or unavailable.

[root@test ~]# systemctl stop initdcsadmin
[root@test ~]# systemctl status initdcsadmin
initdcsadmin.service - Oracle dcs-admin startup
 Loaded: loaded (/etc/systemd/system/initdcsadmin.service; enabled;
vendor preset: disabled)
 Active: inactive (dead) since Wed 2026-02-25 19:21:44 UTC; 1s ago
 Process: 48268 ExecStart=/bin/sh -c . /opt/oracle/dcs/bin/
setupAdminJre.sh;
    $JAVA -Xms128m -Xmx512m -XX:MetaspaceSize=128m -
XX:MaxMetaspaceSize=512m
    -XX:+DisableExplicitGC -XX:Pa>
 Main PID: 48268 (code=killed, signal=TERM)

Feb 25 11:13:58 test su[50418]: (to mysql) root on none
Feb 25 11:13:58 test su[50418]: pam_unix(su:session): session opened
for user mysql by (uid=0)
Feb 25 11:13:58 test su[50418]: pam_unix(su:session): session closed
for user mysql
Feb 25 17:13:59 test su[62504]: (to mysql) root on none
Feb 25 17:13:59 test su[62504]: pam_unix(su:session): session opened
for user mysql by (uid=0)
Feb 25 17:13:59 test su[62504]: pam_unix(su:session): session closed
for user mysql
Feb 25 19:21:43 test systemd[1]: Stopping Oracle dcs-admin startup...
Feb 25 19:21:44 test systemd[1]: initdcsadmin.service: Succeeded.
Feb 25 19:21:44 test systemd[1]: Stopped Oracle dcs-admin startup.
Warning: Journal has been rotated since unit was started. Log output is
incomplete or unavailable.

[root@test ~]# systemctl start initdcsadmin
[root@test ~]# systemctl status initdcsadmin
initdcsadmin.service - Oracle dcs-admin startup
   Loaded: loaded (/etc/systemd/system/initdcsadmin.service; enabled;
vendor preset: disabled)
   Active: active (running) since Wed 2026-02-25 19:23:00 UTC; 2s ago
 Main PID: 75371 (sh)
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
    Tasks: 162 (limit: 79998)
   Memory: 226.0M
   CGroup: /system.slice/initdcsadmin.service
           75371 /bin/sh -c . /opt/oracle/dcs/bin/setupAdminJre.sh;
         $JAVA -Xms128m -Xmx512m -XX:MetaspaceSize=128m -
XX:MaxMetaspaceSize=512m
         -XX:+DisableExplicitGC -XX:ParallelGC>
           75372 /opt/oracle/dcs/java/dcs-admin/jdk-11.0.29/bin/java -
Xms128m -Xmx512m
         -XX:MetaspaceSize=128m -XX:MaxMetaspaceSize=512m
         -XX:+DisableExplicitGC -XX:ParallelGCThre>

Feb 25 19:23:00 test systemd[1]: Started Oracle dcs-admin startup.

[root@test ~]# systemctl restart  initdcsadmin
[root@test ~]# systemctl status initdcsadmin
initdcsadmin.service - Oracle dcs-admin startup
   Loaded: loaded (/etc/systemd/system/initdcsadmin.service; enabled;
vendor preset: disabled)
   Active: active (running) since Wed 2026-02-25 19:24:06 UTC; 2s ago
 Main PID: 76284 (sh)
    Tasks: 161 (limit: 79998)
   Memory: 226.9M
   CGroup: /system.slice/initdcsadmin.service
           76284 /bin/sh -c . /opt/oracle/dcs/bin/setupAdminJre.sh;
         $JAVA -Xms128m -Xmx512m -XX:MetaspaceSize=128m -
XX:MaxMetaspaceSize=512m
         -XX:+DisableExplicitGC -XX:ParallelGC>
           76285 /opt/oracle/dcs/java/dcs-admin/jdk-11.0.29/bin/java -
Xms128m -Xmx512m
         -XX:MetaspaceSize=128m -XX:MaxMetaspaceSize=512m
         -XX:+DisableExplicitGC -XX:ParallelGCThre>

Feb 25 19:24:06 test systemd[1]: Started Oracle dcs-admin startup.
```

g. Check, stop, start, and restart the ZooKeeper service.

```
/opt/zookeeper/bin/zkServer.sh status
/opt/zookeeper/bin/zkServer.sh stop
/opt/zookeeper/bin/zkServer.sh start
```

Example:

```
[root@test ~]# /opt/zookeeper/bin/zkServer.sh status
ZooKeeper JMX enabled by default
Using config: /opt/zookeeper/bin/../conf/zoo.cfg
Mode: follower

[root@test ~]# /opt/zookeeper/bin/zkServer.sh stop
ZooKeeper JMX enabled by default
Using config: /opt/zookeeper/bin/../conf/zoo.cfg
Stopping zookeeper ... STOPPED

[root@test ~]# /opt/zookeeper/bin/zkServer.sh status
ZooKeeper JMX enabled by default
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
Using config: /opt/zookeeper/bin/../conf/zoo.cfg
Error contacting service. It is probably not running.

[root@test ~]# /opt/zookeeper/bin/zkServer.sh start
ZooKeeper JMX enabled by default
Using config: /opt/zookeeper/bin/../conf/zoo.cfg
Starting zookeeper ... STARTED

[root@test ~]# /opt/zookeeper/bin/zkServer.sh status
ZooKeeper JMX enabled by default
Using config: /opt/zookeeper/bin/../conf/zoo.cfg
Mode: follower
```

Display DB system details

1. Display DB system details.

```
/opt/oracle/dcs/bin/dbcli describe-system -d
```

Example:

```
[root@test ~]# /opt/oracle/dcs/bin/dbcli describe-system -d

DbSystem Information
----------------------------------------------------------------
                     ID: 444b444e-c4c4-4fd4-44er-4444ag44444
               Platform: Vmdb
        Data Disk Count: 12
         CPU Core Count: 32
                Created: September 28, 2025 at 9:05:22 PM UTC

System Information
----------------------------------------------------------------
                   Name: test1te
            Domain Name: testpublicsu.testvcn.oraclevcn.com
              Time Zone: UTC
             DB Edition: EeHp
            DNS Servers:
            NTP Servers: 100.200.100.200

Disk Group Information
----------------------------------------------------------------
DG Name                  Redundancy               Percentage
------------------------ ------------------------ ------------
Data                     External                 100
Reco                     External                 100

DcsCli Details
----------------------------------------------------------------
                Version:
            BuildNumber: 2.0.16
              GitNumber:
              BuildTime:

DcsAgent Details
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
------------------------------------------------------------
                Version: 25.4.1.0.0-SNAPSHOT
            BuildNumber: 4
              GitNumber: f2f63fd7ra9bf53a
              BuildTime:

DcsAdmin Details
------------------------------------------------------------
                Version: 25.4.1.0.0-SNAPSHOT
            BuildNumber: 25.4.1.0.0-SNAPSHOT_4
              GitNumber: f2f6edac4a9bf53a
              BuildTime:
```

Note

- On a two-node RAC DB system, the command returns information specific to the local node.
- The `dbcli` command displays the version and other details for key DB system services required for proper operation, such as `DcsCli`, `DcsAgent`, and `DcsAdmin`.

Troubleshoot Backup Failures

Database backups can fail for various reasons. Typically, a backup fails because either the database host cannot access the object store, or there are problems on the host or with the database configuration.

This article includes information to help you determine the cause of the failure and fix the problem. The information is organized into several sections, based on the error condition.

If you already know the cause, you can skip to the topic with the suggested solution. Otherwise, use the *Identify the Cause of Failure* topic to get started.

The following topics are covered in this article:

- #unique_544
- #unique_545
- #unique_546
- #unique_547
- #unique_548
- #unique_549
- TDE Wallet Issues
- Other Causes of Backup Failures
- #unique_552

Tip

You can also create serial console connections to troubleshoot your system in single-user mode. For information on creating a serial console connection in the OCI Console, see *Manage Serial Console Connection to the DB System*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Identify the Cause of Failure

You could perform a validation check to determine if a DB system is properly configured to access the OCI Services Network. The validation check for Identity and Access Management connectivity and Object Storage Service connectivity is described below.

The following topics are covered:

- Validation Check for Identity and Access Management Connectivity
- Validation Check for Object Storage Service (OSS) Connectivity

Validation Check for Identity and Access Management Connectivity

1. SSH to the DB system.
2. Log in as `opc` user.
3. Run the following command:

   ```
   curl https://identity.<region>.oci.oraclecloud.com
   ```

   where, the `<region>` corresponds to the OCI region the DB System is deployed on.

   For more information on the OCI regions, see Regions and Availability Domains.

   For example, if the DB System is deployed in the Ashburn region you need to use `us-ashburn-1` for `<region>`, the `curl` command will be:

   ```
   curl https://identity.us-ashburn-1.oci.oraclecloud.com
   ```

4. If your Virtual Cloud Network (VCN) is properly configured for accessing the OCI Services Network, you will get an immediate response as given below.

   ```
   {
       "code" : "NotAuthorizedOrNotFound",
       "message" : "Authorization failed or requested resource not found."
   }
   ```

5. The SSH session will hang and will eventually timeout if your network is **not** configured for accessing the OCI Services.
6. Depending on your VCN setup, you will need to follow the steps outlined in the action section below to configure access to the OCI Services Network.

Validation Check for Object Storage Service (OSS) Connectivity

1. SSH to the DB system.
2. Log in as `opc` user.
3. Run the following command:

   ```
   curl https://objectstorage.<region>.oraclecloud.com
   ```

   where, the `<region>` corresponds to the OCI region the DB System is deployed on.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For example, if the DB System is deployed in the Ashburn region you need to use `us-ashburn-1` for `<region>`, the `curl` command will be:

```
curl https://objectstorage.us-ashburn-1.oraclecloud.com
```

4. If your VCNis properly configured for accessing the OCI Services Network, you will get an immediate response as given below.

```
{
    "code" : "NotAuthorizedOrNotFound",
    "message" : "Authorization failed or requested resource not found."
}
```

5. The SSH session will hang and will eventually timeout if your network is **not** configured for accessing the OCI Services.
6. Depending on your VCN setup, you will need to follow the steps outlined in the action section below to configure access to the OCI Services Network.

Database Service Agent Issues

Your database makes use of an agent framework to allow you to manage it through the Oracle Cloud platform.

The following topics are covered:

- Resolve Update Failures Caused by a Stopped Agent
- Restart the Database Service Agent
- Resolve Update Failures Caused by an Agent That Needs to Be Updated
- Contact Oracle Support to Update the OCI Database Service Agent

Resolve Update Failures Caused by a Stopped Agent

Occasionally you might need to restart the `dcsagent` program if it has the status of **stop/waiting** to resolve a update failure.

Restart the Database Service Agent

1. From a command prompt, check the status of the agent:

```
initctl status initdcsagent
```

2. If the agent is in the **stop/waiting** state, try to restart the agent:

```
initctl start initdcsagent
```

3. Check the status of the agent again to confirm that it has the **start/running** status:

```
initctl status initdcsagent
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Resolve Update Failures Caused by an Agent That Needs to Be Updated

Update operation can also fail if your agent needs to be updated. The system gives the following error message for this failure:

```
Current DcsAgent version is less than or equal to minimum required version.
```

To resolve this issue, perform the steps in the following section.

Contact Oracle Support to Update the OCI Database Service Agent

1. Confirm that the agent (dcsagent) and DCS Admin program (dcsadmin) are running using the following commands:

   ```
   initctl status initdcsagent
   ```

   ```
   initctl status initdcsadmin
   ```

2. If these programs are not running, use the following commands to restart them:

   ```
   initctl start initdcsagent
   ```

   ```
   initctl start initdcsadmin
   ```

3. Follow the instructions in #unique_557 to collect your DCS agent log files.
4. Contact Oracle Support for assistance with updating the agent.

Oracle Clusterware Issues

The following topics are covered:

- The Oracle Clusterware Is Not Running
- Restart the Oracle Clusterware
- The Oracle Grid Infrastructure (GI) is Not Updated

The Oracle Clusterware Is Not Running

Oracle Clusterware enables servers to communicate with each other so that they can function as a collective unit. The cluster software program must be up and running on the DB system for update operations to complete. Occasionally you might need to restart the Oracle Clusterware to resolve a update failure.

Restart the Oracle Clusterware

1. From command prompt, check the status of Oracle Clusterware:

   ```
   crsctl check crs
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Output:

```
CRS-4638: Oracle High Availability Services is online
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronization Services is online
CRS-4533: Event Manager is online
```

For more detailed status information, you can run `crsctl stat res -t`.

2. If Oracle Clusterware is not online, try to restart the program:

```
crsctl start crs
```

3. Check the status of Oracle Clusterware to confirm that it is online:

```
crsctl check crs
```

The Oracle Grid Infrastructure (GI) is Not Updated

This problem occurs when you try to update a database before you update the DB system of that database. The error description indicates that the Oracle Grid Infrastructure must be updated first. To resolve this issue, update the DB system to latest available version. After you update the DB system, you can retry the database update operation.

To get the current and latest-available GI versions for the DB system, use the following command:

```
dbcli describe-component
```

Object Store Connectivity Issues

The DB system and database updates are stored in OCI Object Storage. Therefore, successful update operations require connectivity between the DB system host and the Object Storage location from which the updates are downloaded.

The following topics are covered:

- Ensure Your Database Host Can Connect to OCI Object Storage

Ensure Your Database Host Can Connect to OCI Object Storage

1. Use the following command to verify the host can access OCI Object Storage:

```
dbcli describe-latestpatch
```

Example output indicating success:

```
componentType   availableVersion
--------------  --------------
gi              12.2.0.1.180417
gi              12.1.0.2.180417
db              11.2.0.4.180417
db              12.2.0.1.180417
db              12.1.0.2.180417
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
oak             12.1.2.11.3
oak             12.2.1.1.0
```

Example output indicating failure:

```
DCS-10032:Resource patch metadata is not found.Failed to download
patchmetadata from objectstore
```

2. If you cannot connect to the Object Store, refer to *Back Up a Database Using the Console* for how to configure Object Store connectivity.

Host Issues

One or more of the following conditions on the database host can cause update operations to fail:

The following topics are covered:

- Database Node Not Running During the Update Operation
- The File System Is Full

Database Node Not Running During the Update Operation

All nodes of the database must be active and running while an update operation is in progress, whether you are updating the DB system or the database home. Use the OCI Console to check that the status of each node is AVAILABLE, and start the node, if needed.

The File System Is Full

Update operations require a minimum of 15 GB of free space in the `/u01` directory on the host file system. Use the `df -h` command on the host to check the available space. If the file system has insufficient space, you can remove old log or trace files to free up space.

Database Issues

An improper database state can lead to update failures.

The following topics are covered:

- Database Not Running During the Update Operation
- Check That the Database Is Active and Running
- Run the `datapatch` Task

Database Not Running During the Update Operation

The database must be active and running for all of the update tasks to complete. Otherwise, you must run the datapatch task manually.

Check That the Database Is Active and Running

Use the following command to check the state of your database, and ensure that any problems that might have put the database in an improper state are resolved:

```
srvctl status database -d <db_unique_name> -verbose
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

The system returns a message including the database instance status. The instance status must be **Open** for the update operation to succeed.

If the database is not running, use the following command to start it:

```
srvctl start database -d <db_unique_name> -o open
```

If the database is mounted but does not have the **Open** status, use the following commands to access the SQL*Plus command prompt and set the status to **Open**:

```
sqlplus / as sysdba


alter database open;
```

Run the `datapatch` Task

Before you run the `datapatch` command, ensure that all pluggable databases (PDBs) are open. To open a PDB, you can use SQL*Plus to execute `ALTER PLUGGABLE DATABASE <pdb_name> OPEN READ WRITE;` against the PDB.

```
$ORACLE_HOME/OPatch/datapatch
```

The `datapatch` command should be run on each database home.

TDE Wallet Issues

The following topics are covered:

- Incorrect TDE Wallet Location Specification
- Check the TDE Wallet Location Specification
- The `ORACLE_UNQNAME` Environment Variable Was Not Set When the Database Was Started Using SQL*Plus
- Pluggable Database Was Added With an Incorrectly Configured Master Encryption Key
- Check Configuration Related to the TDE Wallet
- Missing TDE Wallet File
- Missing Auto Login Wallet File

Incorrect TDE Wallet Location Specification

For backup operations to work, the `$ORACLE_HOME/network/admin/sqlnet.ora` file must contain the ENCRYPTION_WALLET_LOCATION parameter formatted exactly as follows:

```
ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE)(METHOD_DATA=(DIRECTORY=/opt/
oracle/dcs/commonstore/wallets/tde/$ORACLE_UNQNAME)))
```

> **Note**
>
> In this wallet location entry, `$ORACLE_UNQNAME` is an environment variable and should not be replaced with an actual value.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Check the TDE Wallet Location Specification

Use the `cat` command to check the TDE wallet location specification. For example:

```
cat $ORACLE_HOME/network/admin/sqlnet.ora
```

Output:

```
ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE)(METHOD_DATA=(DIRECTORY=/opt/
oracle/dcs/commonstore/wallets/tde/$ORACLE_UNQNAME)))
```

The `ORACLE_UNQNAME` Environment Variable Was Not Set When the Database Was Started Using SQL*Plus

If the database was started using SQL*Plus, and the `ORACLE_UNQNAME` environment variable was not set, the wallet is not opened correctly.

To fix the problem, start the database using the `srvctl` utility:

```
srvctl start database -d <db_unique_name>
```

Pluggable Database Was Added With an Incorrectly Configured Master Encryption Key

In a multitenant environment for Oracle Database versions that support PDB-level keystore, each PDB has its own master encryption key. This encryption key is stored in a single keystore used by all containers. After you create or plug in a new PDB, you must create and activate a master encryption key for it. If you do not do so, the `STATUS` column in the `v$encryption_wallet` view shows the value `OPEN_NO_MASTER_KEY`.

To check the master encryption key status and create a master key, perform the following:

1. Review the the `STATUS` column in the `v$encryption_wallet` view, as shown in the following example:

   ```
   alter session set container=pdb2;


   select WRL_TYPE,WRL_PARAMETER,STATUS,WALLET_TYPE from v$encryption_wallet;
   ```

   Output:

   ```
   WRL_TYPE  WRL_PARAMETER
   STATUS             WALLET_TYPE
   --------- ------------------------------------------------------------
   ------------------ -----------
   FILE      /opt/oracle/dcs/commonstore/wallets/tde/example_iadxyz/
   OPEN_NO_MASTER_KEY AUTOLOGIN
   ```

2. Confirm that the PDB is in READ WRITE open mode and is not restricted, as shown in the following example:

   ```
   show pdbs
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Output:

```
CON_ID CON_NAME   OPEN MODE   RESTRICTED
------ ---------- ----------- -----------
2      PDB$SEED   READ ONLY   NO
3      PDB1       READ WRITE  NO
4      PDB2       READ WRITE  NO
```

The PDB cannot be open in restricted mode (the `RESTRICTED` column must show `NO`). If the PDB is currently in restricted mode, review the information in the PDB_PLUG_IN_VIOLATIONS view and resolve the issue before continuing. For more information on the PDB_PLUG_IN_VIOLATIONS view and the restricted status, review the documentation on pluggable database for your Oracle Database version.

3. Run the following `DBCLI` commands to change the status to `OPEN`:

```
sudo su –
```

```
dbcli list-database
```

```
dbcli update-tdekey -i <database_ID> -n <PDB_name> -p
```

The `update-tdekey` command shown will prompt you for the admin password.

4. Confirm that the status of the wallet has changed from OPEN_NO_MASTER_KEY to OPEN by querying the `v$encryption_wallet` view as shown in step 1.

Check Configuration Related to the TDE Wallet

Several configuration parameters related to the TDE wallet can cause backups to fail.

1. Check that the environment's database unique name parameter (ORACLE_UNQNAME) is set correctly using the following command:

```
srvctl getenv database -d <db_unique_name>
```

For example:

```
srvctl getenv database -d orclbkp_iadxyz
```

Output:

```
orclbkp_iadxyz:
ORACLE_UNQNAME=orclbkp_iadxyz
TZ=UTC
```

2. Check your `sqlnet.ora` settings to confirm that the file has an ENCRYPTION_WALLET_LOCATION parameter with the correct DIRECTORY value. Use the following command:

```
cat $ORACLE_HOME/network/admin/sqlnet.ora
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

For example:

```
cat $ORACLE_HOME/network/admin/sqlnet.ora
```

Output:

```
ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE)
(METHOD_DATA=(DIRECTORY=/opt/oracle/dcs/commonstore/
wallets/tde/$ORACLE_UNQNAME)))
```

3. Confirm that the wallet status is **open** and the wallet type is **auto login** by checking the `v$encryption_wallet` view. For example:

```
select status, wrl_parameter,wallet_type from v$encryption_wallet;
```

Output:

```
STATUS  WRL_PARAMETER
WALLET_TYPE
------- ------------------------------------------------------------
------------
OPEN    /opt/oracle/dcs/commonstore/wallets/tde/example_iadxyz/  AUTOLOGIN
```

For pluggable databases (PDBs), be sure that you switch to the appropriate container before querying `v$encryption_wallet` view. For example:

```
sqlplus / as sysdba


alter session set container=pdb1;


select WRL_TYPE,WRL_PARAMETER,STATUS,WALLET_TYPE from v$encryption_wallet;
```

Output:

```
WRL_TYPE  WRL_PARAMETER                                              STATUS
WALLET_TYPE
--------- ------------------------------------------------------- -------
------------
FILE      /opt/oracle/dcs/commonstore/wallets/tde/tiger_iad3c8/  OPEN
AUTOLOGIN
```

Missing TDE Wallet File

The TDE wallet file (`ewallet.p12`) can cause backups to fail if it is missing, or if it has incompatible file system permissions or ownership. Check the file as shown in the following example:

```
ls -ltr /opt/oracle/dcs/commonstore/wallets/tde/$ORACLE_UNQNAME/ewallet.p12
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Output:

```
-rwx------ 1 oracle oinstall 5680 Apr 18 13:09 /opt/oracle/dcs/commonstore/
wallets/tde/orclbkp_iadxzy/ewallet.p12
```

The TDE wallet file should have file permissions with the octal value "700" (`-rwx------`), and the owner of this file should be a part of the `oinstall` operating system group.

Missing Auto Login Wallet File

The auto login wallet file (`cwallet.sso`) can cause backups to fail if it is missing, or if it has incompatible file system permissions or ownership. Check the file as shown in the following example:

```
ls -ltr /opt/oracle/dcs/commonstore/wallets/tde/$ORACLE_UNQNAME/cwallet.sso
```

Output:

```
-rwx------ 1 oracle oinstall 5725 Apr 18 13:09 /opt/oracle/dcs/commonstore/
wallets/tde/orclbkp_iadxyz/cwallet.sso
```

The auto login wallet file should have file permissions with the octal value "700" (`-rwx------`), and the owner of this file should be a part of the `oinstall` operating system group.

Other Causes of Backup Failures

The following topics are covered:

- Unmounted Commonstore Mount Point
- Check the Commonstore Mount Point
- Confirm That `ora.data.commonstore.acfs` Is Online
- The Database Is Not Properly Registered

Unmounted Commonstore Mount Point

The mount point `/opt/oracle/dcs/commonstore` must be mounted, or backups will fail.

Check the Commonstore Mount Point

Confirm that the mount point `/opt/oracle/dcs/commonstore` is mounted, as shown in the following example:

```
srvctl config filesystem -volume commonstore -diskgroup data
```

Output:

```
Volume device: /dev/asm/commonstore-5
Diskgroup name: data
Volume name: commonstore
Canonical volume device: /dev/asm/commonstore-5
Accelerator volume devices:
Mountpoint path: /opt/oracle/dcs/commonstore
Mount point owner: oracle
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
Mount users:
Type: ACFS
```

Confirm That ora.data.commonstore.acfs Is Online

1. The state for ora.data.commonstore.acfs must be online, or backups will fail. Confirm as shown in the following example:

   ```
   crsctl stat resource ora.data.commonstore.acfs -v
   ```

   Output:

   ```
   NAME=ora.data.commonstore.acfs
   TYPE=ora.acfs.type
   LAST_SERVER=orcl
   STATE=OFFLINE
   TARGET=OFFLINE
   ...
   STATE_DETAILS=admin unmounted /opt/oracle/dcs/commonstore
   ...
   ```

2. List the contents of the commonstore directory to confirm that it is mounted

   ```
   ls -ltr /opt/oracle/dcs/commonstore
   ```

3. If the STATE_DETAILS value is unmounted, mount the file system as shown in the following example:

   ```
   srvctl start filesystem -volume commonstore -diskgroup data
   ```

4. Confirm that the change was successful as shown in the following example:

   ```
   crsctl stat resource ora.data.commonstore.acfs -v
   ```

   Output:

   ```
   NAME=ora.data.commonstore.acfs
   TYPE=ora.acfs.type
   LAST_SERVER=orcl
   STATE=ONLINE on orcl
   TARGET=ONLINE
   CARDINALITY_ID=ONLINE
   ...
   STATE_DETAILS=mounted on /opt/oracle/dcs/commonstore
   ```

5. List the contents of the commonstore directory to confirm that it is mounted, as shown in the following example:

   ```
   ls -ltr /opt/oracle/dcs/commonstore
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Output:

```
total 220
drwx------ 2 root   root     65536 Apr 18 10:50 lost+found
drwx------ 3 oracle oinstall 20480 Apr 18 11:02 wallets
drwxr-xr-x 3 root   root     20480 Apr 20 06:41 pkgrepos
drwxr-xr-x 4 oracle oinstall 20480 Apr 20 06:41 objectstore
```

The Database Is Not Properly Registered

Database backups fail if the database is not registered with the `dcs-agent`. This scenario can occur if you manually migrate the database to OCI and do not run the `dbcli register-database` command.

To check whether the database is properly registered, review the information returned by running the `srvctl config database` command and the `dbcli list-databases` command. If either command does not return a record of the database, contact Oracle Support Services.

For instructions on how to register the database, refer to the following topics:

- Register the Database on the DB System in *Recover a Database from the OCI Classic Object Store*
- Database Commands in *Oracle Database CLI Reference*

Get Additional Help

If your failure occurs in the `post_action` stage, or if you are unable to resolve failures that occur in the `pre_action` stage, do the following:

1. Gather the information listed in the document Diagnostic Data Collection For Oracle Database Cloud Service Instance.
2. Open a Support Request with My Oracle Support. Contact Oracle Support.

Troubleshoot Update Failures

Update operations can fail for various reasons. Typically, an operation fails because a database node is down, there is insufficient space on the file system, or the database host cannot access the object store.

This article includes information to help you determine the cause of the failure and fix the problem. The information is organized into several sections, based on the error condition.

If you already know the cause, you can skip to the topic with the suggested solution. Otherwise, use the *Identify the Cause of Failure* topic to get started.

The following topics are covered in this article:

- #unique_568
- Database Service Agent Issues
- Object Store Connectivity Issues
- Host Issues
- Oracle Clusterware Issues
- Database Issues
- #unique_557

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Tip**
>
> You can also create serial console connections to troubleshoot your system in single-user mode. For information on creating a serial console connection in the OCI Console, see *Manage Serial Console Connection to the DB System*.

Identify the Cause of Failure

You could perform a validation check to determine if a DB system is properly configured to access the OCI Services Network. The validation check for Identity and Access Management connectivity and Object Storage Service connectivity is described below.

The following topics are covered:

- Validation Check for Identity and Access Management Connectivity
- Validation Check for Object Storage Service (OSS) Connectivity

Validation Check for Identity and Access Management Connectivity

1. SSH to the DB system.
2. Log in as `opc` user.
3. Run the following command:

   ```
   curl https://identity.<region>.oci.oraclecloud.com
   ```

   where, the `<region>` corresponds to the OCI region the DB System is deployed on.

   For more information on the OCI regions, see Regions and Availability Domains.

   For example, if the DB System is deployed in the Ashburn region you need to use `us-ashburn-1` for `<region>`, the `curl` command will be:

   ```
   curl https://identity.us-ashburn-1.oci.oraclecloud.com
   ```

4. If your Virtual Cloud Network (VCN) is properly configured for accessing the OCI Services Network, you will get an immediate response as given below.

   ```
   {
       "code" : "NotAuthorizedOrNotFound",
       "message" : "Authorization failed or requested resource not found."
   }
   ```

5. The SSH session will hang and will eventually timeout if your network is **not** configured for accessing the OCI Services.
6. Depending on your VCN setup, you will need to follow the steps outlined in the action section below to configure access to the OCI Services Network.

Validation Check for Object Storage Service (OSS) Connectivity

1. SSH to the DB system.
2. Log in as `opc` user.

Copyright © 2022, 2026, Oracle and/or its affiliates.

3. Run the following command:

   ```
   curl https://objectstorage.<region>.oraclecloud.com
   ```

   where, the `<region>` corresponds to the OCI region the DB System is deployed on.

   For example, if the DB System is deployed in the Ashburn region you need to use `us-ashburn-1` for `<region>`, the `curl` command will be:

   ```
   curl https://objectstorage.us-ashburn-1.oraclecloud.com
   ```

4. If your VCNis properly configured for accessing the OCI Services Network, you will get an immediate response as given below.

   ```
   {
       "code" : "NotAuthorizedOrNotFound",
       "message" : "Authorization failed or requested resource not found."
   }
   ```

5. The SSH session will hang and will eventually timeout if your network is **not** configured for accessing the OCI Services.
6. Depending on your VCN setup, you will need to follow the steps outlined in the action section below to configure access to the OCI Services Network.

Database Service Agent Issues

Your database makes use of an agent framework to allow you to manage it through the Oracle Cloud platform.

The following topics are covered:

- Resolve Update Failures Caused by a Stopped Agent
- Restart the Database Service Agent
- Resolve Update Failures Caused by an Agent That Needs to Be Updated
- Contact Oracle Support to Update the OCI Database Service Agent

Resolve Update Failures Caused by a Stopped Agent

Occasionally you might need to restart the `dcsagent` program if it has the status of **stop/waiting** to resolve a update failure.

Restart the Database Service Agent

1. From a command prompt, check the status of the agent:

   ```
   initctl status initdcsagent
   ```

2. If the agent is in the **stop/waiting** state, try to restart the agent:

   ```
   initctl start initdcsagent
   ```

3. Check the status of the agent again to confirm that it has the **start/running** status:

   ```
   initctl status initdcsagent
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Resolve Update Failures Caused by an Agent That Needs to Be Updated

Update operation can also fail if your agent needs to be updated. The system gives the following error message for this failure:

```
Current DcsAgent version is less than or equal to minimum required version.
```

To resolve this issue, perform the steps in the following section.

Contact Oracle Support to Update the OCI Database Service Agent

1. Confirm that the agent (dcsagent) and DCS Admin program (dcsadmin) are running using the following commands:

   ```
   initctl status initdcsagent
   ```

   ```
   initctl status initdcsadmin
   ```

2. If these programs are not running, use the following commands to restart them:

   ```
   initctl start initdcsagent
   ```

   ```
   initctl start initdcsadmin
   ```

3. Follow the instructions in #unique_557 to collect your DCS agent log files.
4. Contact Oracle Support for assistance with updating the agent.

Object Store Connectivity Issues

The DB system and database updates are stored in OCI Object Storage. Therefore, successful update operations require connectivity between the DB system host and the Object Storage location from which the updates are downloaded.

The following topics are covered:

- Ensure Your Database Host Can Connect to OCI Object Storage

Ensure Your Database Host Can Connect to OCI Object Storage

1. Use the following command to verify the host can access OCI Object Storage:

   ```
   dbcli describe-latestpatch
   ```

   Example output indicating success:

   ```
   componentType   availableVersion
   --------------  --------------
   gi              12.2.0.1.180417
   gi              12.1.0.2.180417
   db              11.2.0.4.180417
   db              12.2.0.1.180417
   db              12.1.0.2.180417
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
oak             12.1.2.11.3
oak             12.2.1.1.0
```

Example output indicating failure:

```
DCS-10032:Resource patch metadata is not found.Failed to download
patchmetadata from objectstore
```

2. If you cannot connect to the Object Store, refer to *Back Up a Database Using the Console* for how to configure Object Store connectivity.

Host Issues

One or more of the following conditions on the database host can cause update operations to fail:

The following topics are covered:

- Database Node Not Running During the Update Operation
- The File System Is Full

Database Node Not Running During the Update Operation

All nodes of the database must be active and running while an update operation is in progress, whether you are updating the DB system or the database home. Use the OCI Console to check that the status of each node is AVAILABLE, and start the node, if needed.

The File System Is Full

Update operations require a minimum of 15 GB of free space in the `/u01` directory on the host file system. Use the `df -h` command on the host to check the available space. If the file system has insufficient space, you can remove old log or trace files to free up space.

Oracle Clusterware Issues

The following topics are covered:

- The Oracle Clusterware Is Not Running
- Restart the Oracle Clusterware
- The Oracle Grid Infrastructure (GI) is Not Updated

The Oracle Clusterware Is Not Running

Oracle Clusterware enables servers to communicate with each other so that they can function as a collective unit. The cluster software program must be up and running on the DB system for update operations to complete. Occasionally you might need to restart the Oracle Clusterware to resolve a update failure.

Restart the Oracle Clusterware

1. From command prompt, check the status of Oracle Clusterware:

```
crsctl check crs
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Output:

```
CRS-4638: Oracle High Availability Services is online
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronization Services is online
CRS-4533: Event Manager is online
```

For more detailed status information, you can run `crsctl stat res -t`.

2. If Oracle Clusterware is not online, try to restart the program:

```
crsctl start crs
```

3. Check the status of Oracle Clusterware to confirm that it is online:

```
crsctl check crs
```

The Oracle Grid Infrastructure (GI) is Not Updated

This problem occurs when you try to update a database before you update the DB system of that database. The error description indicates that the Oracle Grid Infrastructure must be updated first. To resolve this issue, update the DB system to latest available version. After you update the DB system, you can retry the database update operation.

To get the current and latest-available GI versions for the DB system, use the following command:

```
dbcli describe-component
```

Database Issues

An improper database state can lead to update failures.

The following topics are covered:

- Database Not Running During the Update Operation
- Check That the Database Is Active and Running
- Run the `datapatch` Task

Database Not Running During the Update Operation

The database must be active and running for all of the update tasks to complete. Otherwise, you must run the datapatch task manually.

Check That the Database Is Active and Running

Use the following command to check the state of your database, and ensure that any problems that might have put the database in an improper state are resolved:

```
srvctl status database -d <db_unique_name> -verbose
```

The system returns a message including the database instance status. The instance status must be **Open** for the update operation to succeed.

Copyright © 2022, 2026, Oracle and/or its affiliates.

If the database is not running, use the following command to start it:

```
srvctl start database -d <db_unique_name> -o open
```

If the database is mounted but does not have the **Open** status, use the following commands to access the SQL*Plus command prompt and set the status to **Open**:

```
sqlplus / as sysdba
```

```
alter database open;
```

Run the `datapatch` Task

Before you run the `datapatch` command, ensure that all pluggable databases (PDBs) are open. To open a PDB, you can use SQL*Plus to execute `ALTER PLUGGABLE DATABASE <pdb_name> OPEN READ WRITE;` against the PDB.

```
$ORACLE_HOME/OPatch/datapatch
```

The `datapatch` command should be run on each database home.

Get Additional Help

If your failure occurs in the `post_action` stage, or if you are unable to resolve failures that occur in the `pre_action` stage, do the following:

1. Gather the information listed in the document Diagnostic Data Collection For Oracle Database Cloud Service Instance.
2. Open a Support Request with My Oracle Support. Contact Oracle Support.

Troubleshoot Shape Change Failures

This article helps you troubleshoot and fix the issues that might occur when you change the shape of your DB system.

For multi-node RAC DB systems, shape change operations proceed in a rolling fashion. Depending on where in the shape change operation the failure occurs, you may be able to re-try the operation using the console.

This article includes information to help you determine the cause of the failure and fix the problem. The information is organized into several sections, based on the error condition.

The following topics are covered in this article:

- Use the OCI Console to Troubleshoot
- Use `dbcli` to Troubleshoot
- Get Additional Help

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> You can also create serial console connections to troubleshoot your DB system in single-user mode. For information on creating a serial console connection in the Console, see *Manage Serial Console Connection to the DB System*.

Use the OCI Console to Troubleshoot

If your shape change operation fails, a message banner appears on the DB system details page to provide details about the failure. If the failure happens on the first node of a multi-node system, and the operation is rolled back successfully, the **Change shape** button remains available and the system remains online, in the available state. Contact Oracle Support to get additional details about the failure. You can also use the dbcli to learn more about the failure and what issues need to be resolved. After determining that no issues remain, you can try the operation again.

If the failure leaves the system in a state where the operation cannot be rolled back, the system state is *Needs Attention*. In this case, contact Oracle Support as soon as you are aware of the issue so Oracle can help you resolve the issue and complete the shape change operation.

Use `dbcli` to Troubleshoot

The following topics are covered:

- Determine What Stage of the Shape Change Operation Failed
- Troubleshooting Failures That Occur in the Pre_action Stage

Determine What Stage of the Shape Change Operation Failed

1. Login to the DB system as the root user.
2. Navigate to `/opt/oracle/dcs/bin`:

   ```
   cd /opt/oracle/dcs/bin
   ```

3. Update the CLI tool:

   ```
   ./cliadm update-dbcli
   ```

4. List the failed jobs:

   ```
   dbcli list-jobs | grep -i failed
   ```

> **Note**
>
> If the failed job (or jobs) occurred during the `pre_action`, `action`, or `post_action` stage. Also note the `job_id` value of the failed job, which you will need to resolve the issue.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Troubleshooting Failures That Occur in the Pre_action Stage

1. Use the `job_id` value to get more information about the failure:

   ```
   dbcli describe-job -i <job_id>
   ```

2. Search for the error in the `dcs-agent.log` and `dcs-agent-debug.log` files (which are located in the `/opt/oracle/dcs/log/` directory):

   For example:

   ```
   cd /opt/oracle/dcs/log
   grep -ir "DCS-10063:Failed to get node names from olsnodes." *
   ```

3. Using the information about the error recorded in the log file, correct the system configuration if possible.
4. Re-try the shape change operation. If the operation is still not successful, follow the instructions in Get Additional Help.

Get Additional Help

If your failure occurs in the `post_action` stage, or if you are unable to resolve failures that occur in the `pre_action` stage, do the following:

1. Gather the information listed in the document Diagnostic Data Collection For Oracle Database Cloud Service Instance.
2. Open a Support Request with My Oracle Support. Contact Oracle Support.

Troubleshoot Network Connectivity Failures

Network connectivity can fail for various reasons. Typically, the network connectivity fails because of access or authorization issues.

This article includes information to help you to troubleshoot network connectivity issues between the guest VM and OCI Services Network. The information is organized into several sections, based on the error condition.

If you already know the cause, you can skip to the topic with the suggested solution. Otherwise, use the *Identify the Cause of Failure* topic to get started.

The following topics are covered in this article:

- Identify the Cause of Failure
- Resolve Network Connectivity Failures
- #unique_578

> **Tip**
>
> You can also create serial console connections to troubleshoot your system in single-user mode. For information on creating a serial console connection in the OCI Console, see *Manage Serial Console Connection to the DB System*.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Identify the Cause of Failure

You could perform a validation check to determine if a DB system is properly configured to access the OCI Services Network. The validation check for Identity and Access Management connectivity and Object Storage Service connectivity is described below.

The following topics are covered:

- Validation Check for Identity and Access Management Connectivity
- Validation Check for Object Storage Service (OSS) Connectivity

Validation Check for Identity and Access Management Connectivity

1. SSH to the DB system.
2. Log in as `opc` user.
3. Run the following command:

   ```
   curl https://identity.<region>.oci.oraclecloud.com
   ```

   where, the `<region>` corresponds to the OCI region the DB System is deployed on.

   For more information on the OCI regions, see Regions and Availability Domains.

   For example, if the DB System is deployed in the Ashburn region you need to use `us-ashburn-1` for `<region>`, the `curl` command will be:

   ```
   curl https://identity.us-ashburn-1.oci.oraclecloud.com
   ```

4. If your Virtual Cloud Network (VCN) is properly configured for accessing the OCI Services Network, you will get an immediate response as given below.

   ```
   {
       "code" : "NotAuthorizedOrNotFound",
       "message" : "Authorization failed or requested resource not found."
   }
   ```

5. The SSH session will hang and will eventually timeout if your network is **not** configured for accessing the OCI Services.
6. Depending on your VCN setup, you will need to follow the steps outlined in the action section below to configure access to the OCI Services Network.

Validation Check for Object Storage Service (OSS) Connectivity

1. SSH to the DB system.
2. Log in as `opc` user.
3. Run the following command:

   ```
   curl https://objectstorage.<region>.oraclecloud.com
   ```

   where, the `<region>` corresponds to the OCI region the DB System is deployed on.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For example, if the DB System is deployed in the Ashburn region you need to use `us-ashburn-1` for `<region>`, the `curl` command will be:

```
curl https://objectstorage.us-ashburn-1.oraclecloud.com
```

4. If your VCNis properly configured for accessing the OCI Services Network, you will get an immediate response as given below.

```
{
    "code" : "NotAuthorizedOrNotFound",
    "message" : "Authorization failed or requested resource not found."
}
```

5. The SSH session will hang and will eventually timeout if your network is **not** configured for accessing the OCI Services.
6. Depending on your VCN setup, you will need to follow the steps outlined in the action section below to configure access to the OCI Services Network.

Resolve Network Connectivity Failures

Perform the following instructions to resolve network connectivity failures.

The following topics are covered:

- DB Systems Deployed On a Private Subnet
- DB Systems Deployed On a Public Subnet

DB Systems Deployed On a Private Subnet

Perform the following step if you have deployed your DB system on a private subnet.

1. Configure a Service Gateway for use by the DB System to reach the OCI Service Network. For detailed steps, see *VCN and Subnets*.

After you configure your VCN to reach the OCI Services Network, perform the validation check to ensure that you have established connectivity to the OCI Services Network from your DB System.

DB Systems Deployed On a Public Subnet

Perform the following step if you have deployed your DB system on a public subnet.

1. Configure a Internet Gateway for use by the DB System to reach the OCI Service Network. For detailed steps, see *VCN and Subnets*

After you configure your VCN to reach the OCI Services Network, perform the validation check to ensure that you have established connectivity to the OCI Services Network from your DB System.

For more information on managing a service gateway, see Managing a Service Gateway in the Console.

Get Additional Help

If your failure occurs in the `post_action` stage, or if you are unable to resolve failures that occur in the `pre_action` stage, do the following:

Copyright © 2022, 2026, Oracle and/or its affiliates.

1. Gather the information listed in the document Diagnostic Data Collection For Oracle Database Cloud Service Instance.
2. Open a Support Request with My Oracle Support. Contact Oracle Support.

Copyright © 2022, 2026, Oracle and/or its affiliates.